13

PETER CARRUTHERS

The Creative Action Theory of Creativity

Thought is normally the precursor of action. We think first, and then we act. Most models of creative thought and creative activity make a similar assumption. They assume that creative activity is preceded by, and is causally dependent upon, creative thought. This chapter, in contrast, argues for the reverse. It develops a model according to which creative thought is always preceded by, and causally dependent upon, creatively generated action schemata. And it adduces a variety of considerations in support of such a model.

1 Introduction

Creative human thought and activity present cognitive science with two distinct kinds of challenge. One is to model the creative process itself. The goal, here, is to understand how innovative ideas and hypotheses are produced. Can some combination of association, random recombination, conceptual priming, and the use of heuristics for generating novel concept combinations suffice to explain the creative aspect of creative cognition? Although important, this is not the problem that I propose to pursue in the present chapter.

The second challenge for cognitive science to address—and the one on which I will focus here—is to outline the mental architecture underlying creative thought and action. Assuming the existence of some sort of mechanism for generating novel ideas, the problem is to understand how that mechanism fits into the overall "flowchart" of the mind, interacting with other systems in such a way that new and fruitful beliefs and actions can result. What we need, in effect, is an architecture that can implement "geneplore" (for "generate and explore") models of creative cognition (Finke et al., 1992; Finke, 1995; Ward et al., 1999). We need a way for a creative idea generator to be embedded in a wider set of inferential systems in such a way that the implications of a new idea can be developed and evaluated before that idea is believed, adopted, or put into practice.

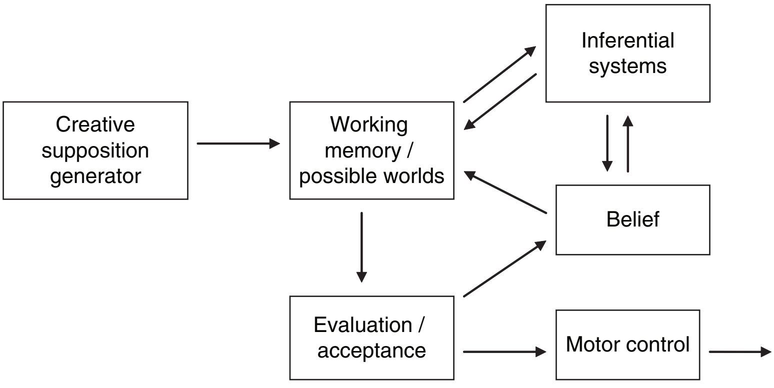

FIGURE 13.1 An architecture for geneplore.

It is widely accepted that creative cognition presupposes a capacity to entertain, and to reason with, hypothetical scenarios, or suppositions (Harris, 2000; Carruthers, 2002; Nichols and Stich, 2003). This is because the "explore" element in geneplore requires a capacity to elaborate an idea and work out its consequences in advance of that idea's being endorsed or accepted, while it is still merely hypothetical. The best-developed model of suppositional reasoning is provided by Nichols and Stich (2003), in the course of their account of the cognitive architecture underlying childhood pretend play.[1] They propose that creatively generated suppositions are held in a working memory system (which they label the "possible worlds box") where those suppositions can be elaborated. The contents of the possible worlds box can be filled out using any of the subject's existing beliefs (screened for consistency, of course—only those beliefs that are consistent with the initial supposition are allowed entry into the possible worlds box). And those contents are also available to any of the subject's inferential systems that normally operate on beliefs, producing new beliefs from old. In addition, the contents of the possible worlds box must be subjected to some sort of evaluative process which decides whether or not the initial supposition should be accepted or implemented in action (see figure 13.1).

I take it that something of this sort would be widely agreed upon. Now the question I want to ask is: What is the format of the representations created by the supposition generator? The orthodox position is that the representations are fully conceptual thoughts or propositions, such as the thought THE BANANA IS A TELEPHONE.[2] On this account, although actions as well as thoughts can be creative, any creative action is always preceded by, and grounded in, some suitably related creative thought. I shall refer to this as the "thought-first" account of creativity. It, or something like it, is assumed

1. The connection is that pretense, too, requires a capacity to entertain a supposition (e.g., that the banana is a telephone) and then to think and act within its scope.

2. I shall follow the usual practice of utilizing small capitals to represent concepts/representations in the language of thought.

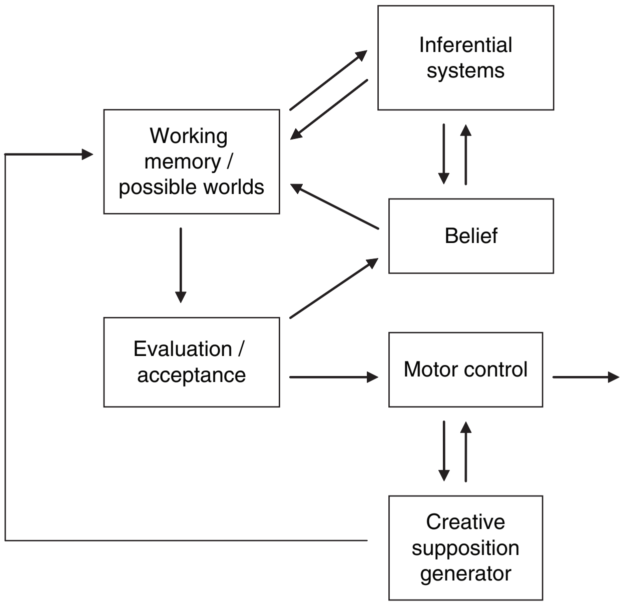

FIGURE 13.2 Action-based creative cognition.

by just about everyone who works on creativity in cognitive science. The contrasting position that I propose to explore in this chapter is that the representations produced by the supposition generator are activated and rehearsed action plans or act schemata. I shall refer to this proposition as the "act-first" account (see figure 13.2). Such an approach is apt to seem mysterious. (This explains, no doubt, why it has practically no adherents.) For doesn't thought precede action? And how could a creative action serve to generate a creative supposition or a creative thought? I shall show, however, that the act-first account is not only possible but plausible; indeed, I shall show that it has significant advantages over the standard thought-first theory.

Here is how I propose to proceed. I shall argue first (in section 2) that creative action can't be reduced to creative thought, and that at least *some* forms of creative action aren't preceded by a creative thought. I shall then briefly argue in section 3 that it is implausible that there should be two distinct and independent sources of creativity—one for action and one for thought. In sections 4 and 5 I shall show how creative thought can be explained in terms of creative action, utilizing known mechanisms including a well-established system for the mental rehearsal of action, and a cognitive architecture for global broadcasting of sensory or quasi-sensory (imagistic) states. In section 6 I shall argue that act-first accounts of creativity have evolutionary precursors, some of them quite ancient. In section 7 I shall show that thought-first accounts of creativity, in contrast, face problems of evolvability, and that they need to assume a heavy explanatory burden in comparison with the act-first account. Finally, in section 8 I shall line up some of the costs and benefits of accepting an act-first account of creativity.

2 Creative Action Without Creative Thought

Can there be creative actions that aren't preceded and/or caused by creative thoughts? Consider a jazz musician who improvises a series of variations on a musical theme.[3] Or consider a dancer who extemporizes a sequence of movements that she may never have made before (and may never make again). These are undoubtedly kinds of creativity. But they seem to be forms of creativity of *action*, rather than creativity of thought. For the novel movements appear to be made "on-line," sometimes extremely swiftly, and without prior reflection or planning—or at least without prior *conscious* reflection or planning.

Someone might pick up on this last concession to argue that jazz and dance improvisation does involve planning—only the thoughts involved occur unconsciously, immediately prior to the execution of the movements in question. Such a view is implausible, however, for a number of reasons. One has to do with the *fineness of grain* that can be present in skilled improvisation. Someone executing a novel sequence of notes on the saxophone, for example, or a novel sequence of bodily movements in a dance, doesn't *just* play those notes or make those movements. For these might, indeed, be actions that the agent has names and/or concepts for ("E flat, followed by F, followed by C flat," or "Up a fourth, down a fifth," and so on). But the agent will also choose a precise length for each note, or a precise speed for each movement, for which there is no name (and probably no concept). Likewise, the agent will add a precise timbre to the playing of the note, or a precise articulation to the movement. Although intentional, these aren't actions that can plausibly be captured fully in any sort of propositional/conceptual description.

In fact there is a strong case for saying that skilled action control has a nonconceptual (or at least an *analog*) aspect, just as perceptual contents are partly nonconceptual or analog in nature. A percept of the precise shades of red in a rose petal has a fineness of grain that escapes any conceptual description that one might attempt to impose on it and that is prior to the application of any concept (Carruthers, 2000; Kelly, 2001). Likewise, a precise movement or sequence of movements, too, has just such a fineness of grain and partially nonconceptual character. In which case skilled creative action can't be fully explained in terms of the creativity of thought. For even if there are (unconscious) conceptual thoughts that precede the action, they can by no means fully determine it; and hence there must at least be an *element* of the creativity displayed by the agent that doesn't reduce to conceptual creativity.

It might be replied that creative action can always be underlain by creative thoughts that are indexical in form. Thus a dancer's thought that precedes and explains a novel set of movements might take the form "I shall move my arms *thus* while moving my legs *so*." But what, on this account, would fix the intentional

3. In the course of his extensive discussion of jazz improvisation, Berliner (1994) outlines a number of different strategies and heuristics that jazz improvisers will adopt to guide and frame their performance. But beyond that, the particular notes and phrases that they play on any given occasion will often strike them with the force of discovery. They are often surprised by their own playing, which seems to them to have a life of its own. I shall return to this point shortly.

content of the two indexicals "thus" and "so"? Since the thought precedes the action, those indexicals can't be grounded in a *perception* of the movement in question, in the way that the indexical in a thought such as "I shall pick up *that* apple" can be grounded in an analogue/nonconceptual percept of the object seen. Thus the only remaining possibility is that the contents of the indexicals in a movement-determining thought are given imagistically. Hence, when the dancer thinks "I shall move my arms *thus*," the content of "thus" will be given by a proprioceptive or visual image of a particular set of fine-grained movements of the arms.

It is implausible that every creative action should be preceded by some such creative thought, however. For one thing, there is evidence that images of movement are themselves caused by activating the appropriate motor schemata, as we shall see in section 4. In which case why shouldn't the schemata sometimes issue in action directly, without first being used to construct an image? Moreover, consider just how *fast* creative actions can be. A jazz improviser can be playing at full speed, piecing together and recombining previously rehearsed phrases and patterns, when he suddenly finds himself playing a sequence of notes that he has never played before, and which surprises him (Berliner, 1994). For example, Charlie Parker was famous for being able to play his improvised solos at amazing speed—some of them at 400 beats per minute (Owens, 1995). Most of us would have trouble even tapping our feet to such a tempo. And even though Parker's solos were mostly composed out of arrangements and rearrangements of formulaic fragments—ranging from two- or three-note patterns to clusters of a dozen notes—it is difficult to believe that there was time in which to form a conceptually driven but fully detailed imagistic representation of each such fragment in advance of activating the motor schema for it.

Let me now return to the point noted in passing above: that jazz improvisers are often surprised by their own products. This is direct evidence in support of the view being proposed here: that actions can be creative without prior creative thought. For surprise is the emotion that we feel when something *unexpected* happens. But the expectations in question don't have to be consciously entertained. On the contrary, events can be most surprising when they violate tacit expectations that it would never have occurred to us to formulate consciously otherwise. So when a jazz improviser is surprised by the sequence of notes that he hears himself play, this is evidence that he didn't have a prior expectation (whether conscious or unconscious) that he would play just those notes. At the very least it follows that the creative thought that is alleged to have preceded the action must have occurred within some subsystem that is cut off from access to globally broadcast perceptions (in this case, of sound). But the suggestion that there exists such a subsystem has nothing to support it.

There is reason to think, then, that not all creativity reduces to the creativity of thought. At least some forms of creative activity would appear to be spontaneous, occurring in the absence of prior creative thought.

3 How Many Sources of Creativity?

I have argued that the creativity of action can't be reduced to the creativity of thought. But how plausible is it that there should be two distinct and independent sources of creativity—one for action and one for thought? Although possible in principle, any such view assumes a heavy explanatory burden. For we would need to tell two distinct evolutionary stories about the emergence of these two forms of creativity, and we would need to describe two distinct cognitive mechanisms underlying them. It is therefore preferable to explain the creativity of thought in terms of the creativity of action, if we can. Although this seems initially unpromising—indeed, mysterious, for how do new *actions* create novel *thoughts?*—I believe that it is defensible. A view of just this sort will be explained and elaborated in sections 4 and 5.

I shall suggest, in fact, that all creativity reduces to the creative generation of action schemata. Sometimes these schemata are used to bring about novel actions directly. But sometimes they are use to generate visual or other images, which are globally broadcast in the manner of perceptual states generally (Baars, 1988, 1997), and received as input by the myriad inferential and motivational systems. And in the special case where the novel action schema that gets created is a linguistic one, its mental rehearsal results in a sentence in "inner speech," which when processed by the language comprehension system, will present a new propositional thought to the various inferential systems for elaboration and further processing.

It might be objected that there is only a "heavy explanatory burden" imposed on the view that there are two or more distinct sources of creativity if we think of creativity as being some sort of *process* (such as sentence parsing). But why can't creativity be a *property* or *manner* that a variety of events and processes could instantiate? Why can't creativity be more like *stealth* or *haste?* There is no temptation to think that there must be a single system or capacity underlying haste. On the contrary, almost any activity or cognitive process can be conducted in haste, utilizing just the resources that are normally involved in that activity or process itself. Might it not be so with creativity?

It seems to me plain, however, that creativity can't be just a *manner* in which familiar events or processes are conducted, precisely because creativity involves the introduction of *novelty*. Thus it makes sense to ask at what point or points within cognition novelty can be introduced. And for each such "point" that is proposed, it looks like some sort of evolutionary explanation can be demanded. We will return to this topic in section 7, when we examine a competing thought-based account of creativity in more detail.

In addition to reducing our explanatory burden by half, the view that all creativity reduces to the creative generation of action schemata has other virtues, too, as we will see in more detail later. In particular, it enables us to envision how creativity might have evolved quite easily by adapting and utilizing mechanisms that were already in place, that evolved initially for other purposes. But first I need to outline how an act-first account of creativity might work.

4 Mental Rehearsal of Action and Global Broadcast

There are a number of components of the act-first theory of creativity, each of which is independently warranted. The first is the two-systems theory of vision (Milner and Goodale, 1995; Jacob and Jeannerod, 2003; Glover, 2004), also replicated in other

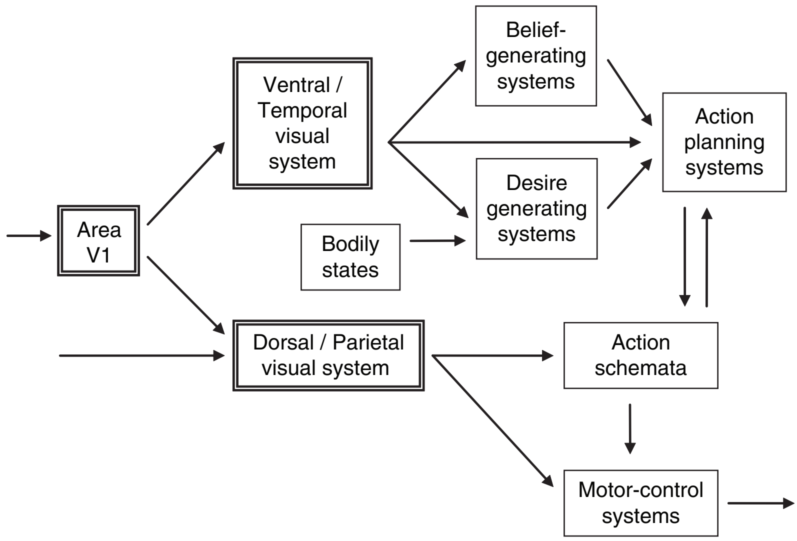

FIGURE 13.3 Two visual systems.

sense modalities (Michel and Peronnet, 1980; Paillard et al., 1983; Rossetti et al., 1995). It has long been known, of course, that the visual system contains a number of semi-independent subsystems (e.g., for color processing and for face recognition). But research since the 1980s has demonstrated that vision divides, at the highest level of analysis, into two functionally distinct systems. One of these is located ventrally in the temporal lobes, and the other is located dorsally in the parietal lobes. The ventral system is comparatively slow, uses allocentric spatial coordinates, gives rise to medium- and long-term memories, and is concerned with object recognition and planning. Its outputs are characteristically conscious, given appropriate levels of attention, and such outputs are globally broadcast to a wide range of systems for drawing inferences, for forming memories and emotions, and for practical reasoning about what to do in relation to the perceived environment. (For evidence of the global broadcasting of the outputs of the ventral visual system, see Dehaene and Naccache, 2001; Dehaene et al., 2001, 2003; Baars, 2002, 2003; Baars et al., 2003; Kreiman et al., 2003.) The dorsal system, in contrast, is fast, uses body-centered or limb-centered spatial coordinates, and has a memory window of just two seconds. It isn't involved in conceptualizing its inputs, and its outputs (which are unconscious) are used in the on-line guidance of bodily movement in relation to the perceived environment. Each of these two systems receives its primary input from the retina via area V1 at the back of the brain; the dorsal system also receives a separate stream of input via the superior colliculus in the midbrain (see figure 13.3).

As well as operating in a feed-forward manner, each of these two systems contains substantial back-projecting neural pathways, whose functions are now beginning to be well understood. In the ventral system they are used in the process of object recognition, directing attention to aspects of the incoming information and

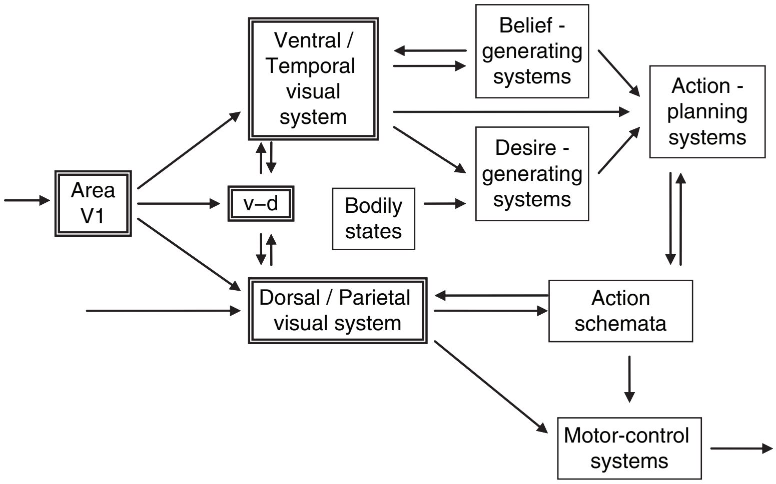

FIGURE 13.4 Two visual systems with back-projecting pathways.

also "querying" degraded or ambiguous input (Kosslyn, 1994). The querying process works somewhat like this: Candidate concepts are activated and used to project representations of their instances back down through the ventral system, where they are matched against the incoming percepts. This same system can then also be used "off-line" to generate visual imagery: An activated conceptual representation (of a horse, say) is used to create activity early in the ventral system similar to that which would occur if a horse were actually being perceived. This is then processed in the usual way, giving rise to a quasi-percept *as of* a horse (Kosslyn, 1994).

The back-projecting pathways in the dorsal system, in contrast, exist to help monitor and fine-tune the on-line guidance of action (Wolpert and Ghahramani, 2000; Wolpert and Flanagan, 2001; Wolpert et al., 2003). It works like this. Whenever a motor schema is activated, not only are commands sent to the muscles necessary to control the intended action, but "efferent copies" of those commands are at the same time created and used to generate a representation of the perceptions (not only visual, but also proprioceptive) that are to be predicted as resulting from the execution of that motor schema. (This probably requires that there should exist one or more separate "emulator systems" which take efferent copies as input and are capable of generating predictions about the likely future positions of the limbs and body, perhaps utilizing some sort of model of the kinematics of the body; Grush, 2004.) The predicted sequence is then compared with the actual sensory input received, and the detailed further execution of that action (or its replacement by another one) is determined accordingly.

Although the ventral and dorsal visual systems subserve different functions (object recognition and action guidance, respectively), it is important to realize that they are nevertheless significantly connected with one another via a region of ventrodorsal cortex. This is probably best thought of as a common functional *component* of each. (See figure 13.4, in which "v-d" stands for "ventrodorsal." The areas in question are the superior temporal sulcus and area FP in the rostral part of the inferior parietal lobule. These are strongly interconnected with each other, and also with area F5 in premotor cortex. See Rizzolatti, 2005.) This common component constitutes part of the "mirror neuron" system (Gallese et al., 1996; Rizzolatti et al., 2000), whose functions are also beginning to be well understood. As a bridge between the ventral and dorsal visual systems, the mirror neuron system is well placed to map conceptual representations of the actions of another person, categorized within the ventral system, onto corresponding motor schemata of the appropriate type, via the dorsal system's intimate connections with motor cortex. This enables imitation of the other person's actions to occur. But the mirror neuron bridge is also well placed to map one's own actual or supposed movements into the ventral system, giving rise to corresponding conceptualized and globally broadcast visual representations of those movements and their immediate consequences.

Consider, for example, the action of grasping the handle of a coffeepot, lifting the pot, and pouring a cup of coffee. An abstract action schema for the movement is activated, and rendered successively more determinate in the light of perceptual input. Motor commands are then issued, and an efferent copy is projected back through the dorsal visual system and mapped across into the ventral system to generate a prediction of the way that the intended action should look. And these perceptual images, when received by the various inferential systems that interact with the ventral system, can be further elaborated to include some of the predicted consequences of the action, too (such as precisely where the coffee will land when poured). These are matched against the incoming perceptual data as the action unfolds. If discrepancies are found (for example, the pot is heavier, and thus rises more slowly than expected, or the coffee begins to pour closer to the edge of the cup than expected), then the motor program is adjusted accordingly.

It is important to note that the ventrodorsal bridge probably also plays an important role in the deliberate *transformation* and *movement* of conscious visual images, which is driven by activity in motor cortex. Kosslyn (1994) argues on the basis of a variety of (mostly behavioral) data that motor cortex and premotor cortex are active whenever subjects transform visual images—for example, when they are rotating an imagined figure. The idea is that subjects will get their image to rotate by activating a motor schema for an act of rotation linked to the imagined figure. In effect, the idea is that we get the image to rotate by imagining ourselves *acting* in such a way as to *cause* the imagined object to rotate. This claim has been further confirmed by later research.

For example, Ganis et al. (2000) find that interfering with the activity of motor cortex via direct electrical stimulation has a significant effect on response times for people engaged in imagery rotation tasks. Similarly, Turnbull et al. (1997) report that people with lesions in ventrodorsal cortex have problems in recognizing objects that seem to require transformations of imagery, such as recognizing an object seen from an unusual perspective. And Kosslyn et al. (2001) report the results of a brain imaging study in which subjects watched an object being rotated either by hand or by a machine before undertaking a mental rotation task involving a similar sort of object. They found that primary motor cortex was active only in the first ("by hand") condition; but that premotor cortex was activated in both conditions—suggesting that motor schema planning of some sort lies at the bottom of all imagery transformation. (See also Richter et al., 2000; Lamm et al., 2001.)

Just as the back-projecting pathways in the ventral visual system can be used "off-line" to create an elaborate visual imagery, so the corresponding pathways in the dorsal system, together with the ventrodorsal bridge, can be used for the mental rehearsal of action prior to and independently of any actual movements being made. This has obvious utility, enabling us to test out and examine the likely consequences of proposed actions in advance. Here is how it works. A conceptual representation of a proposed action, produced by the practical reasoning system, is used to construct a corresponding motor schema. This is then projected back through the dorsal visual system in the form of an efferent copy, where it is transformed into visual representations of the movements involved. The latter are mapped across into the ventral system via the ventrodorsal bridge, where they are globally broadcast in the manner of a conscious perception. The various inferential systems that receive such broadcasts then set to work figuring out the likely consequences. These can also be added to the broadcast image by utilizing the back-projecting pathways in the ventral system. The resulting images are received, in turn, by the various emotional and motivational systems, which respond somewhat as they would to visual input. We then monitor our resulting bodily and/or hedonic reactions, and the desirability of the original action gets adjusted up or down as a result (Damasio, 1994; Schroeder, 2004).[4]

Consider a particular example. Looking at my monthly credit card statement, I realize that I need more money. After reviewing some options, I hit upon the idea of going to ask my boss for a raise. I mentally rehearse the action of walking into his office and broaching the question of salary. The resulting images are globally broadcast, and are elaborated to include my boss's likely response (the glowering face, the harsh words). The result is that I feel fear and disappointment. And that leads me to abandon any thought of asking for a raise, and returns me to considering other options. We spend much of our waking lives, as adults, in mental rehearsals of this sort, often to good effect. Initially promising plans can turn out to be disastrous when rehearsed; and plans whose success at first seems implausible can turn out to be much more likely to succeed.

5 The Act-First Account of Creativity

There is little doubt that the mental rehearsal of action takes place pretty much as I have just described. And the capacity for such rehearsals, and the resulting cycles of globally broadcast representations of proposed actions, may very well be common to other primates as well as to human beings (Carruthers, 2006, chap. 2). Moreover,

4. There are significant differences between these two authors that aren't germane to my purposes here. Damasio believes that imagined actions give rise to emotional reactions which in turn cause bodily changes that we monitor using the somatosensory systems. Schroeder believes that the imagined actions give rise to unconscious forms of punishment or reward, which are represented and monitored in the frontal lobes in the form of pain or pleasure.

because of the obvious utility of such action rehearsals, the whole arrangement may have evolved over significant periods of time. But there is nothing here, as yet, that requires creativity. The starting point for a mental rehearsal of action can be something that the agent is about to do or is considering doing, arrived at via normal, uncreative processes of practical reason.

Suppose, however, that a capacity for creative generation and activation of action schemata were to be added into the mix. (I shall treat this capacity as a "black box" for the moment, returning to consider how it might work, and to discuss some of its precursors, in section 6.) Then everything would be in place for an act-first account of creative cognition to operate. A creatively assembled action schema is activated and rehearsed, giving rise to an imagistic representation of the action in question. The latter is then broadcast to the various inferential, belief-generating, and motivational systems, which further elaborate and evaluate it.

Recall that the main puzzle about the act-first account of creativity is how *action* can give rise to *thought*. This question is now readily answered. For the mental rehearsal of an action schema will give rise to an imagistic thought representing the action in question, occurring at a point in the overall architecture of the mind where it can be further elaborated by inference, give rise to emotional reactions, and enter into our practical reasoning (see figure 13.4). Moreover, in the special case where the action schema in question is a *speech* action schema, its mental rehearsal will give rise to an imagistic representation of the corresponding utterance in "inner speech," which, when received and processed by the language comprehension system, will result in the global broadcast of a propositional thought.[5] The interpreted propositional content of the utterance will be broadcast alongside the imagined sounds, just as happens when we hear another person speak. So this will be a thought that is caused by a creatively generated and rehearsed action schema.

Compare figure 13.4 with figure 13.2, which we used to represent an act-first version of the geneplore model of creativity. The working memory system/possible worlds box can be identified with the iterated global broadcast of perceptual images, elaborated in the light of interactions between those images and the subject's beliefs, and utilizing any inferential resources that are normally available to process perceptual input. And the evaluative system can (in part) be identified with the responses of the emotional and motivational systems when they receive the globally broadcast images as input. (A rather more nuanced story needs to be told when what is at issue isn't the evaluation of a proposed action, but rather the evaluation of a rehearsed thought as true or false. See Carruthers, 2006, chap. 6.) This evaluative system is depicted in figure 13.5 (using the version defended in Damasio, 1994, rather than that in Schroeder, 2004).

5. Recall that a bifurcation of functions similar to that of the two visual systems exists in other sense modalities, including hearing. And there is evidence that both the language production and language comprehension areas of the cortex are active during inner speech. (See Paulescu et al., 1993; Shergill et al., 2002).

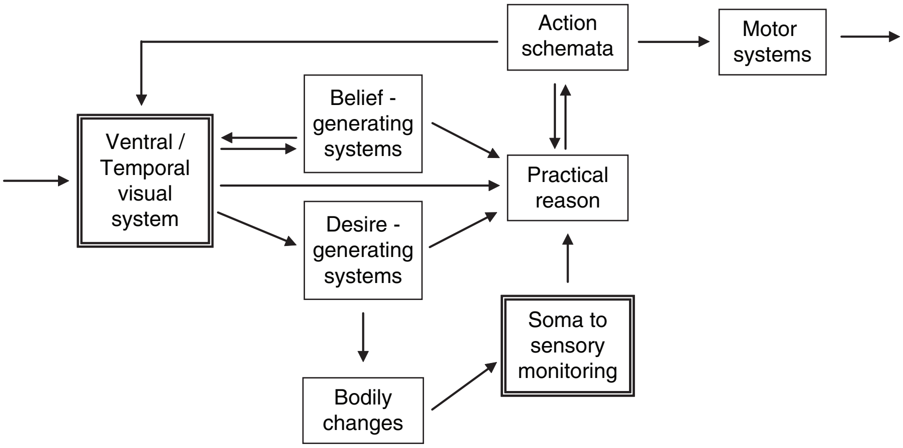

FIGURE 13.5 Mental rehearsal and somatosensory monitoring.

6 Precursors: Ancestral Forms of Creative Action

It appears that an act-first account of creativity will be well warranted, provided that action schemata can be creatively activated independently of any prior creative thought. In section 2, I argued just this, by focusing on creative activity in jazz and dance. But it is also worth noting that simple kinds of creative action (which almost certainly aren't guided by creative thought) are actually quite widespread in the animal kingdom, in the form of "protean" erratic behavior (Driver and Humphries, 1988; Miller, 1997). Let me briefly elaborate. When a moth is hit by bat ultrasound, for example (signaling a predator's approach), it will start to loop and tumble in a fashion that seems genuinely random and is wholly unpredictable; this is a much more effective evasion technique than mere passive tumbling or a predictable (but faster) straight flight away (Roeder and Treat, 1961; Roeder, 1962; May, 1991). Such randomized escape behaviors are extremely common in the animal kingdom, and for good reason. For the best way to make your behavior unpredictable to a predator is to somehow make it genuinely unpredictable, period. It was for just this reason that submarine commanders in World War II would throw dice to determine the timings and directions of their zigzag travel paths, to make themselves unpredictable to submarine-hunting surface vessels.

There is another important factor in the generation of actions among mammals and birds that should be mentioned here, which Gallistel (1980) calls "the principle of autonomous buildup of action-specific potentiation." This, too, seems to give rise to simple forms of creative action generation without prior creative thought in some species. Consider actions that are normally performed only as components of a larger action schema and in the service of another goal, such as running (in rats) or pouncing (in cats). Normally rats run when exploring a novel environment or in search of food; and cats pounce in the course of predatory behavior. But when these sorts of actions have not been performed for some time, there is a buildup of an intrinsic disposition to do so—in effect, creating a novel intrinsic desire.

Thus rats can, of course, learn to work by running on a wheel in order to obtain water or food. But equally they will learn, if they are prevented from running, to drink in order to run. And a cat that hasn't pounced for some time will work with considerable resourcefulness at patterns of action that finally yield something to pounce upon. Likewise, Lorenz (1950) describes the case of a starling that had been caged for some weeks, and had thus been unable to hunt. When released from its cage it exhibited the entire repertoire of actions in its insect-hunting behavior. But there were no insects present, and the bird itself was fully fed.

The result is that many animals—and especially monkeys and apes—will spend significant amounts of time engaged in seemingly aimless mixing of behavioral fragments. We often describe this as "play," although it is common in adult animals as well as in infants. One function might be to keep the actions themselves fine-tuned and efficiently performed, in readiness for their co-option into the service of a biologically important goal. But Schiller (1957) suggests another function, based on his analysis of the problem-solving abilities of the apes famously studied by Köhler (1927). This is that the jumbling of action components can create novel combinations, which can sometimes generate a reward of some sort. The latter can then stabilize the new combination in the animal's behavioral repertoire.

It seems likely, then, that a capacity for the creative generation of actions independently of prior creative thought would have been part of our animal inheritance. And when combined with a capacity for mental rehearsal of action, this would have resulted in at least a limited capacity for creative thinking, in the manner outlined in sections 4 and 5. The stone toolmaking abilities of earlier species of hominids reveals the existence of just such a capacity, I believe, as I shall now briefly explain.

A number of authors have stressed the cognitive difficulties involved in making the symmetrical hand axes and blades that were being produced by members of *Homo ergaster* from about 1.4 million years ago, probably long before the evolution of language (Gowlett, 1984; Pelegrin, 1993; Mithen, 1996, 2002; Schlanger, 1996; T. Wynn, 2000). Many of these items possess a fine three-dimensional symmetry that was plainly intended (T. Wynn, 2000). And it is often the case that hand axes from the same assemblage or from the same region conform to a similar pattern. So it is evident that their makers started out with a clear idea of the intended product, and that in some sense planning was involved. Moreover, we now know quite a lot about how hand axes were made, both from the testimony of contemporary knappers who have succeeded in reproducing them (Pelegrin, 1993), and from processes of painstaking reconstruction in those rare instances where a completed hand ax has been found together with the waste flakes resulting from its manufacture (Schlanger, 1996).

What we know is that it is impossible to produce such an artifact by reasoning purely "analytically," without creative thought. Even if the whole production process is well practiced and familiar, there is generally no way to work back from the desired finished product to what one should do first (nor to what one should do at many of the intermediate stages, either). For the stone cores from which the production process starts are always (to some degree) unique in shape, size, and the details of their material composition; hence each core presents a unique challenge. Even an experienced knapper must pause at a number of different stages in the production process to visualize the next goal (such as the preparation of a small ledge, or "striking platform," which can be used to dislodge a larger flake from the core), and he has to try out in imagination various potential strikes, or sequences of strikes, that might achieve that goal.

In effect, the knapper must continually be entertaining thoughts of the form "Suppose I struck it *there* like *that*; or suppose I struck it *here* instead, like *this*."[6] He will thus be rehearsing a number of different action schemata and monitoring the predicted results, sometimes then rotating his visual image of the resulting core to see what it would look like from the other side (T. Wynn, 2000). So we can conclude, then, that at least some limited capacity for creative thought would have been present in the minds of our hominid ancestors prior to the evolution of *Homo sapiens*. Specifically, our ancestors had the capacity to "try out" a number of different action schemata in visual (and other forms of) imagination, hence making those supposed actions available as input to the full range of belief-generating, desire-generating, and action-selecting systems, and thus recruiting the activity of those systems into the service of the intended goal.

With limited, task-specific forms of creative mental rehearsal in place, what then had to happen was for humans to start generating such rehearsals much more widely and more often? I have argued elsewhere that this may be the proper function of childhood pretend play (Carruthers, 2002, 2006). By being disposed to generate and explore creative suppositions, in pretense children will both strengthen their disposition to think creatively and begin to develop a set of heuristics for selecting creative combinations of action schemata for rehearsal. But the main point to emphasize here is that with all the groundwork already prepared (especially a basic capacity for creative action generation together with capacities for mental rehearsal of action), it would have required only some relatively minor changes for full-blown creative human thought to make its appearance.

7 Against the Opposition: Problems for a Thought-First Account

The act-first account appeals to processes that we already have reason to believe in, then. There are good reasons to think that perceptual and quasi-perceptual (imagistic) states are globally broadcast to a wide range of inferential systems for forming memories, for creating new beliefs and emotions, and for practical reasoning (Baars, 1988, 1997, 2002, 2003; Dehaene and Naccache, 2001; Dehaene et al., 2001, 2003; Baars et al., 2003; Kreiman et al., 2003). And there is good reason to think that motor schemata can be used to create and transform such visual images (Kosslyn, 1994;

6. Let me stress that I don't mean to imply that such thoughts must be entertained in natural language. On the contrary, I believe that creative thought is possible in the absence of language through image-creating rehearsal of action schemata, even if it is greatly enhanced and extended by the presence of language. Rather, since language is unlikely to have been present among *Homo ergaster*, the thoughts in question will have been realized in mental rehearsals of actions from the agent's repertoire.

Turnbull et al., 1997; Ganis et al., 2000; Richter et al., 2000; Kosslyn et al., 2001; Lamm et al., 2001). And there are also good reasons, I claim, for thinking that a limited capacity to reason with suppositions in the form of mental rehearsals of potential actions might have long predated the evolution of human beings. What humans have in addition is a disposition to generate such rehearsals creatively. There would therefore have needed to be only a *small* evolutionary benefit accruing from creativity in order for the novel disposition to generate suppositions that aren't so directly related to the actual environment to emerge. For all of the basic connections and systems would already have been in place in our great ape/hominid ancestors.

A thought-first account, in contrast, can't similarly build on the existence of known systems and processes. The conceptual supposition generator would somehow have to be built ab initio, as would a specialized propositional working memory system (a "possible worlds box"). And input and output connections would need to be constructed from the possible worlds box to each of the belief-generating and desire-generating systems with which it is to interact. For there is no reason to think that any of this would already have been in existence, waiting to be co-opted into the service of creative thinking when the latter began to make its appearance. This has two implications for the evolution of our capacity for supposition in general, and of our capacity for creativity in particular. One is that there would then need to have been some very significant selection pressure at work. For evolution would need to build a whole new cognitive system, with myriad input and output links to other systems. And the other implication is that it therefore becomes difficult to tell a sensible story about piecemeal evolvability. Does the supposition generator get connected to other inferential systems one by one, for example? If so, how?

Notice, too, that it isn't enough just to build connections between the possible worlds box and other systems, since there would also have to be corresponding adaptations within each of the inferential systems to which the possible worlds box feeds its representations. The adaptation would be: When you receive input from the possible worlds box, you *only* pass your output back to that box, not onward as a belief to any other system with which you may be connected; or if you do pass it on, you pass it on tagged so as not to be believed. Something like this is necessary to explain how inferential systems that get to work on a supposition don't issue in beliefs. This makes it even more difficult to see how the possible worlds box could have evolved. Granted, it is of considerable use to be able to engage in conditional and suppositional reasoning. But how could such a capacity ever get started on the model under consideration here? For in the absence of the adaptations to each of the receiver systems, serving to insulate suppositions from issuing in belief and action, the creation of a possible worlds box would have been disastrous. The only option I can see would be that the possible worlds box evolved in conjunction with the required corresponding adaptation in its consumer systems one at a time. That box first emerged linked to just one inferential system, and that system also happened to be altered in such a way as to pass on any subsequent output tagged so as not to be believed. And then this adaptation was copied into each of the other inferential systems as new input connections were built to them.

The act-first account sketched above, in contrast, builds upon the cognitive architecture according to which the mind already contains a capacity for *action* rehearsal. Motor schemata can be activated in suppositional mode for the purpose of testing the consequences of actions. Here I think one can get a handle on how such a system might evolve, because one can see the significance and usefulness of adapting just a single inference system in such a way that, when an act rehearsal is received as input, the output of the inference is limited in its effects. Consider, for example, imagining that I strike one stone with another, or imagining that I make a specific type of movement within someone else's sight. And with that system in place, it would have required only small changes to begin the *creative* generation and rehearsal of action schemata.

By the same token, the act-first account can help us to bridge the divide between ourselves and other animals. For virtually all of the systems implicated in action-based creativity would have long preexisted human beings. This is true for some of the processes that generate actions creatively, for the mental rehearsal of action schemata (generating imagery of the actions being undertaken or under consideration), and for the global broadcast of some of the images so generated. Granted, humans are, in many ways, rather special. But it is surely a huge plus in favor of a theory if it can explain (as the act-first theory of creativity surely can) how that specialness emerged against a backdrop of animal capacities, which collectively provided most of the main elements of the ensuing human-specific system.

8 Concluding Thoughts: Benefits and Costs

I have argued that an act-first account of creative cognition has much to recommend it. In particular, it provides us with a plausible account of the way in which the geneplore model of creativity is implemented in the mind, for the most part utilizing systems and capacities that we already have good reason to believe in. At the very least, the act-first account deserves to be taken seriously by cognitive scientists, who should begin exploring it and testing its implications.

Honesty requires us to note that there are also costs attending the act-first account, however. The main such cost is that standard models of speech production will need to be significantly modified. According to these models, a speech action always begins with the formulation of a thought-to-be-communicated, whereupon lexical, syntactic, and phonological resources are recruited and assembled in such a way as to express that thought in an utterance (Levelt, 1989). But if creative action is prior to creative thought, then creative sentence generation will somehow have to be autonomous, not starting from a preexisting thought. The claim will have to be that action schemata for items of speech can be assembled in the absence of any prior thought content for them to encode, but for purposes of supposition. We can *try out* saying things, either out loud or to ourselves in inner speech, using various heuristics for the generation of such sentences, without previously entertaining in thought the contents of the things that we say. New contents are thereby created which might go well beyond anything that could ever have been produced as the output of our various inferential and belief-forming systems, whether singly or in combination.

How implausible is it that standard models of speech production should be modified in some such way as this? Well, everyone agrees that speech is a kind of action. And when we reflect on some of the ways in which actions in general can be generated, I believe we can see some reasons to expect that the standard model should be overturned. For processes of action production will generally proceed from an abstract—often highly schematic—representation of a desired behavior, through progressively more concrete and fine-grained implementations of that schema (guided partly by details of the context), until issuing finally in a fully detailed action. Consider a nonlinguistic example first, by way of illustration.

Suppose that I am thirsty, and that I form the intention of having a drink. Then, depending on the circumstances in which I find myself and on what there is to hand (and perhaps on my knowledge of what there is in the fridge), I might form the plan of carrying *this* empty glass to the kitchen to get water from the faucet. I then have to pick the glass up with one hand or the other, and with one sort of grip or another; I have to choose a precise route to the kitchen; and I have to reach to turn the faucet with one hand while holding out the glass with the other, monitoring the level of the water as it rises. And so on. At each stage choices among more detailed motor schemata have to be made, often influenced by their relative ease of implementation (e.g., there is already an empty glass on the desk before me), but sometimes selected at random.

Likewise, I suggest, when speech is recruited to the service of some goal, such as strengthening a friendship. On seeing an old friend again, I might form the intention of saying something flattering. This is a highly nonspecific utterance schema, which then needs to be made progressively more precise, influenced by features of the context (does her hair look recently done? do her clothes look new? and so forth), as well as by background knowledge of my friend's beliefs and values, and perhaps also utilizing mental rehearsal of some my options to pull in the inferential resources of my mind-reading capacity. Here, too, at each stage choices need to be made, from the general (should I comment on her clothes or her seeming youthfulness?) to the precise (should I use an active or a passive sentence? should I use this word or that?), partly influenced by factors similar to those noted earlier. But language production is also constrained and partly guided by considerations of *relevance*, in the technical sense of Sperber and Wilson (1995). That is to say, speakers have the standing goal of achieving significant cognitive effects in the other, while at the same time minimizing the processing effort required from the audience for those effects to be achieved.

This way of seeing speech as a form of action suggests a perspective on sentence production that is even farther away from the standard model (Levelt, 1989), which begins with a conceptual representation of the message-to-be-communicated. Often, no doubt, speech production is like this when one's goal is simply testimony (telling someone that P, for some particular *P*). But often it isn't. Often the starting point is a specification of a kind of utterance (such as saying something flattering), which then needs to be made more precise. In such cases the content communicated might come into existence only near the end point of the process of speech production, rather than at the beginning. And even many cases of normal conversation that *look* like they might fit the standard model (I tell someone what I did on the weekend; she tells me what she plans to do tomorrow evening) probably don't really do so. For the actual information exchanged is often incidental to the activity, which is really governed by such goals as keeping a conversation going (finding *something* to say), being pleasant to an acquaintance, or whatever.

Of course there are large issues here for cognitive science to address.[7] My intention has just been to acknowledge an implication of the act-first account of creativity (namely, that when mental rehearsal of a creative *speech* act is used to produce a novel thought, the language production process in question cannot conform to the standard model), and to draw some of the sting from that acknowledgment. For when we remind ourselves that speech actions are *actions* which can be undertaken for all sorts of different purposes, perhaps we should be inclined to reject or modify the standard model in any case.

The bottom line, however, is this: The act-first account of creative cognition has many strengths, and there are many respects in which it has the advantage over the standard thought-first account. It deserves to be taken seriously.

Author's Note

The ideas in this paper are drawn from Carruthers (2006, chapter 5), but they have been modified in such a way as to render them independent of the massive modularism that is the main topic of that book.

7. See Baker (this volume, chap. 13) for discussion of the difficulties that stand in the way of an adequate explanation of the creative aspect of language use. My hope is that the framework presented in the present chapter might make those difficulties significantly more tractable.