AI, CREATIVITY, REPRESENTATIONAL REDESCRIPTION, INTENTIONALITY, MENTAL LIFE: An Emerging Picture

TERRY DARTNALL[1]
Computing and Information Technology
Griffith University
Nathan Qld 4111 Australia

Abstract. This paper outlines a general theory of creativity and locates the concept of creativity with respect to the concepts of thought, knowledge and intentionality. The picture that I provide is outrageously simplified, but, as Dennett said recently, such idealisation is the price we must sometimes pay for synoptic insight (Dennett 1992). I first provide an account of a significant and central type of creativity. Then I show that there is independent empirical evidence that we are natively endowed with this kind of creativity. The emerging picture throws light on what it is to have a mental life, and on what it is for a representation to *mean* something to a system. This in turn gives a broader perspective, including a general theory of creativity that, on the one hand, distinguishes between creative and non-creative systems, and, on the other, accounts for individual differences between creative systems.

The paper is structured as follows. Section 1 argues that a significant and central type of creativity is the ability to 'break out' of a ruleset, or out of some other constraining framework of ideas. Section 2 looks at 'breaking out' in more detail. We break out when we access certain representations in such a way that we can reflect upon and modify them. But these are not any old representations. They are declarative structures that express the procedural knowledge that had hitherto driven our behaviour. In articulating these structures we acquire a mental life through discovering what it is that we know, think and believe. The beaver knows how to build a dam, but it cannot access this procedural knowledge and express it as a declarative structure: it does not know what it thinks, and consequently has no thoughts. Thus there is an intimate relationship between being creative and having a mental life.

Section 3 turns to Annette Karmiloff-Smith's Representational Redescription Hypothesis (RRH) (Karmiloff-Smith 1986, 1990, 1991, forthcoming a, forthcoming b, et al.; Clark & Kamiloff-Smith, forthcoming; Clark, forthcoming). The RRH is a theory of cognitive development that tries to explain how the human mind goes beyond domain-specific constraints. It maintains that we are endogeneously driven to redescribe our implicit procedural knowledge as explicit declarative knowledge, and to continue to redescribe our knowledge in increasingly abstract terms. Thus we are natively endowed with the kind of creative ability outlined in Sections 1 and 2. This empirical evidence is independent of the earlier argument. Creativity therefore takes centre stage in our efforts to understand human cognition, because we have grounds for believing that it is in our natures (and may be unique to *our* natures) to be creative.

This package of creative ability, mental life and native endowment appears to be more than a coincidence, and we might wonder whether there is a unifying factor that we have missed. This hunch is spurred by the fact that we would be surprised to find that a creative system could not, even in principle, have thoughts, or that a system with a mental life could not, even in principle, be creative. If this is correct then we have the following: (1) a system is creative if and only if it can redescribe, (2) a system is creative if and only if it has thoughts; therefore a system has thoughts if and only if it can redescribe. Thought, creativity and redescription begin to emerge as facets of the same thing. Can we throw light on this?

In Sections 4, 5 and 6 I suggest that the missing link is that it is only through redescription that we can have (genuinely intentional) mental states. I suggest that we initially acquire intentional states by being causally situated in the world. Initially we do not have access to these states, which constitute our implicit knowledge. Later we redescribe this knowledge in terms of explicit, accessible structures - that express our knowledge, thoughts and beliefs.

Section 7 returns us to creativity. This can now be seen as the struggle to articulate what it is to be (causally situated) in the world. The beaver stays on its plateau of behavioural mastery, but we go beyond this and transform our procedural knowledge into thoughts and theories that we can reflect upon and change.

This gives us a general theory of creativity that consists of a phylogenic component, distinguishing people from slugs and Sun4s, and an ontogenetic component that accounts for individual differences in terms of our abilities to deploy our redescriptive powers. Section 8 suggests that rare individuals, such as Mozart and Picasso, may have direct access to their redescriptive powers. If machine learning theorists can implement this ability then machines may become more creative than people.

[1] I wish to thank Andy Clark, Margaret Boden and others for helpful discussions of this paper.

1. Boden on 'Breaking Out'

Margaret Boden (1990, forthcoming) argues that creativity is the exploration and transformation of conceptual space—or more correctly, of conceptual spaces. A conceptual space is a space of structures generable, that is, defined, by the rules of a generative system. We explore such a space by studying the structures within it, and we transform it by modifying the rules of the system to produce different types of structures.

Boden argues for this by first considering the notion that creativity involves the novel combinations of old ideas. Creativity does involve novelty, but Boden observes that such accounts do not tell us *which* combinations are novel, nor how novel combinations can come about. Her main—and related—criticism is that many creative ideas not only *did* not occur before, but *could* not have occured before. Previous thinking was trapped in a framework, relative to which new ideas were impossible. Consequently, creativity requires us to 'think the impossible', to have ideas that are impossible in the present framework: we need to 'break out' of a conceptual space by changing the rules that define it. Kekule, for instance, broke out of the space defined by the rules of nineteenth-century chemistry, and opened up the new space of aromatic chemistry.

2. Breaking Out and Having a Mental Life

Let us look at a couple of cases of 'breaking out'. First, suppose that you are a Newtonian physicist who believes that light travels in straight lines. There is now a sense in which it is unthinkable for you that light waves should bend, since it is *logically* impossible for Newtonian physics to be true and for light waves to bend: necessarily, *if* the laws of Newtonian physics are true *then* light does not bend. Whilst you remain within the Newtonian framework you are trapped by this necessity. Or imagine that you are following the rules of Euclidean geometry. Within the framework defined by these rules it is *logically* impossible for the sum of the angles of a triangle to be other than 180 degrees. We can break out of these frameworks by representing the rules to ourselves. Once we have done so we will see that the necessity lies in the hypothetical '*If* these rules *then* X'—*not* in X itself.

This shows us how we are limited by the assumptions that we make, but it does not explain why we find it so hard to articulate the assumptions. I suggest that we can throw light on this by distinguishing between rule *users* and rule *followers*. Rule followers are subject to rules. They do what the rules say. Computer programs follow rules when we set them the task of (say) generating proofs in Euclidean geometry. Rule users, on the other hand, can access the rules as a declarative structure.

Sometimes we are rule followers—when we follow the rules of grammar, for instance—and we are trapped to the extent that we are merely followers. We break out when we become users—when we articulate the knowledge implicit in the rules we had been following.

We can cast light on this by imagining a 'cognitive ladder'. At the bottom of the ladder are marshflies, hoverflies and ants. The marshfly is exasperatingly persistent and follows us about despite our attempts to swat it away. Hoverflies meet and mate in midair. Ants exhibit apparently complex behaviour. These creatures, however, have little or no knowledge, either procedural or declarative. Their behaviour is principally driven by information that is in the environment, not in the organism (Dreyfus 1979). The marshfly follows the carbon dioxide that we emit. Hoverflies are hardwired to transform a specific signal into a specific muscular response (Boden 1990). Ants can detect contours, and have routines for selecting the most level path (Simon 1969; see discussions in Pylyshyn 1979/1981, Winograd 1981; also Dreyfus & Dreyfus 1987.)

When we discover these things we should abandon any suspicions that we might have had about these creatures having a mental life. They are not merely mindless: they have no knowledge worth speaking about. They are not even rule *followers*.

Midway up the ladder is the beaver. It has structured, procedural knowledge about how to build a dam. However, it has no access to this knowledge. It is trapped in a procedural framework, doing what its rules tell it to: first you put in stones, then you put in big logs, then you put in small ones. If a storm sweeps away the small logs, the beaver demolishes what's left of the dam and starts again. (This may not be empirically true of beavers—if it is not, run the argument with a more stupid creature.) In this respect the beaver is like a computer program: it is a rule follower rather than a rule user. It does not represent the rules as an accessible structure. If it did it would be able to directly access the rule for putting small logs on top of big ones.

Nor does the beaver have any thoughts. We see it scurrying about trying to put the log *there*. The log keeps slipping out, and the beaver keeps putting it back. We want to say that the beaver *thinks that the log should go there*—but we do not want say that it *has the thought* 'the log should go there'. This is a distinction that we commonly draw with people: the tennis player is 'thinking what she's doing' when she plays intelligently, but this does not mean that she has explicit thoughts about what she is doing.

(Adrian Cummins (1990) draws a similar distinction between thoughts that have, or do not have, conceptual content. If Jo believes that Bill is a bachelor, then Jo has grasped the concept 'bachelor'. But when we say that Fido 'believes that the noise (or the scent) comes from the south', we do not believe that Fido has grasped the concept 'south'! Cummins believes that Fido has no structured or contentful thoughts. I believe that he has no *thoughts* at all.)

We tie these facts together when we say that the beaver 'doesn't know what it thinks'. It thinks that the log should go there, but it doesn't know that it thinks this. It does not have access to its procedural knowledge.

Now consider Le Penseur (Rodin's sculpture of someone engaged in deep thought—head bent, brow furrowed). Le

Penseur's human equivalent has thoughts by virtue of being able to express his knowledge as a declarative structure, reflect upon it, and change it if necessary. In doing this he is exercising an ability that all of us have.

Enter the empirical evidence—that we are *natively endowed* to redescribe our procedural knowledge as declarative knowledge (and to continue to redescribe it at increasingly abstract levels). We are endogenously driven to go from being rule followers to rule users, thereby 'breaking out' in the sense described.

3. Representational Redescription

The RRH is a theory of cognitive development proposed by Annette Karmiloff-Smith that tries to account for the way in which the human mind goes beyond domain-specific constraints. It maintains that the mind is endogenously driven to go beyond behavioural mastery and to represent its knowledge to itself in increasingly abstract forms. It does this, as it were, under its own steam—without need of exogenous pressure. Initially, the system's knowledge is embedded in procedures. It is implicit in the system's abilities. It is not available to other procedures, nor to the system as a whole. As Karmiloff-Smith puts it, it is *in* the system, but not available *to* the system. Later the system redescribes it as declarative knowledge, which is available to other procedures. The system continues to redescribe its knowledge on increasingly abstract levels, all of which are retained and may be accessed when necessary.

The most commonly cited evidence for this is the 'funny men' pictures (see Figure 1).

We start with a group that has achieved mastery of a task. Children between the ages of four and six, for instance, are able to draw houses and stick figures of people. But when they are asked to draw a funny man or a funny house, they can do very little. They are locked into their house-drawing or man-drawing procedures: first you draw the head, then you draw the body, then you draw the legs etc. To begin with they can only modify parts of the procedure, such as drawing a square head. But slowly they acquire the ability to access the procedures as declarative structures, whereupon they can apply the rules in any order they wish. Children between eight and ten are quite fluent at this, and can draw hands on the end of legs and feet on the end of arms. They can even mix procedures together, and draw centaurs, and heads that have windows.

'But', you might say, 'the younger children could have done things like that—it just didn't occur to them to do so'. To meet this objection, Karmiloff-Smith asked the 4-6 year-olds to draw a man with two heads. Commonly, the children drew a man, drew a second head, and then went on to draw a second body-and-legs: they were locked into a man-drawing procedure.

The plot so far is this. A central and significant type of creativity is the ability to break out of a ruleset, to go from being subject to rules to being able to articulate and use them. With this transcendence come thoughts and the mental life. And there is independent evidence that we are natively endowed with such an ability.

We are bound to ask more. We are bound to ask whether it is *only* redescribers who are creative and who have mental lives. In the next section I argue that this is indeed the case, since our declarative structures can only be *about anything* if they are grounded in implicit knowledge brought about by our being causally situated in the world, and are then redescribed as accessible structures.

4. Intentionality and Redescription

The problem of intentionality is the problem of how mental states, symbol structures or artifacts can be *about* anything. I have argued that a system is creative if and only if it can access its knowledge structures. But what is it for a structure to represent knowledge for a system? The standard computational account is that the structure is in the mind's 'knowledge bin' or "thought bin" and that the mind manipulates it according to formal rules (Dennett1986 calls this "High Church Computationalism"; see also Richardson 1981, Fodor 1987). The Knowledge Representation Hypothesis (see esp. Brian Smith 1985) says that a system contains knowledge if and only if it contains a syntactic structure that means something to us, and that causes the system to behave in an appropriate way. E.g., the system knows that tigers bite if and only if it contains a structure such as 'Tigers bite' that causes it to get up trees in the presence of tigers.

It is not clear how locating a structure within a system can enable it to know what the structure means. And understanding a sentence surely involves more than being propelled by its morphology. Rather than pursuing these points I shall outline an account of intentionality that gives a major role to representational redescription, and that throws light on creativity.

5. The Information Theoretic Account of Intentionality

My account is based on the Information Theoretic account of intentionality. This is most commonly associated with Fred Dretske, though John Heil and K.M. Sayre have developed their own versions (Dretske 1980, 1981, Heil 1983, Sayre 1986). The Information Theoretic account exploits the mathematical theory of information advanced by Shannon and Weaver, which says that one state carries information about another just to the degree that it is lawfully dependent on that other state. Dretske realised that this sheds light on intentionality, since "Any physical system, then, whose internal states are lawfully dependent, in some statistically significant way, on the value of an external magnitude... qualifies as an intentional system." (1980, p 286). Thus *aboutness* is not a unique feature of mental states but is found in all causal relationships.

Now, however, Dretske faces the problem of cognitive error (of how we can have false beliefs, etc). He maintains that information is distorted by the cognitive system. His critics reply that this commits him to saying that we can gain knowl-

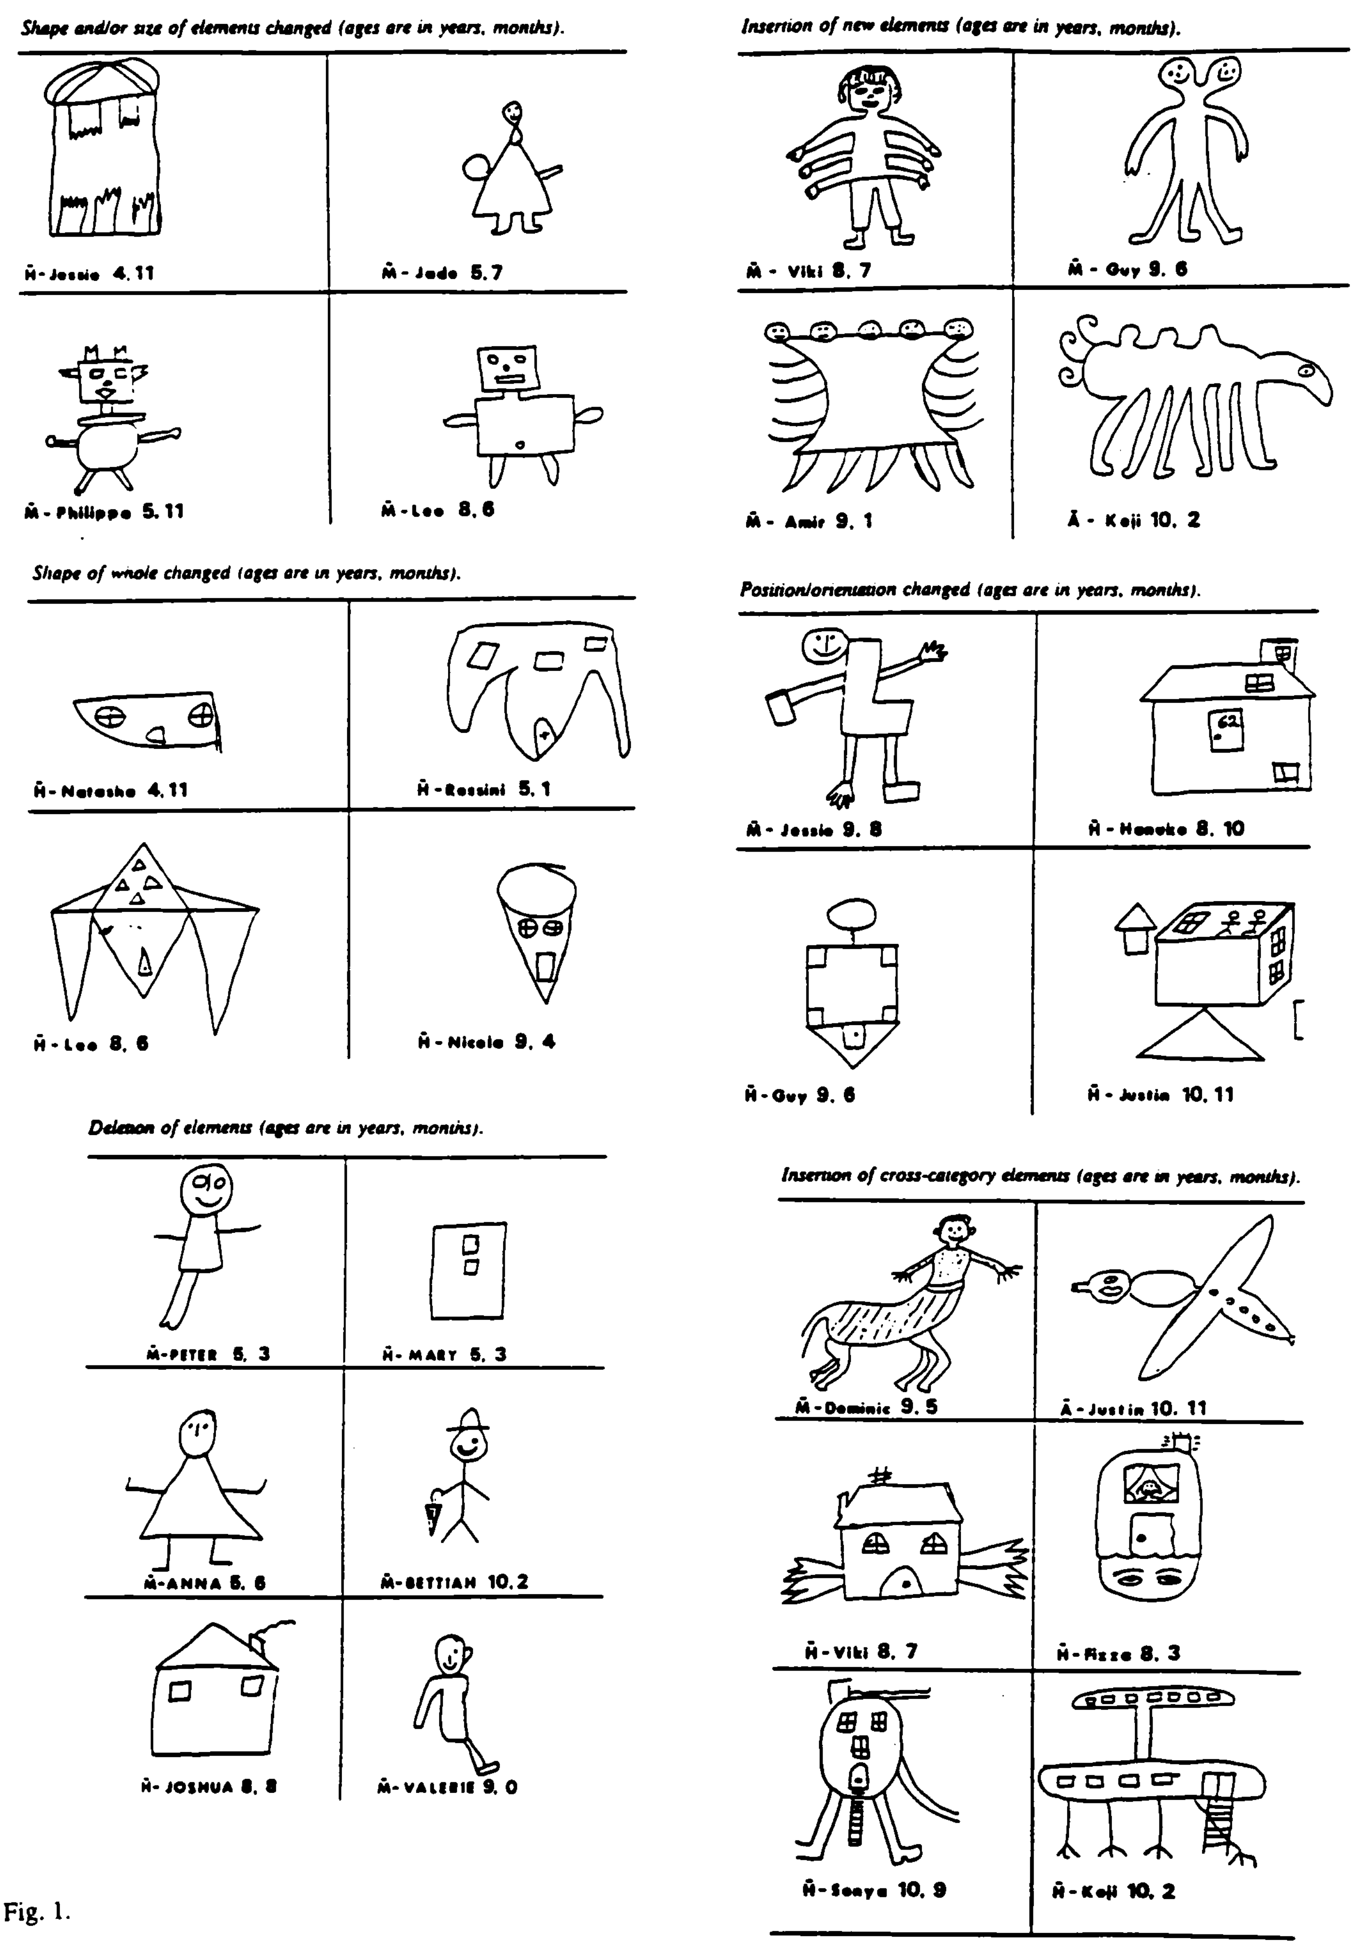
Shape and/or size of elements changed (ages are in years, months).
H-Jessie 4.11
M-Jade 5.7
M-Philippe 5.11
M-Lee 8.6
Shape of whole changed (ages are in years, months).
H-Natasha 4.11
H-Rossini 5.1
H-Lee 8.6
H-Nicole 9.4
Deletion of elements (ages are in years, months).
M-PETER 5.3
H-MARY 5.3
M-ANNA 5.6
M-BETTIAN 10.2
H-JOSHUA 8.8
M-VALERIE 9.0
Insertion of new elements (ages are in years, months).
M - Viki 8.7
M - Guy 9.6
M - Amir 9.1
A - Koji 10.2
Position/orientation changed (ages are in years, months).
M - Jessie 9.8
H - Hanako 8.10
H - Guy 9.6
H - Justin 10.11
Insertion of cross-category elements (ages are in years, months).
M - Dominic 9.5
A - Justin 10.11
H - Viki 8.7
H - Fizza 8.3
H - Sonya 10.9
H - Koji 10.2

Fig. 1.

edge by removing the errors from our cognitive systems. This implies that cognitive systems merely distort information—and this runs counter to our belief that more complex systems can gather more and better information. I shall return to this.

6. Combining Accounts

Let us run the story with a redescriber. The redescriber is placed in the world. Its peripheries are bombarded. It develops internal states that constitute procedural knowledge: it can build dams, draw pictures, pronounce words, etc. This is knowledge *in* the system that is not available *to* the system. It consists in the system being in a certain state. Because this state was causally determined by the world in a lawlike way, it is knowledge *about* the world.

Consider NETtalk. When it is bombarded with phonemes, NETtalk develops internal states that are about phonemes. It settles into internal states that (to use Dretske's words) are "lawfully dependent, in [a] statistically significant way, on the value of an external magnitude". These states account for its ability to pronounce words. Of course, it requires, e.g., cluster-analysis to identify the states that it has settled into, so that, as Clark and Karmiloff-Smith (forthcoming) say, the states are not available to NETtalk. They are only just available to us!

We can provide a straightforwardly causal account of the intentionality of such a system. Moreover, its knowledge is relatively undistorted, since its internal state is a direct reflection of its environment. On the other hand, this knowledge is limited, inflexible and inaccessible.

Back to our redescriber. The redescriber now redescribes its implicit knowledge as explicit knowledge—as knowledge available *to* the system. This is about the world because it is a redescription of knowledge that came about by *being situated* in the world. Its intentionality lies in its pedigree.

With redescription comes risk and the possibility of cognitive error, which is the price that we pay for increased abstraction, flexibility and an improved ability to gather information. This resolves Dretske's problem of how cognitive systems can apparently go wrong. They do not go wrong—but there is a trade-off between risk and flexibility. (Margaret Boden has pointed out to me that skills and abilities are fallible as well—see her chapter on Hoverflies and Humans in Boden 1990. Hoverflies can miss one another in mid-air. Beavers can build faulty dams by putting in the small logs first. The point I am making is that at least some of the fallibility that is associated with higher cognitive systems—e.g. false belief—is a natural consequence of redescribing our knowledge in increasingly abstract terms.)

Thus redescribers can acquire intentional states and gain knowledge through being situated in the world. They then redescribe this knowledge in terms of accessible, modifiable structures that constitute genuine representations, such as thoughts and beliefs. They are genuine representations, moreover, *because* they are redescriptions. They have intentionality because they are redescriptions of states that, as it were, have 'prime' intentionality—intentionality that is due to direct causal influence. Genuine representations are *grounded* by virtue of the system 'being in the world'.

Are *all* genuine representers redescribers? Consider the beaver on the one hand, and a Sun4 on the other. The Sun4 contains datastructures but has no mental states. Its datastructures do not have any intentionality for it, do not express knowledge for it, do not 'mean anything to it'. Let us suppose that this is because they are not redescriptions of prime intentionality. The beaver (certainly the marshfly) is at the other extreme. It has prime intentionality, but doesn't know what it thinks. Let us suppose that this is because it hasn't redescribed its knowledge into accessible structures. Given these assumptions, the Sun4 and the beaver have no mental life because they are not redescribers. If having no mental life is always due to one of these two impediments, then *only* redescribers have a mental life. (Kant said that "concepts without percepts are empty, percepts without concepts are blind". We might say that 'ungrounded representations are empty, prime states are blind'.)

This, however, only restates the problem. Consequently I shall resort to an onus of proof argument—sometimes referred to as 'the best bad argument in philosophy'. Let us assume that (despite some contemporary philosophical opinions) we really do have mental states. These are intentional states. Now, what grounds have we got for thinking that intentional states can arise in any way other than through causal stimulation from the environment? If we have no grounds then we must believe that this is how intentional states come about in us. The RRH then tells its redescriptive story. Surely the point about being a *physical* symbol system, rather than a virtual one, is that it can causally acquire intentional states. Then, if the architecture is right, it can redescribe the knowledge implicit in these states in terms of explicit symbolic representations. Cognitive scientists and machine learning theorists must show us how this is done.

7. A General Theory of Creativity

We now have the rudiments of a general theory of creativity. Such a theory contains a phylogenic component and on ontogenetic component. The phylogenic component concerns the difference between creative systems (such as people) and non-creative systems (such as slugs and Sun4s). The ontogenetic component is concerned with individual differences in the ability to redescribe.

The phylogenic component locates creativity at what Clark & Karmiloff-Smith (forthcoming) call "a genuine joint in the natural order". One of the great evolutionary divides appears to be between systems that can redescribe their procedural knowledge as accessible structures, and that have thoughts and a mental life, and systems that cannot and do not.

Now look at ontogeny. The child redescribes its procedural knowledge as declarative knowledge. Here ontogeny recapitulates phylogeny. We then continue to update our framework of beliefs, struggling to give expression to what it is to be causally situated in the world.

Of course, such frameworks are not everywhere grounded in redescription. They are only grounded as a whole. Quine and Davidson talk about a loosely grounded *web* of belief. The RRH says that there are *layers* of (re)description, grounded (if my account is correct) in the bottom layer of prime intentionality. No layers are lost or destroyed, and all can be accessed. And our articulated beliefs, as well as a fluctuating environment, cause us to continually restructure and update the framework.

This is most obvious in the Arts. Art tries to express what we know or believe by articulating it as a declarative structure: Munch's Scream, El Greco's Revelation of St John, Bach's Passions, Owen's war poetry—all articulate by externalising. We sometimes say that we do not know what we believe until we have written it down. In writing *Sons and Lovers* D. H. Lawrence articulated his complex attitude to his parents, that had driven him to regard his father as a despot and tyrant. Having voiced his attitude he was able to evaluate it and change it. "We shed our sicknesses in books," he said.

The theory of redescription explains what is going on here, and shows why it is creative. Redescription does not give expression to beliefs that were already there: *it actually brings them into existence*. Hitherto we *believed that*... such and such (just as the beaver believes that the log should go there, but does not have the thought 'the log should go there'). We were in states that drove our behaviour. Redescription expresses the dispositions inherent in these states as accessible representations. This is not creating something out of nothing—but it is the next best thing.

This resolves an apparent problem with Boden's account. There is, I think, a danger of her account begging the question by saying that routine creativity is the *creative* search of conceptual space (or the search of a limited, 'creative' space). Since generative rules only determine wellformedness, we need to know why Mozart was consistently more successful (that is, more creative) in searching through the musical space of his day than Salieri was.

I suggest that the answer lies in distinguishing between three kinds of art. The first merely explores a space of structures—say, musical structures. We might be inclined to regard this as little more than a technical exercise. At the other extreme we have art that breaks out of the confines of the genre. The third kind remains within the genre, but breaks out of another space by giving voice to feelings and attitudes that had not been voiced before, either by the individual or at all (Boden's distinction between P-creativity and H-creativity).

8. Are there Two Types of Creative Ability?

For most of us the creative process is a slow and painful one. We engage in a cycle of externalisation and evaluation, in which we externalise something, evaluate it, readjust our goals, and repeat the process (Edmonds, forthcoming; Sharples, forthcoming). Attempts to produce computerised art and computer-assisted art have found precisely this. In developing his figure-drawing program AARON, Harold Cohen found it necessary to continually externalise and evaluate AARON's ability, in order to find out what its procedural knowledge enabled it to do. Edmonds (forthcoming) quotes him as saying "we externalise in order to find out what it is that we have in our heads... It is through this externalising process that we are able to know what we believe about the world" (Cohen, 1983; see also McCorduck, 1991).

A few rare individuals, however, seem to create with consumate ease. Mozart said that he would experience a composition all at once—in a moment—both before and after he wrote it down: "Aah, what a feast is there," he said. We may be sceptical, but the manuscripts contain no errors. There is no apperceptive agonising here, no evidence of an externalisation/evaluation cycle. Picasso, too, was unable to put a foot wrong: "I do not seek—I find", he said. Mozart and Picasso seem to have directly and unproblemmatically achieved what the rest of us blindly struggle for. Mozart's talk about instantaneousness suggests that he had direct and spontaneous access to a declarative structure. In 'A Conversation with Einstein's Brain' the tortoise/Doug Hofstadter invites us to imagine what it would be like to experience a piece of music instantaneously by looking at the side of a long-playing record: "since all of the music is on the *face* of the record, why don't you take it in at a glance, or at most a cursory once-over? It would certainly provide a much more intense pleasure". And "Why don't you paste all the pages of the written score of some selection upon your wall and regard its beauties from time to time, as you would a painting?... instead of wasting a full hour listening to a Beethoven symphony, on waking up some morning you could simply open your eyes and take it all in, hanging there on the wall, in ten seconds or less, and be refreshed and ready for a fine, fulfilling day?" (Hofstadter 1981).

Was it like this for Mozart? (Remember that he claims to have experienced it all at once *before he wrote it down!*) The RRH says that we have the ability to spontaneously redescribe, so perhaps Mozart had (almost?) full control of this ability.

If so, there are two types of creativity ability. The first involves the externalisation/evaluation cycle: we construct theories about our abilities by observing them in action. The second rests on the fact that abilities are grounded in states. Balls are able to bounce because of their molecular structure. NETtalk is able to pronounce words because it has settled into a subsymbolic state. Our abilities are similarly grounded in states. Now suppose that we can redescribe the knowledge that is implicit in these states. Then we would no longer need to construct theories about what we know and believe by observing our actions. The RRH tells us that we have exactly this ability. Mozart, it seems, was especially good at exercising it.

9. Conclusion

The RRH is a philosopher's stone for creativity research: it solves problems, generates plausible explanations, and leaps mighty buildings at a single bound. We pay a price for this:

the belief that the mind is endogenously driven to spontaneously redescribe—that to some extent human knowledge, like the objective knowledge of Hegel's Overmind or Absolute, 'evolves under its own steam'. This phrase is Popper's (1972), who vigorously objects to the notion. The claim, however, is now an empirical one. Assuming that the data and analysis of human cognition are correct, cognitive scientists and machine learning theorists must now try to implement this ability.

References

Boden, M.A.: 1990, *The Creative Mind, Myths and Mechanisms*, Weidenfeld & Nicolson, London. Revised edition, 1992, Cardinal, London.

Boden, M.A.: forthcoming, Creativity and Computers, *in* Dartnall,T. H., forthcoming.

Clark, A.C. & Karmiloff-Smith, A.: forthcoming, The Cognizer's Innards, *Mind and Language*.

Clark, A.C.: forthcoming, Creativity and Cognitive Development, *in* Dartnall, T. H., forthcoming.

Cohen, H.: 1983, Catalogue. Tate Gallery. Cited in Edmonds, E., forthcoming.

Cummins, A.: 1990, The Connectionist Construction of Concepts, *in* Boden, M. A. (ed.) *The Philosophy of Artificial Intelligence*, Oxford.

Dartnall, T.H.: forthcoming, *Artificial Intelligence and Creativity*, Kluwer Academic Publishers, Dordrecht.

Dennett, D.: 1986, The Logical Geography of Computational Approaches: A View From the East Pole, *in* Brand, M. & Harnish, R. M. (eds,) *The Representation of Knowledge and Belief*, University of Arizona Press, Tucson.

Dennett, D.: 1992, Commandos of the Word, *Times Higher Education Supplement*, 27/3/92.

Dretske, F.: 1980, The Intentionality of Cognitive States, *Midwest Studies in Philosophy*, 5.

Dretske, F.: 1981, *Knowledge and the Flow of Information*, MIT/Bradford, Cambridge Mass.

Dreyfus, H.L.: 1979, *What Computers Can't Do: The Limits of Artificial Intelligence*, 2nd ed., Harper and Row, New York.

Dreyfus, H.L. & Dreyfus, S.E.: 1987, Mind Over Machine. The Power of Human Intuition and Expertise in the Era of the Computer, The Free Press, New York.

Edmonds, E.: forthcoming, Cybernetic Serendipity Revisited, *in* Dartnall,T. H. (forthcoming).

Fodor, J.: 1987, *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*, MIT Press, Cambridge, MA.

Heil, J.: 1983, *Perception and Cognition*, University of California Press, Berkeley.

Hofstadter, D.R.: 1981, A Conversation with Einstein's Brain, *in* Hofstadter,D. R. and Dennett, D.C. (eds) *The Mind's I*, Basic Books, New York.

Karmiloff-Smith, A.: 1986, From metaprocess to conscious access: evidence from children's metalinguistic and repair data, *Cognition*, **2**: 3.

Karmiloff-Smith, A.: 1990, Constraints on representational change: evidence from children's drawing, *Cognition*, 34.

Karmiloff-Smith, A.: 1991, Beyond Modularity: Innate constraints and developmental change, *in* Carey, S. and Gelman, R. (eds), *Epigenesis of the Mind: Essays in Biology and Knowledge*, Laurence Erlbaum Associates.

Karmiloff-Smith, A.: forthcoming, a, *Beyond Modularity: a developmental perspective on cognitive science*, MIT/Bradford, Cambridge Mass.

Karmiloff-Smith, A.: forthcoming, b, PDP: Preposterous Developmental Postulates? *Connection Science*.

McCorduck, P.: 1991, *Aaron's Code*, W.H. Freeman and Company, New York.

Popper, K.R.: 1972, *Objective Knowledge: An Evolutionary Approach*, Clarendon Press, Oxford.

Pylyshyn, Z.W.: 1979/81, Complexity and the Study of Artificial and Human Intelligence, *in* Ringle, M. (ed.) *Philosophical Perspectives in Artificial Intelligence*, Humanities Press, Atlantic Highlands, NJ.

Richardson, R.C.: 1981, Internal Representation: Prologue to a Theory of Intentionality, *Philosophical Topics*, 12.

Sayre, K.M.: 1986, Intentionality and Information Processing: An Alternative Model for Cognitive Science, *Behavioural and Brain Sciences*, 9.

Sharples, M.: forthcoming, Computer Support for Creativity, *in* Dartnall, T. H., forthcoming.

Simon, H.A.: 1969, *The Sciences of the Artificial*, MIT Press, Cambridge.

Smith, B.: 1985, Prologue to 'Reflection and Semantics in a Procedural Language', *in* Brachman, R. and Levesque, H. (eds), *Readings in Knowledge Representation*, Morgan Kaufmann, Los Altos, California.

Winograd, T.: 1981, What Does it Mean to Understand Language?, *in* Norman, D. (ed.), *Perspectives on Cognitive Science*, Ablex, NJ.