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

March 20, 2007

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). With regard to scientific goals, the theory has been extremely successful, explaining a broad range of creative behaviour and stimulating debate in diverse disciplines. For practical purposes, however, it has been less productive. Used as a means of classifying behaviour, Boden's definition of creativity is analytically non-deterministic. This has limited the utility of the theory for the design of systems which express or support creativity. Attempts have been made to progress the theory so as to eliminate the non-determinism but these have generally involved refining the underlying definition. In contrast, the present paper takes steps to *reconstruct* it from first principles, using concepts of conceptual complexity. In this re-working, Boden's definition is redeveloped as a qualitative abstraction relating to a formally defined complexity tradeoff. With its ambiguity eliminated, the definition has the capacity to underwrite fully deterministic, process classifications and to offer clearer theoretical guidance for systems-building work.

Keywords: computational creativity, generative creativity, complexity, representation, theoretical cognitive science.

1 Introduction to Boden's theory

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 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 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 founded on conceptualisation or the realisation of a point within a particular 'conceptual space'.[1] 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 rules as are other, familiar ideas. A genuinely original, or creative, idea is one which cannot.' [2] (Boden, 1990, p. 40) A new concept, then, is

[1] This is a common interpretation of Boden's approach. However Ritchie (2001, 2006) generalises Boden's conceptual space to be a space of concrete or abstract *artifacts*. It is the ability of Boden's theory to provide an explanation of the creative content of artifacts which is significant for Ritchie and he identifies a range of factors which are significant for this issue.

[2] For reasons that will be explained, it can be assumed that mental representation of a

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 a clear characterisation of the process of creativity. However, consideration of further evidence compels her 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 ... familiar conceptual spaces.' (ibid., p. 108) However, it is not Boden's intention to broaden the definition to include *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 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 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,

---

conceptual space must be generative in character.

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.

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 Ritchie notes, the distinction between *combinational* and *exploratory* is 'hard to pin down' (Ritchie, 2001, p. 7). 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.

Insofar as there is a definition underpinning the theory, then, it has to be the 'strong' definition. But this raises a difficulty. Transformational operations can always be regarded as exploratory operations carried out in a transformation space. Whether we classify a particular thought process (or computer system) as creative or not would then seem to depend on the viewpoint we take; more specifically, it must depend on how we choose to label internal data structures. Boden's definition, then, appears to be intrinsically *subjective*.

Concerns with the ambiguity of the definition were raised by many of the original commentators, including 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 noted 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). Others noted that transformation is a form of meta-level search.[3] Schank and Forster describe application of the definition as a 'mare's nest' (Schank and Forster, 1995, p. 138) while Turner (1995, p. 152) notes that although Boden wishes 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 response to this criticism, Boden accepted that the definition was insufficiently clear. 'Most of the reviewers point out, quite rightly,' she commented, 'that my definition of creativity ... was vague.' (Boden, 1994b, p. 559) Her defence was that the necessities of writing for a general audience dictated

[3]This has subsequently been demonstrated formally by Wiggins (2006b).

the limiting of technical content. In later writings, she conceded the existence of an inconsistency but suggested that it resulted 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' (Boden, 1995, p. 163)

Unfortunately, neither of Boden's defences really address the difficulty. The problem with the definition is not that it is *vague*. On the contrary, it is clearly stated, makes use of well-established computational concepts and has generally been readily understood. Indeed the indications are that the definition has been well understood even in fields where computational knowledge must be the exception rather than the rule.

Neither is the problem that the definition has been mis-applied, although it is possible that this may have occurred. The problem with the definition is that it is analytically non-deterministic. Used as a means of classification, it produces subjective and potentially contradictory results. According to Boden's rule, any process which is deemed to be creative on the grounds that it entailed the construction of a new conceptual space, may also be regarded as *un*creative, on the grounds that the construction of that space involved mere search.

In some cases, the non-determinism can be safely ignored. Even though no theoretical discrimination can be made between exploration and transformation, the position can be taken that the latter is different in some important way that cannot be formally defined. On this basis, discursive analysis and explanation can proceed and, as Boden has demonstrated, in a wide-ranging and fruitful way. But the definitional ambiguity becomes more of a difficulty in any practical context where the determination of the degree of creativity, or its likely origin, is a key issue.

Computer-based creativity practitioners want to know how to better build systems that express or support creativity. But the guidance Boden's theory offers is hampered by the definitional ambiguity and restricted to a fairly small domain of application. Defining creativity in terms of transformational/exploratory operations carried out in conceptual spaces, Boden's theory explicitly invokes notions of search. For systems-builders whose work revolves around such ideas, there is the identification of interesting parallels between creative thought and search-based functionality. For practitioners whose work is located elsewhere, the theory is less informative.

Even for systems-builders operating within a specifically search-oriented paradigm, the framework offers a relatively vague message. As Wiggins has shown (Wiggins, 2006a), Boden's definition is essentially the proposition that genuine creativity involves search at a conceptual meta level. On this basis, systems which exploit some form of meta-reasoning may be deemed genuinely creative. But exactly how such creativity should be enhanced or developed is not addressed.

There is also the concern that Boden's definition is too inclusive with regard to computational function, seeming to place any mechanism which embodies transformational functionality on an equal footing. As an illustration of the implications, consider the case of the digital watch which responds to the first day of January in the year 2000 by introducing a new number representation system in order to accommodate a greater range of dates. Treating the representation of complete days as concepts, the application of Boden's definition to this example would result in the conclusion that the watch is genuinely — even *radically* — creative. Even for systems-builders committed to exploitation of search, the degree of lassitude here is awkward.

Even so, it is certainly this group which has been best placed to make use of Boden's theory and its members have developed a range of strategies for dealing with the problems. Ritchie, for example, deals with the problem of over-inclusivity by suggesting that Boden must have intended that transformational functionality should be considered merely a *necessary* condition for creativity, rather than necessary and sufficient (Ritchie, 2006).

Wiggins has proposed a formalisation of Boden's theory which represents its key concepts in terms of set-theoretic constructs. A useful feature of this contribution is the way it distinguishes different forms of transformation. Under Wiggins' approach, it becomes clear that we need to discriminate transformation of the concept rules (i.e., the set of rules defining the conceptual space) from transformation of the transformation rules. This is an important point because, as Wiggins notes, it seems likely that transformation of transformation rules may represent a particularly important aspect of genuine creativity, since it is a potentially recursive operation. Commenting that "[this] kind of transformation ... more naturally applies to the creative individuals' *modus operandi*' Wiggins notes that this form is not discriminated within Boden's framework. (Wiggins, 2006a)

For present purposes, however, Wiggins formalisation consolidates rather than resolves difficulties. The suspected equivalence of exploration and transformation is confirmed mathematically and the lack of detail about the implementation of transformation is highlighted. In a way, Wiggins extends the problem by identifying a second category of operations that stands in need of clarification.

Directed towards the refinement and elaboration of Boden's framework, the work of Wiggins, Ritchie and others retains her definition as a theoretical foundation. But the position of the present paper is that, being analytically non-deterministic, the definition is in need of reconstruction. In its existing form, the definition cannot provide the basis for formal classification and therefore cannot provide clear guidance for practical work. The aim of the present paper is to reconstruct the definition in a way that retains its character while better serving the needs of classification and systems-building work.

The approach will build on Bundy's suggestion (Bundy, 1994) that Boden's definition could be reconstructed in terms of some notion of complexity. A possible avenue for this approach would be to make use of notions of algorithmic complexity. (Kolmogorov, 1965, 1968; Chaitin, 1966; Chaitin, 1987) This line of attack might proceed along the lines that Schmidhuber has explored (Schmidhuber, 1997), i.e., accounting for creativity in terms of the Kolmogorow complexity of created artifacts. But this approach has some pitfalls in the present context. The view has to be taken that the creative agent is a form of computer program while estimates of creativity can only be made on an artifact-by-artifact basis. This is incompatible with the emphasis on conceptual *spaces* that is central to Boden's framework. Finally, there is the problem that algorithmic complexity cannot generally be measured in practice.

The present paper explores an alternative avenue, seeking to measure the complexity of a creative process in terms of the number of concepts and conceptual spaces that can be generated. On the basis of there being a duality in concept-construction (separating composition and categorisation) formulae are introduced which identify the limits of conceptual growth. Taking into account that conceptual spaces can only exist within an over-arching conceptual universe (cf. Wiggins, 2006a), the value of a concept is then estimated in terms of its coverage of sub-concepts. This leads to the introduction of a *net conceptual complexity* measure and the demonstration that Boden's transformation/exploration distinction can be seen as abstracting a significant cut-point in a complexity tradeoff. Reconstructed in this form, the definition offers clearer guidance for practical work.

The plan for the remainder of the paper will be as follows. The next section (Section 2) will analyse the structure of generative representation. The notion of concept duality introduced in Section 3 will then be used (Section 4) for purposes of establishing a formal measure of maximal concept growth. Bringing to bear Wiggins' observation the conceptual spaces must logically exist within conceptual *universes*, Section 5 will then introduce the key concepts of gross and net conceptual complexity. Section 6 completes the groundwork for the reconstruction by showing that 'sponge zones' (defined as sub-spaces exhibiting reduced net-conceptual-complexity) form an analogue of Boden's conceptual spaces. It is then the task of Section 7 to articulate the details of the reconstruction and justify the re-development of Boden's definition of transformation as a complexity-based abstraction. Section 8 evaluates the reconstruction in terms of the way it is likely to be viewed by commentators on the theory and by Boden herself. It also examines the degree to which the proposed fleshing-out of the transformation/exploration distinction is in agreement with the systems evidence Boden presents and reviews some of the reconstruction's explanatory possibilities. Finally, Section 9 sets out a summary.

2 Generatively represented conceptual spaces

The key notion in Boden's theory is that of the *conceptual space*. While no formal definition has been provided, it is common to interpret the phrase literally, taking the conceptual space to be a space of *conceptualisations* or concept representations.[4] In support of this, we have Boden's observation that a computational search space is 'one example ... of a conceptual space" (Boden, 1992, p. 77)

[4]More generally, Ritchie (2006) takes the conceptual space to be a space of abstract or concrete artifacts.

As for the representation of a conceptual space, we know that it must be *generative* in nature. Any form of explicit, point-by-point representation 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:

(1) To each number between 0 and 9 apply the next 2 instructions

(2) To each number between 0 and 9 apply the next instruction

(3) 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 to specify a generative system. 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 simple grammar:

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

Det -> the | a | one
```

```
N -> man | woman
V -> ran | walked | sang | laughed
```

The first two lines here 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 different ways of constructing an 'S' object. These are '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.

3 Concept duality

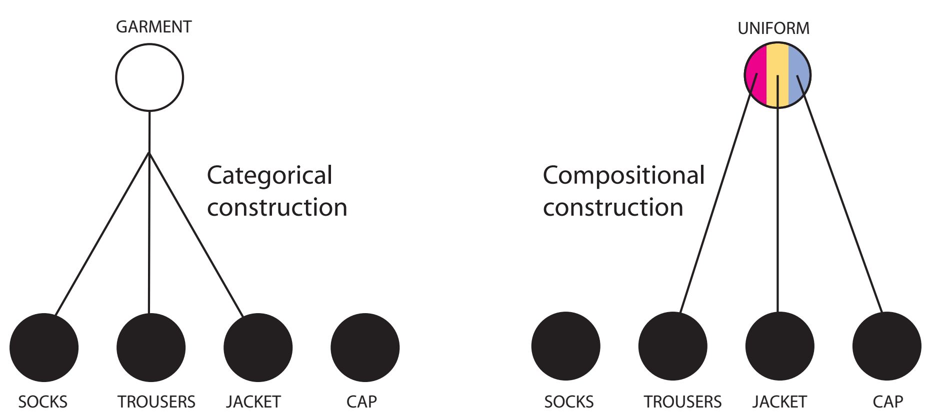

Figure 1: Concept-construction methods.

A generative representation of a conceptual space must consist, 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 formed *compositionally*, in which case the existing concepts become constituents of a new whole; or it may be formed *categorically*, in which case they are brought together to form a class. There is thus a *duality* in concept formation, in the sense that it may entail composition or categorisation.

Figure 1 illustrates the distinction between compositional and categorical construction. The constituent concepts here are articles of clothing: SOCKS, TROUSERS, JACKET and CAP, each one represented by a filled circle. The first three may be combined to form the categorical concept GARMENT in which each constituent forms an equal instance. This concept is represented by an empty circle in the left side of the figure. The second three may be combined to form the compositional concept UNIFORM (as in 'MILITARY UNIFORM') in which each constituent forms a component of the whole. This concept is represented by the filled circle in the right half of the figure. Note the way in which the arcs connect to the internal structure of UNIFORM, reflecting compositionality. In contrast the arcs connecting to GARMENT combine, reflecting the fact that the concept is a category (i.e., class) in which the constituents are alternatives.

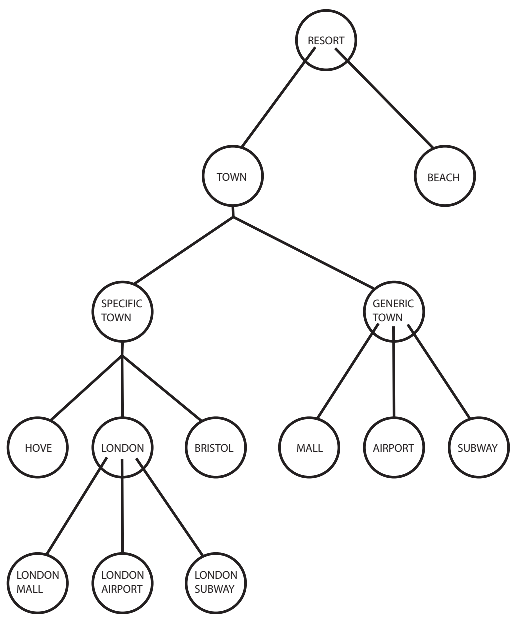

Figure 2: Schematic concept hierarchy for the RESORT concept.

Of course, constructions may be combined in arbitrary ways, producing complex hierarchical structures. Figure 2 presents a toy example, being a sample concept hierarchy for the concept RESORT. Lack of space precludes incorporation of realistically-sized constructions. But the alternation between categorical and compositional construction is not implausible. Note that categorical construction is again discriminated here using arcs which combine, differentiating it from compositional construction in which the arcs remain separate.

In the language of logic, categorisation is simply the arrangement of the constituents as a *disjunction.* Composition, on the other hand, is the arrangement of the constituents in a relationship (i.e., a function) such that each one plays the role of a part or component. In a simple case this might mean the formation of a straightforward *conjunction* (i.e., introduction of the AND relationship). But the number of compositional possibilities depends on how many relationships can be applied. With more relationships available, there are more ways to form compositional concepts.

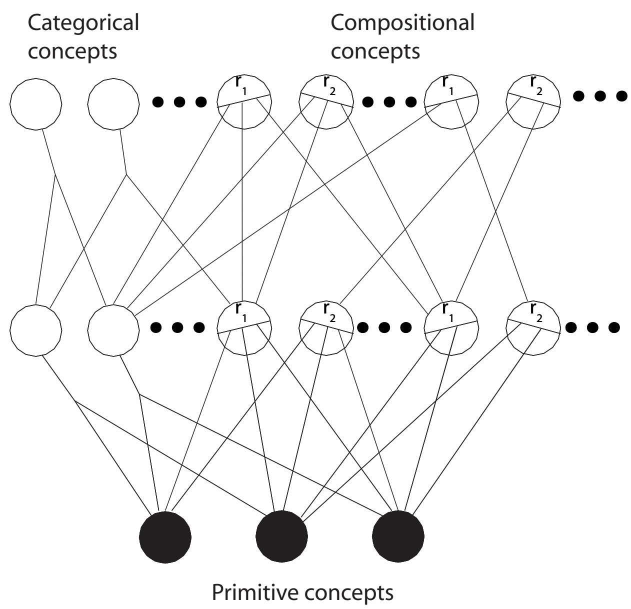

Figure 3: Possible avenues of concept development.

Even with a very small number of relationships, the number of potential concepts grows rapidly under hierarchical development. Figure 3 illustrates the effect. This schematic follows the convention of discriminating categorical construction using arcs which combine. However here the compositional arcs abut a line, the slope of which represents the character of the relevant relationship. In addition to categorical construction, this conceptual expansion is making use of two forms of compositional construction, one invoking the relationship r_1 and the other the relationship r_2.

In deciding whether a particular concept is categorical or compositional, we must first identify its constituents (i.e., the relevant sub-concepts). Once this is done, classification can be accomplished by determining whether the sub-concepts are alternatives or constituents, with categorical construction being indicated in the former case and compositional construction in the latter. For example, taking sub-concepts to be representations of keys on a piano keyboard, the concept BLACK-NOTE would be deemed categorical while ARPEGGIO would be compositional. The note-representing concepts which form the constituents in BLACK-NOTE are alternatives but the note-representing concepts within ARPEGGIO are constituents. Similar example can be given in other domains. Taking sub-concepts to represent buildings, the concept SUBURBAN-SEMI would be categorical while VILLAGE would be compositional. Taking sub-concepts to represent individual playing cards, the concept SPADE would be categorical and the concept FULL-HOUSE compositional.

3.1 Shallow vs. deep conceptual space

The task in hand is to place an upper bound on the number of concepts a given agent can construct in a particular scenario. The means of doing this will be to identify the possible concept-construction methods that the agent has at its disposal and to then calculate how many different ways they can be applied. An initial expectation is that there will be a link between concept-construction methods and concepts, i.e., more construction methods should lead to the production of more concepts. However, it turns out that the key issue for conceptual growth is not the number of construction methods but whether they can be applied *hierarchically*.

Consider first the extreme case of the agent that has no concept-construction methods at all. Clearly, there can be no conceptual growth here. The agent remains conceptually undeveloped indefinitely. At the other extreme, there is the agent that is fully equipped with construction methods, including some which can be applied to the non-primitive concepts that the agent itself constructs. Here, there is the potential for *deep* conceptual growth: the agent is capable of building arbitrarily large conceptual hierarchies, as illustrated in Figure 3.

In the intermediate case, the agent is able to engage in concept-construction but only of a non-hierarchical variety. This will be the result, for example, if only categorical methods of construction are available. Categorical construction alone cannot support hierarchical development since 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 4 is a case in point.) In some cases, compositional construction methods may also be limited to non-hierarchical application. Whatever the reason, without hierarchical capabilities, the conceptual growth that can be achieved by the agent is necessarily *shallow*.

The distinction between deep and shallow conceptual growth is qualitative rather than quantitative. As a rule, deep conceptual development will furnish more concepts than shallow conceptual development. Generally, the number of concepts obtained through deep development 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

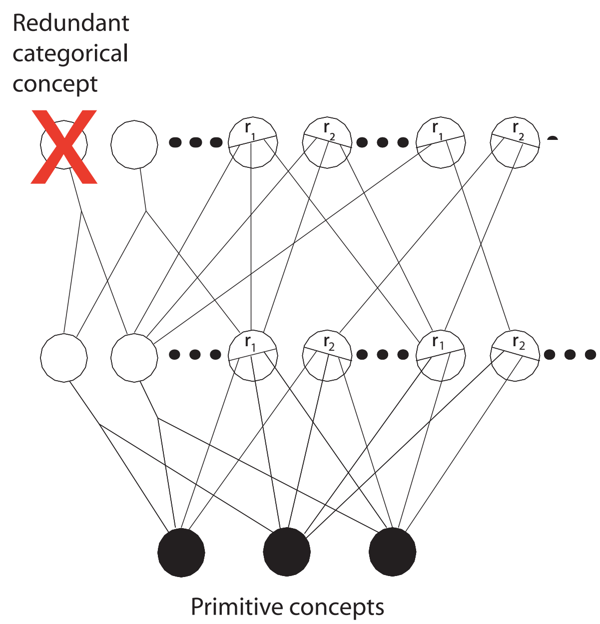

Figure 4: Development of a redundant categorical concept.

that it provides 'deep' conceptual development on the grounds that it makes use of compositional rules which can be applied to produced concepts. However, constraints are in place which limit the degree of recursion. The development is therefore finite. Were one to introduce a new rule so as to allow for infinite recursion (N -> S 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 is based on construction methods which cannot be applied to non-primitive concepts. 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' development has the potential to be larger than the 'deep' development defined by the grammar.

4 The complexity of conceptual growth

We have seen that the production of new concepts depends on (a) the available constituents (i.e., the primitive or non-primitive concepts) and (b) the available construction methods. To identify the limits of conceptual growth in a particular system we must now combine these two factors so as to identify the number of concepts that can be generated in a given case.

Initially, there are just the primitive concepts themselves. Let k represent the number of these. Each directly derived concept must combine some subset of the primitives. So their number must be related to the total number of subsets:

$$2^k$$

Each subset provides the basis for one categorical construct and, for each available relationship, one compositional construct. Letting r represent the number of accessible relationships, the total number of concepts which can be directly constructed is thus

$$2^k(r+1)$$

Applying categorical and compositional construction to the first layer of derived concepts generates a second layer of concepts. And we might expect that the number at this second layer will also be $2^k(r+1)$, with k set to the number of concepts at the first level. However, categorical concepts which subsume lower-level categorical concepts are redundant (as seen in Figure 4): they implicitly represent a class of instances already represented at the layer below. So from the first layer up, we should use

$$2^k r$$

The number of concepts k_i which can be constructed at any level i of the hierarchy may then be determined using

$$\begin{aligned} k_0 &= \text{the number of primitive concepts} \\ k_1 &= 2^{k_0}(r+1) \\ k_{i+1} &= 2^{k_i} r \end{aligned}$$

This, then, is the *growth formula* for conceptual development. It tells us the maximum number of concepts that any agent could generate using the given number of primitives and relationships; more specifically, it is the maximum number of concepts which may be developed up to a particular level of construction.[5]

As may be clear, deep conceptualisation of the unconstrained variety has the potential to create new concepts at a prodigious rate. There are two distinct exponential forces at work: in addition to the explicit exponential value, there is a second multiplicative effect implemented through the recursion. After construction has proceeded above the initial few levels, the number of potential concepts becomes astronomically large. In any real context, it can be assumed that only a small minority will have any value. But all are syntactically feasible.

[5] One might expect there would a necessity to sum evaluations for different levels. However, the inclusion of singleton elements means that every level in the hierarchy contains a proxy for all concepts represented below. The formula as presented thus defines the total number of concepts.

The strength of the effect can be illustrated with 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 eight concepts at the first level of construction and 256 concepts at the second level. At the third level, there would be more than 10^{77} to take into consideration and at the fourth level more than $2^{(10^{77})}$.

5 Net and gross conceptual complexity

Bundy's suggestion that Boden's definition might be reconstructed in terms of complexity concepts is the inspiration for the present approach. On the basis of the conceptual-growth quantification, it is now possible to formulate a plausible approach. In general terms, the complexity of a concept for an agent is the amount of work that the agent needs to do in order to build or obtain the relevant representation. The growth formula allows us to place an upper bound on this. Provided we can identify the level of the hierarchy at which the concept exists and the number of primitive concepts the agent has to begin with, we can then determine the maximum number of constructive operations required by evaluating the formula for the relevant value of i, setting k_0 equal to the number of primitive concepts and r equal to the number of relationships that the agent can bring to bear.

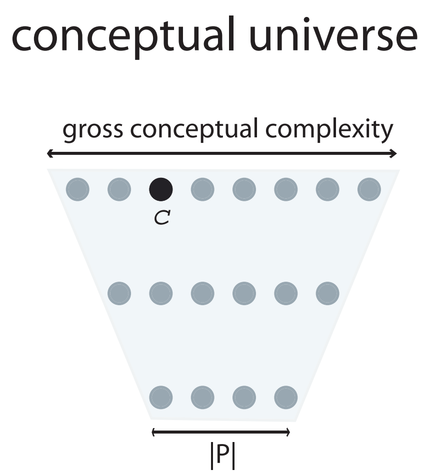

Figure 5: Gross conceptual complexity schematic.

Let this evaluation of a concept be termed the *gross conceptual complexity* and be denoted G. Taking $c(P, r)$ to be the level at which concept c may be constructed in terms of primitives P, under categorical construction and compositional construction based on r rules we have

$$G(c, r, P) = k_{c(P,r)} : k_0 = |P| \quad (1)$$

The gross conceptual complexity of any concept is the number of concepts which must be developed from the given foundation to guarantee its construction. The derivation of the measure is illustrated schematically (with regard to concept c) in Figure 5.

The complementary quantity, the *net conceptual complexity*, can then be defined as the entailed amount of constructive work less the amount of work that bringing the concept into existence *eliminates*. To calculate how much this is, we must look at the extension or coverage of the relevant concept.[6] 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 goal of concept formation. 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. Thus the representational gain is the size of X.

This leads directly to the following formulation for *net conceptual complexity*, to be denoted N.

$$N(c, r, P) = k_{c(P,r)} - |E(c)| : k_0 = |P| \quad (2)$$

Here $c(P, r)$ is again the level at which concept c may be constructed in a hierarchy based on r composition rules and primitives P, while $|E(c)|$ is the size of c's extension. (The derivation is illustrated schematically in Figure 6.)

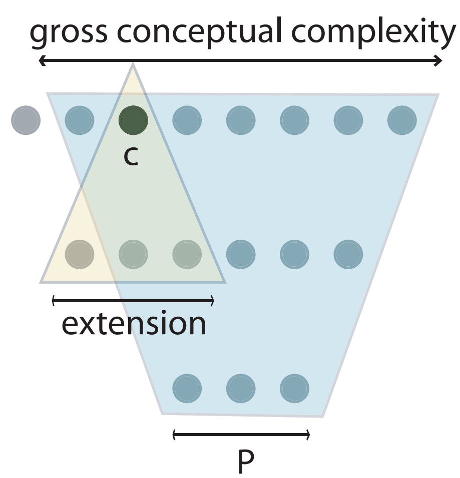

Figure 6: Derivation of net conceptual complexity.

[6] In this case, instances of a concept are themselves concepts.

If $E(c)$ is measured in terms of primitives P then the evaluation is *objective* in the sense that the representational power of the concept is being measured in units representing the ground elements of the concept hierarchy. However, it may be useful in some cases to formulate subjective measures, i.e., to differentiate the concept hierarchy under development from the entities in terms of which representational power is to be measured. In not restricting the way in which $E(c)$ is measured, the formula allows for this.

Values of N may be high or low or *negative*, the latter occurring in any situation where the representational benefit of bringing a concept into representation outweighs the cost. Net conceptual complexity is thus an implicit measure of the representational *gain* which results from bringing a particular concept into representation.

Might this particular complexity measure be sufficient to play the role envisaged by Bundy? Where Boden asserted that genuine creativity entails transformation, we might now consider asserting that it requires achievement of a value of N below some significant threshold. But aside from the problem of deciding how to fix the threshold, this would also leave open the issue of how values of N relate specifically to exploration and transformation. These are key elements in Boden's framework. To fulfil the objective of *reconstructing* the definition, a connection needs to be established.

6 Conceptual spaces as sponge zones

As Wiggins has noted (Wiggins, 2006b), the notion of conceptual space entails the assumption of there being some enclosing **conceptual universe**. This is the set of all concepts which, by definition, embodies all spaces. In terms of the analysis of generative concept-representation, an agent's conceptual universe may be defined in terms of the minimal set of primitive concepts and concept-construction methods from which all of the concepts may be generated. In developmental terms, the minimal set is the foundation that must be established in order for the agent to be capable (in principle) of developing all the embodied conceptual spaces.

Any concept in a conceptual universe may be labelled with its net-conceptual-complexity value. But given that the definition of a conceptual universe implies an exhaustive enumeration of all construction possibilities, it can be assumed that most such concepts are likely to have little value. Treating representational power as a measure of this, the net-conceptual-complexity for any such concept will be positive and relatively high, indicating that the cost of conceptualisation considerably outweighs any representational benefit. Other values of N 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 values of N, there are flatter areas where concepts show low values of N, and where concepts show negative N, there are marked depressions. The landscape thus shows regions of high, low and negative N.

Where a set of concepts form the primitive constituents of an explanatory theory, we would expect to find a marked depression in the area of the hierarchy immediately dominating the constituents. This 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* overall complexity of the system. Putting it graphically, the depression serves as a 'sponge', soaking up complexity from the hierarchy.

A conceptual universe is thus a landscape in which we expect to find pronounced areas of low or negative complexity — 'sponge zones' as they will be termed — located in parts of the hierarchy which immediately dominate any set of concepts forming the constituents of an explanatorily useful conceptual framework. By accessing concepts in these zones, the agent moves into representational profit, reaping a representational reward that exceeds the cost of development.

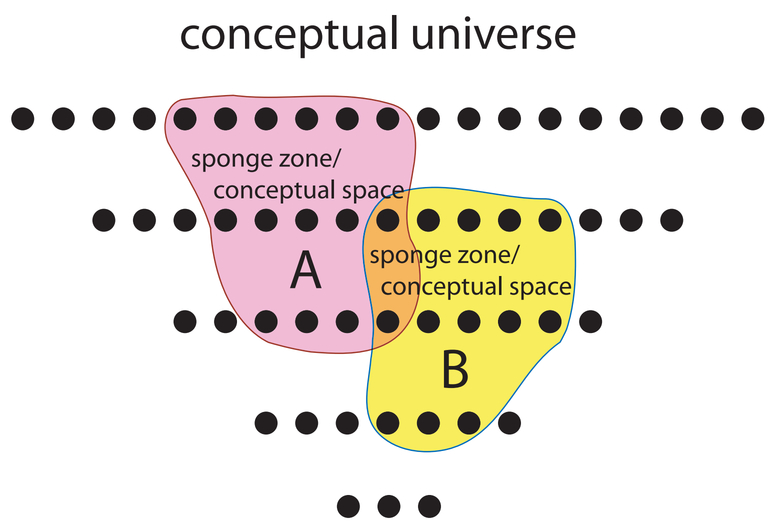

Figure 7: Sponge zones/conceptual spaces in a conceptual universe.

In Boden's framework, every conceptual space is based on primitives which, by definition, provide the basis of an explanatory framework or theory. There is thus a correlation between conceptual spaces and sponge zones (illustrated in Figure 7). Every conceptual space *is* a sponge zone with the explanatory power of the space corresponding to the capacity of the concepts in the zone to eliminate conceptual complexity.

7 The definition reconstructed

Interpreting conceptual spaces as sponge zones permits reconstruction of the creativity definition. The application of exploratory operations can be understood, now, to involve ordinary concept construction. The application of transformational operations, on the other hand, can be understood to involve exploration which specifically leads to the construction of sponge-zone primitives. More generally, exploration may be seen as the attempt to focus conceptualisation *within* zones of representational dividend. Transformation, conversely, can be understood as the attempt to open up new zones in which exploration may then be profitably applied.

Reconstructed in this way, Boden's definition is able to better fulfil practical requirements. The foundational distinction between transformation and exploration is given a definition which permits quantitative analysis (via the growth formula). Sponge zones are defined as complexity depressions within a conceptual universe. Transformation involves the construction of sponge-zone primitives. Any activity of this type is thus determined to be transformational in character.

A useful side-effect of this reconstruction is the generalisation of the all-or-nothing property of the definition, regarded by many as implausible. On the basis of measuring the reduction of N values within potentiated sponge zones, transformational creativity can be determined on a continuous basis. On the other hand, the distinction between negative and positive values provides the basis for identification of a quantitatively significant cut-point. This offers the potential for consolidating demarcation between 'radical' and 'ordinary' creativity.

In its use of a *net* complexity measure, the reconstruction also builds in an evaluative element. Whereas Boden's original definition allowed any transformationally active system to be deemed genuinely creative (including representation-changing watches), the reconstructed definition only extends the classification to systems which achieve a representational advantage. On this basis, transformation alone is not sufficient for creativity: it must also be advantageous.

This does not fully resolve the problem of how to deal with artifactual value, however. As Wiggins notes (2006b), there is increasingly widespread recognition that the value we expect to be associated with creative output may take many forms. While it is sensitive to a key element of intellectual value, net-conceptual-complexity ignores other types of value (e.g., associative, aesthetic and practical) completely. Even with regard to the measurement of intellectual value, it offers a rather specific view, being responsive specifically to explanatory 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 representational advantage and therefore creativity (Bundy, 2005).

8 Evaluation

Some consideration will now be given to the degree to which the proposed reconstruction meets more general requirements. The aim has been to devise a reconstruction of Boden's definition which answers the needs of systems-building work. The evaluation problem being only partly resolved, it can be assumed that the reconstructed definition must still be over-inclusive to some degree. However it fulfils the requirements of practical work better than the original, answering the need for a definition offering deterministic classification and a well-grounded account of transformational and exploratory behaviour.[7]

In resolving the ambiguity of the definition, the reconstruction addresses the concerns of authors who have noted this to be a difficulty, including Fetzer (1994), O'Rourke (1994), Bundy (1994), Ram *et al* (1994), Ram *et al* (1995) Treisman (1994), Weisberg (1994), Schank and Forster (1995) and Turner (1995). In some cases, the reconstruction can be shown to address other, more specific concerns. Ram *et al*'s requirements (1995, p. 122) are a case in point. 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 scope of the original account. In a comment which suggests they might view the reconstruction positively, Ram *et al.* 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.

Turning to Boden herself, the picture is mixed. There is some evidence that Boden would not approve the *method* of reconstruction. She has resisted 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)

Boden's concern is that any reconstruction of the definition which entails the assumption of a creativity continuum must implicitly deny the multi-dimensional

[7] Enabling Boden's distinction to be reconstructed in terms of operations carried out in a uniform search process, the reconstruction effectively eliminates any necessity to invoke notions of 'object' and 'meta' level search to explain the operation of exploration and transformation. In this model, any concept is at the object-level when considered in relation to super-concepts, and at the meta-level when considered in relation to sub-concepts.

nature of creativity. Yet it is common practice for judgements to be made about degrees of creativity without there being any implication that creativity *itself* is therefore one-dimensional. Boden's own definition is just such a judgement, albeit one in which the implied continuum has only two points. Indeed, there are many situations where we usefully bring to bear a one-dimensional measure in dealing with a multi-dimensional phenomenon. A familiar case involves the measurement of distance. The assertion that it is about five minutes walk from the station to the beach does not imply a denial that the space traversed is three-dimensional.

Other comments from Boden on the possibility of a complexity-based reformulation are more equivocal. Shortly after publication of the original theory, she listed some requirements (Boden, 1994a) that any such approach 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.' (ibid., p. 113)

She goes on to say that a satisfactory measure should somehow return a 'structured distance' and be sensitive to 'interesting differences' between artifacts of different styles or genres.[8] 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 reconstruction stand up to these requirements? Expressed specifically in terms of the way different modes of operation in a conceptual universe activate sponge zones, the reconstruction provides a detailed account of the relationship between transformational and exploratory creativity. It thus does 'relate' one process to the other. It also handles the tradeoff between conceptual 'depth' and 'number' in the sense that 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, then, the requirements from (Boden, 1994a) seem to be met reasonably well.

[8]She gives the example of a comparison of the degree of creativity in the *Mona Lisa* as opposed to the *Demoiselles d'Avignon*.

8.1 The systems evidence

While recognising the non-determinism of Boden's definition of creativity, it is important to take stock of the fact that it was never proposed as a self-sufficient entity, applicable on a stand-alone basis. In fact, Boden recognised the definition's deficiencies from the start.[9] However, her proposal was that it could be rendered precise through use of examples. To this end, she presented thumbnail sketches of a broad range of computational systems whose actions, she felt, illustrated the concepts on which the definition was based.

That this move failed to achieve its goal more fully may be due to the fact that the examples were presented only in fairly general terms. No formal specifications were offered and no algorithms were listed. The knowledge that complex computational systems typically embody considerable elements of poorly understood behaviour engendered the feeling that Boden's examples could not really clarify what the terms 'exploration', 'transformation' and 'conceptual space' meant. The fact that none of the example systems appeared to trade explicitly in these entities deepened scepticism. Nevertheless, it is clear that Boden's intention was that the example systems should be the means of making the definition precise. There is a need then to consider whether Boden's systems evidence is consistent with the proposed framework.

The computational examples she provides cover a wide territory. Working forwards from her Chapter 6 (Boden, 1994a, pp. 112-216), we see that creative, conceptual-space operations may involve processes ranging from the 'pattern-completion' and 'contextual memory' of connectionist networks (Rumelhart *et al.* 1986) to the algorithmic search for qualitative, physical laws featured in GLAUBER (Langley *et al.* 1987).

All of the systems she cites are unique in the sense that they embody a specific approach and terminology and implement a distinct form of processing. None of them offers any *explicit* account or illustration of Boden's concepts. In a sense, they are all special cases. However, a fair proportion of the systems are, or may be understood to be, *concept-learning* systems. The essential process in such systems is the generation of new concepts and in all relevant cases, structures are hierarchical. These systems are thus readily understood to be implementing conceptual growth of the type indicated in the reconstruction. The most obvious cases are AQ-11 (Michalski and Chilausky, 1980) and ID3 (Michie and Johnston, 1984, pp. 110-12) but also to be included are the mathematical concept learning system AM (Lenat, 1977) and its close relation EURISKO (Lenat, 1983).

On the understanding that analogy-formation is a type of concept construction, the list can be extended to include the 'analogical thinking' of ACME (Holyoak and Thagard, 1989) and the analogy-mapping of ARCS (Holyoak *et al.* 1988). Discovery and rule-formation programs which engage in the construction of representations interpretable as concepts can also be grouped under the

[9] The characterisation is 'vague' she writes, and no more than 'intuitive talk'. 'Anyone hoping for a scientific explanation of creativity must be able to discuss mental spaces, and their exploration more precisely.' (Boden, 1990 p. 73)

same heading.

But what of the systems cited by Boden which cannot be interpreted in terms of concept learning? Cases include the 'pattern-completion' and 'contextual memory' of connectionist networks (Rumelhart *et al.* 1986) and 'reasoning' of traditional problem-solving systems. Boden also refers to the combining of 'general and specific knowledge' and 'constraint satisfaction' featured in Harold Cohen's AARON system (Cohen, 1981).

Regarding these systems, a functional equivalence cannot be straightforwardly assumed, except perhaps in the case of Johnson-Laird's jazz improvisation program (Johnson-Laird, 1988) whose operations on 'complex, nested, hierarchical formulations' would appear to have a recognisable counterpart in the reconstruction. On the other hand, it would be a mistake to think that there is any fundamental inconsistency. All the systems which Boden cites are computational. Since concept growth (in the present model) is essentially the formation of a series of many-to-many mappings between symbolic entities, it is capable of replicating any function or structure of functions we like. Therefore it is capable in principle of replicating the behaviour of any computational system.

Consider for example the computation of the conjunctive truth function (i.e., boolean AND). This is a simple but fundamental computational process whose behaviour can be summarised using the following truth table.

A	B	$A \wedge B$
T	T	T
T	F	F
F	T	F
F	F	F

The computation of the function is effectively the implementation of a 2-to-1 mapping between boolean values. Implemented through the medium of concept generation, this computation takes the form of two constructions, one compositional, one categorical (see Figure 8). Regarding those of Boden's examples which cannot readily be regarded as performing some form of concept learning, there is no *fundamental* inconsistency, then. There is always the potential for the behaviour of such systems to be comprehended in terms of concept-construction processes.

9 Explanatory implications

While this paper's goals are practical in nature, Boden's theory is primarily oriented towards explanation. There is then the question of how the proposed reconstruction impacts the explanatory content of her theory. Those difficulties which arise specifically as a result of the definitional ambiguity — such as the ones noted by Ram *et al* (1995) — are addressed. Aside from this, what are the implications?

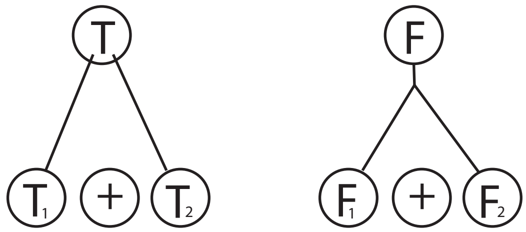

Figure 8: Computation of conjunction viewed as conceptual growth.

In mapping out her explanatory goals, Boden introduced the distinction between *P-creativity*, which is creativity in the personal context, and *H-creativity*, which is creativity in the historical context. Processes may exhibit P-creativity without necessarily being H-creative and questions about the psychological constitution of creativity can then be separated from questions about its constitution within artifacts. Following Wiggins (2006b), I take Boden's transformation/exploration framework to be essentially a model of P-creativity and the reconstruction framework, insofar as it has explanatory content, is in the same vein. That is to say, the reconstruction of the definition of creativity as a distinction between styles of conceptual growth must refer to *processes* of creativity rather than to any property embodied in specific artifacts.

Like conceptual spaces, conceptual universes and sponge zones are intrinsically *subjective* entities and, in principle, there is no basis for expecting any subjective-to-objective connection, i.e., no basis for thinking that P-creativity will necessarily align with H-creativity. But a potential benefit of the reconstruction, however is that it may allow for a narrowing of the gap.

9.1 Thought years

Creative agents may differ in terms of their realised conceptual hierarchy or — if they access different primitive concepts or construction methods — in terms of their *potential* conceptual hierarchy. For normal agents showing deep conceptual growth, there are infinitely many ways of showing the former type of difference but relatively few of showing the latter type. It is not unreasonable to think, then, that in some contexts a group of related agents might not show *any* differences of the latter type; i.e., they will share identical primitive conceptualisations and construction methods. In this context, there is the possibility of making creativity judgements which are objective across the group, effectively creating an identity between P-creativity and at least one aspect of H-creativity. If all agents in the group share the same potential conceptual universe, then net conceptual complexity measures are common for all concepts for all members of the group. On this basis, the determination that a constructive act entails a certain degree of P-creativity necessarily defines one dimension of its H-creativity.

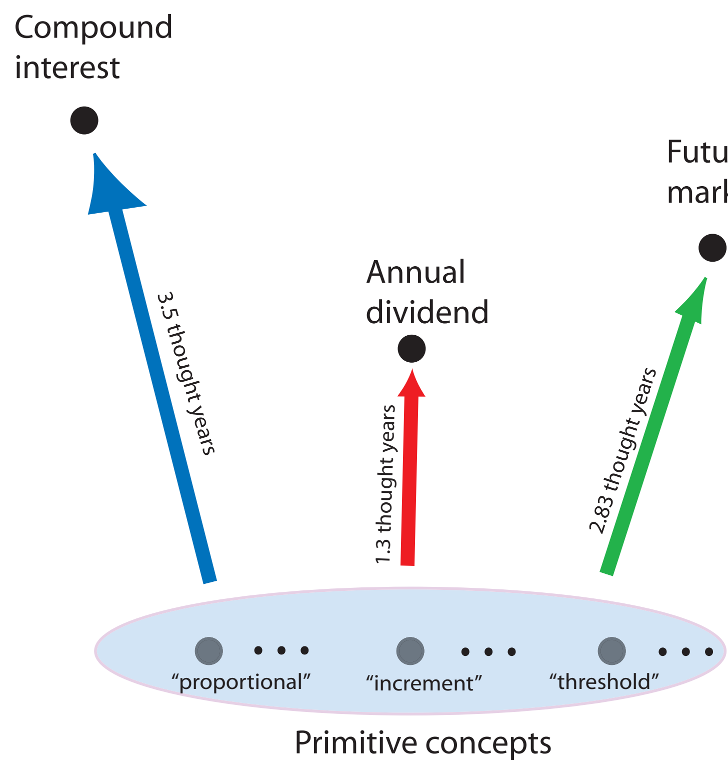

Figure 9: Imaginary thought-year distances for some financial concepts.

Should it transpire that humans constitute such a group there would be the prospect of a creativity metric that partially hybridises H- and P- components. This might allow the creative content of a concept to be measured, perhaps, in *thought years* — analogous to light years — where the value would be determined in terms of the total conceptualisation entailed in the concept's construction. Being a measure of total conceptualisation effort, gross conceptual complexity might fulfil this role. (Figure 9 illustrates the notion of 'thought year' in the context of some financial concepts.)

9.2 Boden and Koestler

Boden went to some lengths to differentiate the explanatory content of her theory from that of Koestler's earlier work. However, the reconstruction seems to hint at a reconciliation. The key difference 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 good model of any kind of concept-leaning process, particularly where there is utilisation of structured concept representations. To the extent that Koestler's analogy like 'matrix bisociation' is within this category, there is common ground.

There is also the potential for explaining one of the notable areas of agreement between the two theories. 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.' (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.

On the basis of the reconstruction, the reason why both Boden and Koestler see creativity as closely integrated in other processes of cognition is because it is functionally indistinguishable from the process of conceptual growth: 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 originating concepts or construction methods, different concept-generators may produce different results. But, taking the explanatory implications of the reconstruction at face value, the deed is done in the same way.

On this radicalised version of the integrative view shared by Boden and Koestler, 'creative masters' use the same mental equipment as 'common-or-garden thinkers'. While they may produce different results, this is a result of different experiences and starting conditions. Kekulé's dream-inspired innovation of the benzene ring concept may have been achieved by exactly the process used in the imaginary crossing-sweeper's innovation of the 'nicely swept crossing' concept. Indeed, if we make the not implausible assumption that some animals generate concepts using the same mechanism, there might be the implication that Kekulé's ring concept was derived by the same mechanism as would be used, say, by a dog 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 little difference.

10 Summary

While its commitment to the use of computational language suggested Boden's creativity theory might serve to provide detailed guidance for the design of creative systems, in practice the ambiguity and inclusivity of the foundational definition has obstructed such applications. Its orientation to a specific processing paradigm has also limited its technological utility, particularly in view of the increase over the last decade in the *variety* of work being done on computer-based systems to support or express creativity.

As a possible means of addressing this problem, Bundy suggested reconstructing Boden's definition using concepts of complexity. The present paper has shown one way this can be done. Utilising the principle of concept duality, the logical structure of conceptual growth has been determined and formulae presented which allow maximal growth rates to be calculated. Two complexity concepts have been introduced: *gross conceptual complexity*, which refers to the total cost of construction for a particular concept, and *net conceptual complexity*, which refers to the net cost, i.e., the representational gain less the construction cost. Consideration of the way net conceptual complexity will vary across conceptual universes, leads to the concept of the sponge zone and the observation that this is an analogue for Boden's conceptual space. Boden's distinction between transformation and exploration can then be viewed as discriminating different modes of concept growth and her definition of creativity re-rendered in a form which both eliminates the ambiguity and offers an assimilation of 'all-or-nothing' and 'continuous' perspectives.

The reconstruction fulfils many requirements that arise in the context of systems-building work and appears to meet all of those which Boden herself has listed for any complexity-based reformulation. The framework also answers (implicitly) requests for a theory which would clarify the basis of creativity in 'ordinary thought' and to deal more directly with the question of evaluation. It also supports some of the fundamental themes of Boden's account, including the stress placed on the notion that creativity should be regarded as an integral part of ordinary cognition. In sum, the reconstruction re-iterates much of what Boden has argued, but allows it to rest, definitionally-speaking, on firmer ground.

Acknowledgements

I would like to thank Alan Bundy, Maggie Boden, Andy Clark, Geraint Wiggins, Graeme Ritchie and Jon Bird for helpful comments on this material.

References

Boden, M. (1990). *The Creative Mind: Myths and Mechanisms*. London: Weidenfeld and Niicolson.

Boden, M. (1992). *The Creative Mind*. London: Abacus.

Boden, M. (Ed.) (1994a). *Dimensions of Creativity*. MIT Press.

Boden, M. (1994b). Reply to reviewers. *Behavioral and Brain Sciences*, *17*, No. 3 (pp. 558-570).

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 Niicolson.

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.*

Chaitin, G. (1966). On the length of programs for computing finite binary sequences. *Journal of The Association of Computing Machinery*, *13* (pp. 547-569).

Chaitin, G. (1987). *Algorithmic Information Theory.* Cambridge: Cambridge University Press.

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.

Kolmogorov, A. (1965). Three approaches to the quantitative definition of information. *Prob. Inf. Trans.*, *1* (pp. 1-7).

Kolmogorov, A. (1968). On the logical basis of information theory and probability theory. *IEEE Trans. Inform. Theory IT-14* (pp. 662-664).

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).

Ritchie, G. (2001). Assessing creativity. *Proceedings of AISB Symposium on AI and Creativity in Arts and Science*. York.

Ritchie, G. (2006). The transformational creativity hypothesis. *New Generation Computing*, *24* (pp. 241-266).

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).

Schmidhuber, J. (1997). Low-complexity art. *Leonardo, Journal of the International Society for the Arts, Sciences, and Technology*, *30*, No. 2 (pp. 97-103). MIT Press.

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*. Oxford University Press.

Weisberg, R. (1994). The creative mind versus the creative computer. *Behavioral and Brain Sciences*, *17*, No. 3 (pp. 555-557).

Wiggins, G. (2006a). A preliminary framework for description, analysis and comparison of creative systems. *Knowledge-Based Systems* (pp. xxx-xxx). Elsevier B.V.

Wiggins, G. (2006b). Searching for computational creativity. *New Generation Computing*, *24* (pp. 209-222). Ohmsha, Ltd and Springer.