How Thinking Inside the Box can become Thinking Outside the Box

Chris Thornton
Department of Informatics,
University of Sussex,
Brighton,
BN1 9QH,
UK
c.thornton@sussex.ac.uk

Abstract

While it remains a central reference for work in computational creativity, Boden's exploration/transformation model can be interpreted as having different meanings at the application level. As a result, programmers developing creative systems may have difficulty in realising the theory's practical value. The paper develops a formalisation which recasts the key elements in quantitative terms while reinterpreting the exploration/transformation distinction as a continuum. This has the effect of making the framework more amenable to practical application in the system-building context.

1 Introduction

In Boden's original model (Boden, 1990) creativity is seen as taking two forms:[1]

- guided search in existing conceptual spaces (termed *exploration*) and
- creation of new conceptual spaces (termed *transformation*).

For Boden, the second type is the more important, being the origin of 'true originality' (Boden, 2003, p 40) and for many there is an echo here of the intuition that thinking 'outside the box' can be more creative than thinking 'inside the box'.[2] However, for those interested in the building of creative systems, the identification of these categories raises questions about what sort of guidance and transformation will be most advantageous. The practitioner may want to ask 'is my system starting out with the right conceptual spaces?''is search being guided in a productive way?' and 'is the system generating the right sort of transformations?'

There may also be practical difficulties with the distinction between exploration and transformation. Since transformation can be seen as search performed in a space of conceptual spaces, it can be regarded as a form of exploration. (Wiggins, 2001, Ritchie, 2006) Wiggins has in fact taken the step of proving the formal equivalence of transformation and exploration by showing that transformation can always be viewed as meta-search. (Wiggins, 2006b, Wiggins, 2006a) The system-builder may need some way of deciding when transformation should really be treated as exploration, or vice versa.

One strategy is to treat Boden's theory as strictly explanatory and not seek any particular processing recipe, or at least nothing more than the general notion that creativity involves guided search in conceptual spaces.[3] Another approach is to seek a formalisation of the theory capable of meeting the needs of the systems builder. This is the approach pursued here.

2 Concept duality

Boden's observation that conceptual spaces must be generatively represented (Boden, 1990, p. 78) implies that identifying a new concept in a conceptual space must involve *construction* of the concept. This process must presumably make use of existing concepts. So exploration of conceptual space must involve processes of concept construction in which new concepts are constructed from existing concepts. But what are the possible forms of this process? In what ways can sub-concepts be combined to form a new concept?

Two cases can be discerned. First there is the case where the sub-concepts are treated as instances. Second there is the case where they are treated as constituents. The former type of construction is entailed in the construction of class-based, category-based and property-based concepts. The latter is entailed in construction of function-

[1] In Boden's revision of the model (Boden, 1998, 2003), a third form of creativity — combinational creativity — is identified. However, within the formalisation all exploratory creativity is combinational and for present purposes the two processes are therefore regarded as equivalent.

[2] Boden talks about the creative process mainly in terms of conceptual spaces explored by humans (and computers) but the model can be seen as covering search in *any* representational system. Ritchie, for example, interprets Boden's model in terms of search in a space of generic artifacts, subsuming concepts. (Ritchie, Forthcoming).

[3] This is the approach taken by Wiggins in his recent, information-based work on composition (Wiggins, 2007).

and relation-based concepts. For present purposes, construction in which the components are treated as instances will be termed **categorical** while construction in which the components are treated as constituents will be termed **compositional**.

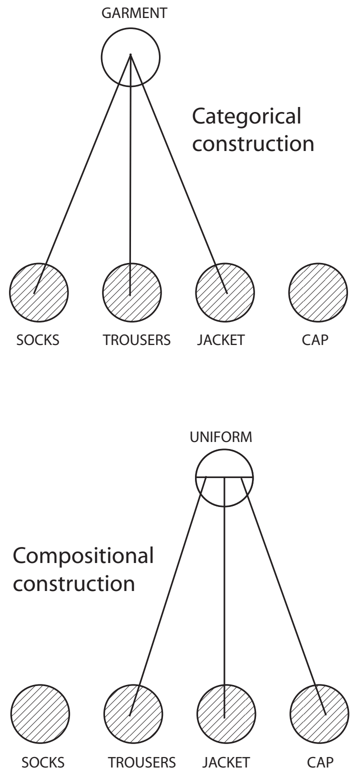

Figure 1: Concept-construction methods.

As an illustration of the distinction between categorical and compositional construction, consider Figure 1. This shows constructions for two clothing concepts. In the diagram, concepts are represented as circles with primitive concepts being cross-hatched. Three of these — SOCKS, TROUSERS and JACKET — may be combined to form the concept GARMENT. This is a categorical construction since it treats the sub-concepts as instances. In contrast, TROUSERS, JACKET and CAP may be combined to form the concept UNIFORM (as in 'MILITARY UNIFORM'). But here the construction is compositional since the components are treated as constituents of a new whole.

Note how the arcs connect to the internal structure of UNIFORM, reflecting compositionality while the arcs connecting to GARMENT combine, reflecting the fact that the concept is a class in which the components are alternative instances. (This convention is followed throughout.)

Whether a particular concept can be constructed in a particular way depends on the concepts and, in the case of compositional construction, the relations that can be applied. The fact that a particular concept can be constructed in one way from certain components does not mean that it cannot also be constructed in another. Nor does it in any way limit the ways in which the components can be used. The UNIFORM concept is here shown in a compositional construction. But it could also have been shown in a categorical construction, using sub-concepts such as ARMY UNIFORM, POLICE UNIFORM and SCHOOL UNIFORM. And while the compositional construct makes use of the 'and' relation, using this relation in a different way, or using a different relation altogether, other constructs could be formed from the same primitives.

3 The general form of conceptual development

Being able to construct new concepts endows an agent with the ability to 'explore' a particular conceptual universe. But what can we say about this universe? How big is it? What is the structure? If the agent is solely capable of categorical construction, only a finite number of new concepts may be constructed and these must correspond to the possible subsets of primitive concepts. If the agent is capable of compositional construction, then there is the possibility of an infinite expansion of concepts. But the rate at which new concepts may be developed depends on the relationships that may be applied.

For example, consider Figure 2. This illustrates conceptual development from three primitives (the cross-hatched circles) using categorical construction, and compositional construction with two relationships (labelled 1 and 2). As before, categorical construction is indicated using arcs which combine. But here compositional construction is indicated using arcs which connect with a bar labelled with the relationship invoked.

Initially, there are just the three primitives. For every way of grouping these, there is the potential for a categorically constructed concept and two compositionally constructed concepts — one for each available relation. These initial constructions generate concepts which are first-order with respect to the primitives. For every way of grouping the first-order concepts, the same situation applies with the result being a layer of second-order concepts. The number of potential constructs thus grows multiplicatively, with each level containing concepts of higher order.

But what is the maximum rate of growth? Initially, there are just the primitive concepts themselves. Let k represent the number of these. Each directly derived concept must combine some of the primitives. So the number of derivable concepts must be related to the 2^k possible subsets, but discounting the empty set and all singleton sets.[4] That is to say, the number of subsets on which new concepts can be constructed must be

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

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

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

[4]Not discounting them would allow construction of nonsensical and duplicate concepts.

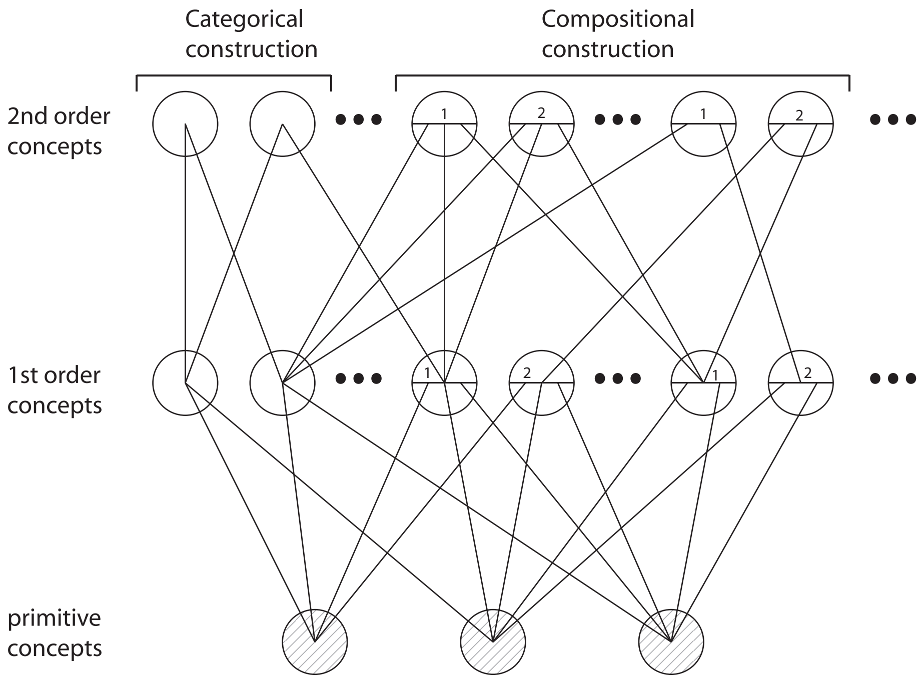

Figure 2: Potential concept constructions using 3 primitives and 2 relations.

Applying categorical and compositional construction to the first layer of derived concepts generates a second layer of concepts, and so on. The number of concepts which can be constructed at any level of the hierarchy thus depends on the number at the level below. This allows the rate of growth to be defined recursively. The number of concepts which can be constructed at any level i of the hierarchy may be determined using

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

Considering this growth formula, there can be no doubt that the number of possible constructions grows exponentially fast as development progresses. The number of constructions at a specific level of the hierarchy is multiplicatively related to the number at the level below both through the explicit exponentiation of k and through the multiplication with r. The end-result is that an unmanageably large number of potential constructions is reached very rapidly regardless of the initial base. For example, assuming a base of just three primitives and a single relationship, there are

- 8 potential concepts at level 1,
- 494 at level 2 and
- more than 10^{149} at level 3.

As a general rule, exhaustive expansion of a concept hierarchy beyond two levels of construction is intractable.

4 The Complex Extension

Examination of the mechanisms of concept construction reveals how conceptual spaces must be explored and answers some of the questions raised in the application of Boden's model. Exploration of conceptual space must proceed on the basis of categorical or compositional concept-construction. For any conceptual space, there must be an initial set of primitive concepts and, if compositional construction is used, a set of applicable relations too. In the case where only categorical construction is applied, the space is finite. If compositional construction is also available there is the potential for infinite development of the space. However, in this case, the rate of growth in any unrestricted process of construction is such that exploration of the space beyond the low-order concepts (i.e., 2nd or 3rd order) is prohibitively costly.

Understanding the general form of conceptual-space exploration, however, does not answer the critical questions about *which* forms of exploration are likely to be most advantageous. The system-builder needs to know something about the general principles of heuristic guidance. This is the 'nuts-and-bolts' end of the evaluation problem, of course — the general problem of how to discriminate concepts which have genuine value.

While evaluation does not figure in Boden's core model (except in the sense that transformation is deemed to lead to more 'radically' creative conceptualisation) it does figure considerably in her commentary and illustrative examples. A notion which seems particularly significant is that of *explanatory value*, i.e., the ability of one concept to account for, generalise or explain several others. This is to the good from the perspective of formalisation since generality has a well-defined meaning: the generality of a concept is *by definition*, the number of cases or instances which it generalises.

We can express the notion as

$$g(c) = |e(c)| \tag{1}$$

where $g(c)$ is the generality of concept c and $e(c)$ is its extension, i.e., its set of instances.

Can we incorporate a notion of explanatory value based on this equation? To do so we will need a mechanism for computing the extension of *any* concept. In the simple case of a categorical construct built directly on primitives, there is no difficulty. The extension is just the set of primitive concepts used in the construction. But how to compute the extension in the case where the concept is not defined directly in terms of primitives, or in the case where the construct is compositional?

In general, the instances of a concept are its possible manifestations and each distinct way of constructing a concept offers an alternative manifestation. Thus, alternative forms of construction *are* alternative forms of instantiation. Instantiation recapitulates construction. We can compute the extension of any concept, then, be evaluating the set of ways in which it can be constructed from the relevant primitives.

Illustrating the general idea, Figure 3 shows a concept hierarchy whose highest-level concept c is categorically constructed in terms of two compositional concepts which themselves are categorically constructed in terms of primitives. The extension of c contains its possible manifestations and these are identical to its possible constructions (as shown in the lower part of the figure).

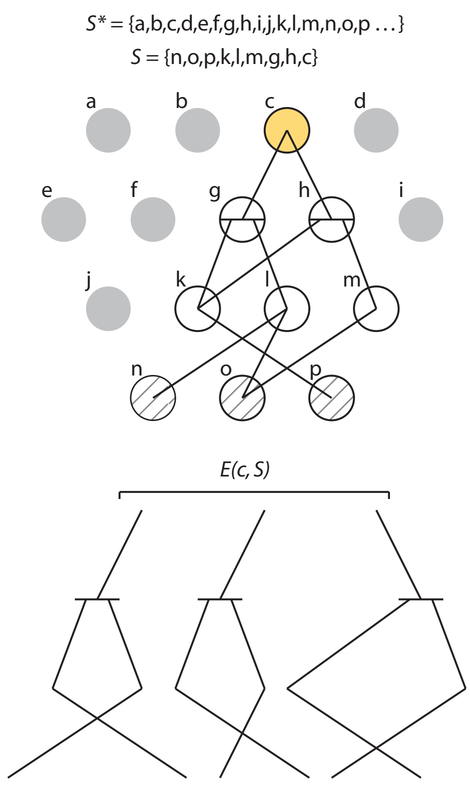

Figure 3: Derivation of conceptual complexity values.

There is a complication however since two forms of extension can be differentiated. On the one hand we have possible forms of construction made in terms of *existing* concepts of the space. On the other we have possible forms of construction made in terms of *potential* concepts of the space. In Figure 3 the set S is defined as containing all existing concepts of the space while $S*$ is defined as containing all potential concepts. Possible forms of construction for c in terms of already constructed concepts is therefore notated as

$$E(c, S)$$

with E being used instead of e to indicate the instances are conceptual structures rather than atomic entities.

In any conceptual space, the set of developed concepts must be a subset of the set of potential concepts. Thus

$$S \subset S*$$

and by the same token

$$E(c, S) \subset E(c, S*)$$

5 Full-house illustration

Where a single concept is constructed both categorically and compositionally, we have one extension for each construct. But though the contents must differ, the size of the extension (and by implication the generality of the concept) must be identical. Consider Figure 4. This illustrates compositional and categorical construction of the FULL HOUSE poker concept, whose extension we know to contain exactly 156 cases (because there are exactly 156 instances of a five-card hand meeting the definition of 'full house').

The primitive concepts in this example represent specific cards, with 1d = 'ace of diamonds', 2h = 'two of hearts' etc. Concept construction invokes one relationship and this is simple conjunction (shown here as '&'). Possible categorical concepts at the first level are subsets of cards and these include the suit concepts HEART, SPADE, CLUB and DIAMOND. Compositional concepts are conjuncts of cards and these include all instances of pairs (e.g., 1d & 1c), three-of-a-kind, etc.

At the second level of development, there is the potential for categorically constructed PAIR and THREE OF A KIND concepts. There is also the potential for compositional constructs combining instances of both PAIR and THREE OF A KIND.

Finally, at the third level of development, there is the potential for both a categorical and a compositional FULL HOUSE construct, with the latter combining the second-level PAIR and THREE OF A KIND.

Although it cannot be determined from the diagram, it should be clear that the complex extension of this concept will contain the same number of cases whether it is calculated from the categorical construct or from the compositional construct. The situation with the categorical construct is straightforward. It utilises 156 components, each one of which corresponds to a unique full-house hand. The compositional construct, on the other hand, is built from just two categorical constructs but the possible constructions for these combine to produce the same overall total.

6 Hill-climbing in the generality landscape

Conceptual space development is subject to 'guidance' favouring explanatory value just in case it shows a preference for concepts of greater generality. If we envisage

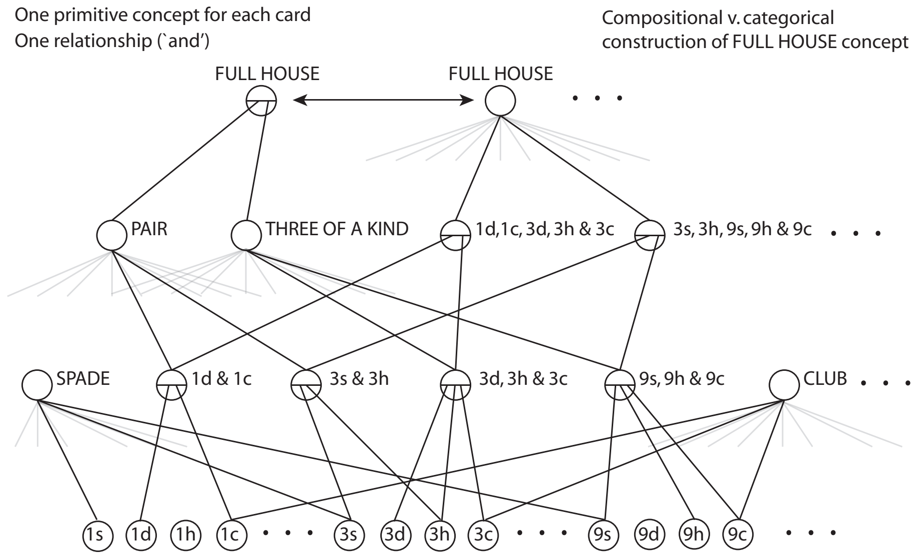

Figure 4: Dual developmental trajectories leading to FULL HOUSE.

a conceptual space as a landscape of generality levels, we can take this to imply that the process is essentially a form of hill-climbing search. Figure 5 illustrates the point. Taking the more darkly filled circles to represent concepts of higher generality, explanation-oriented conceptualisation should construct concept c before a, b, or d because it has a higher level of generality. An operational definition of Boden's explanation-oriented exploration is thus that it is hill-climbing search carried out in a concept-generality landscape. But where does this leave the process of transformation? Where does the creation of new conceptual spaces fit in?

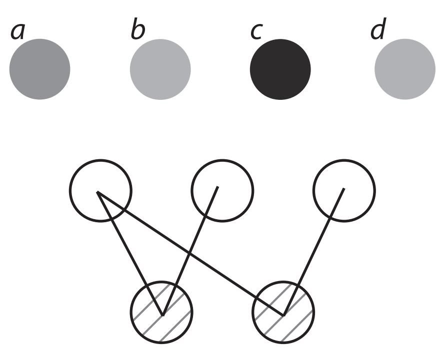

Figure 5: Explanation-oriented conceptualisation must hill-climb in the generality landscape.

Recall that generatively-represented conceptual spaces can only be explored through concept construction and that the components used in this process are existing concepts. On this basis any newly constructed conceptual space must exist within an overarching conceptual universe.[5] A conceptual space then forms a *subspace* in an enclosing conceptual universe. But taking into account Boden's point that the concepts in a conceptual space are connected through explanatory content, we can refine the definition and say that a conceptual space must comprise a subset of concepts whose extensional properties overlap.

In terms of the formalisation, then, transformation may be seen as *any* conceptual development which has the effect of creating a set of concepts with overlapping extensional properties. On this basis, all conceptual development involves exploration but only some has transformational impact. The separateness and significance of transformation is thus fully upheld within the formalisation. But the link with exploration is more clearly delineated.

The general effect of this aspect of the formalisation is to re-cast Boden's all-or-nothing distinction between exploration and transformation as a continuum. But this seems not to disrupt the main content of the theory in any serious way. Indeed, some commentators have argued that an exploration/transformation continuum might be preferable (cf. Treisman, 1994, Weisberg, 1994), while others have made proposals that implicitly assume its existence (cf. Bundy, 1994, Koestler, 1964). There is also some cause for thinking that a continuum may be more compatible with psychological observations of creative activity (cf. Ram *et al.* 1995, Perkins, 1981)

7 Transformation distinctions

Viewing transformation as a type of exploration allows some sub-classifications of the process to be introduced. The transformation process is defined as being any development which has the effect of creating a set of concepts with overlapping extensional properties. But, in practice, this might mean two different things. It could mean the *construction* of the required concepts or it might just mean the *enhancement* of their apparent generality, since this will in any case have the affect of making construction more probable. There is a distinction to be made, then, between transformation which *creates* a set of new con-

[5] Wiggins (2006b) draws the same conclusion on formal grounds.

cepts and transformation which merely helps to *potentiate* them by enhancing their apparent generality.

Further to this, there are two ways in which potentiating transformation may occur. On the one hand, there is the simple case where the constructed concepts *are* potentiated concepts. On the other, there is the case where the potentiated concepts are in a different part of the hierarchy altogether. These two forms of the process are tentatively termed ‘direct’ and ‘indirect’.

An an illustration of direct transformation, consider the example of Figure 6. This sketches the creative process underlying the innovation of the ‘reality TV’ concept. This has been seen as a novel combination of media genres such as the soap opera, the game show and the human-interest documentary.[6] The diagram reflects this, showing the REALITY-TV concept as a compositional construct using the GAME SHOW, SOAP OPERA and HUMAN DOCUMENTARY concepts.

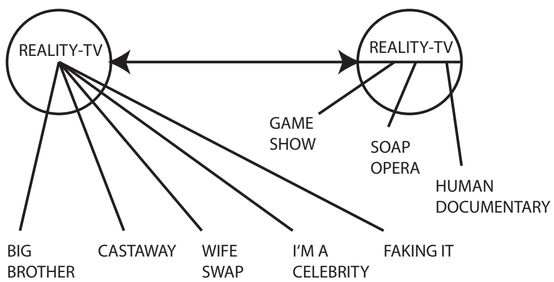

Figure 6: Illustration of direct transformation.

Prior to the compositional construction of REALITY-TV, concepts which might figure in its categorical construction (such as BIG BROTHER and CASTAWAY) are assumed to have negligible or nil generality. But once construction is complete, their generality is increased. All concepts within the category acquire an enhanced level of generality and the compositional construction therefore generates *potentiating* transformation. Intuitively, the innovation of REALITY TV has value due to the fertility of the concept. In terms of the model, it has value due to its potentiating, transformational impact.

For an example of indirect transformation, see Figure 7. This features concepts relating to performances of the ‘crossover’ Jazz pianist Jamie Cullum. Cullum is widely praised for his lively and imaginative improvisations which may involve use of pianos and other stage equipment (not to mention people) as percussive instruments. During intervals between passages of conventional keyboard wizzardry, Cullum may engage in unconventional percussive activity, e.g., drumming on the music stand, banging the lid of the piano up and down and head-butting the microphone.

Figure 7 envisages Cullum’s mould-breaking activities in terms of potentiating transformation. But unlike the previous case, the transformation here is indirect — the constructed and potentiated concepts being well separated. The compositional construct in the bottom-left

[6] The confrontation-prompting chat show, such as Jerry Springer, is another plausible constituent.

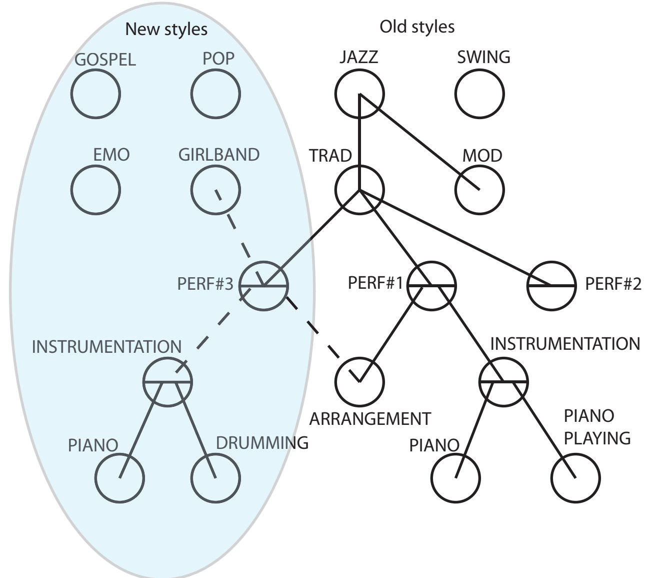

Figure 7: Illustration of indirect transformation.

corner represents the combination of PIANO and DRUMMING, i.e., it represents the idea of using a piano as a percussive instrument. Initially, this concept and all other concepts in the shaded area are unrealised. (The situation might correspond to an early stage in Cullum’s artistic career.) However, all concepts outside the shaded area are in existence and there is a well defined subspace of 4th and 5th-order concepts representing his early styles of performance (e.g. JAZZ and SWING).

With the innovation of the PIANO+DRUMMING composition, there is an immediate impact on the generality of concepts which trace their construction through this construct. Assuming DRUMMING has a variety of instantiations (which would seem to be the case) the generality of all these potential concepts is significantly increased. The innovation of the PIANO+DRUMMING concept thus has a potentiating transformational effect, enhancing the generality of a set of new 4th nd 5th-order concepts (GOSPEL, EMO, POP etc.). In performance terms, the innovation opens up a whole space of performance variation involving combination of novel percussive activity with previous unvisited musical territory.

8 Summary

A formalisation of Boden’s creativity model has been put forward which aims to answer the more detailed questions which can arise in applications work while at the same time staying, true to the original account. Observations about basic mechanisms of concept construction were used to derive a mathematical model of conceptual development. This was extended so as to ground the notion of explanatory value and allow explanation-oriented creative conceptualisation to be understood as hill-climbing search. On the assumption of conceptual spaces being collections of concepts with overlapping extensional properties, transformation was interpreted to be any exploratory conceptualisation serving to create or accentuate such a set of concepts. Several sub-classifications of the process were then distinguished. i.e., activating v. potentiation transformation and direct v indirect transformation.

A immediately apparent limitation of the formalisation is the restricted notion of value brought to bear. The only way in which creativity can be guided in the formalisation is through the generality measure. Therefore the only form of evaluation which is accommodated is that of explanatory power. While this may be appropriate in the context of scientific or intellectual creativity, it is presumably less relevant in other areas. Extending the formalisation so as to deal better with non-explanatory types of value is thus an important goal for future work.

A possible approach to the problem would be to look at the degree to which miscellaneous types of value can be effectively 'inherited' via the generality criterion. Recall that explanatory value is determined (for present purposes) in terms of extension size, i.e., its coverage of other concepts. And since these other concepts may themselves have *any* type of value, the valuation of the constructed concept measured purely in terms of coverage of subconcepts must reflect the types of value which those subconcepts have. So evaluation in terms of coverage might, under certain circumstances, implicitly support evaluation of a *non*-explanatory property.

For example, consider the cook who first discovered that sausages and onions are a good food combination. The artifact created is 'sausages & onions' and this presumably has elements of practical value (it satisfies hunger), associative value (it smells nice) and perhaps aesthetic value too (it looks nice). To the extent that these elements of value can be understood as originating in types of value associated with the *components* — sausages and onions treated independently — a generality-based model of evaluation might suffice.

However, while this approach may deal satisfactorily with inherited components of value, there is still the problem of accounting for that element of value attributable to the construction itself. Regarding this issue the most promising avenue may be to extend the formalisation so as to incorporate notions of conceptual blending (Fauconnier and Turner, 2005, Pereira and Cardoso, 2002) or combinatorial creativity (Butnariu and Veale, 2006, Veale and O'Donogue, 2000).

References

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

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.

Bundy, A. (1994). What is the difference between real creativity and mere novelty?. *Behavioral and Brain Sciences*, *17*, No. 3 (pp. 533-534).

Butnariu, C. and Veale, T. (2006). Lexical combinatorial creativity with gastronaut. *Proceedings of the Computational Creativity Workshop, ECAI 2006*.

Fauconnier, G. and Turner, M. (2005). *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. Basic Books.

Koestler, A. (1964). *The Act of Creation*. London: Hutchinson.

Pereira, F. and Cardoso, A. (2002). Conceptual blending and the quest for the holy creative process. *Proceedings of the 2nd Workshop on Creative Systems: Approaches to Creativity in AI and Cognitive Science, ECAI 2002, Lyon Frane*.

Perkins, D. (1981). *The Mind's Best Work*. Cambridge, Mass.: Harvard University Press.

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

Ritchie, G. (Forthcoming). Some empirical criteria for attributing creativity to a computer. *Minds and Machines*.

Treisman, M. (1994). Creativity: myths? mechanisms?. *Behavioral and Brain Sciences*, *17*, No. 3 (p. 554).

Veale, T. and O'Donogue, D. (2000). In S. Coulson and T. Oakley (Eds.), *Cognitive Linguistics* (special issue on Conceptual Blending), *11* (pp. 3-4).

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.

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.

Wiggins, G. (2007). *The Components of a Creative System: The Architecture of a Computer Composer*.