cept of a rule (let alone a metarule) is only applicable as an emergent property of the interactions of the system.

One criticism of EURISKO is that, although it did exhibit interesting behaviour in a number of domains, within elementary number theory it was unable to go any further than its predecessor, AM. But the whole point of EURISKO was to help AM out of ruts by synthesising new heuristics. However, it seems that applying metarules to a set of rules does not necessarily significantly alter a system's behaviour.

We can get some feel for why this might be so from the following simple examples. Consider the metarule M1 which states when a certain rule is applicable:

M1: if X then use rule R1

and

R1: if A then B

These are equivalent to the following single rule:

if X and A then B

Similarly, consider a metarule for suggesting the negation of a rule under certain conditions:

M2: if Y then negate rule R2

and

R2: if C then D

These are again equivalent to a single rule:

if Y and C then NOT(D)

At best, metarules might provide templates for expressing a number of rules in a concise form.

The structure of COPYCAT is entirely different from this. Analogies are generated through the interactions of low-level structures, without any central control. Any patterns and regularities are emergent phenomena, not built explicitly into the system. This gives COPYCAT a great deal of flexibility in its behaviour, seen from above, arising from stochastic activity below.

Genetic algorithms also exhibit this property, particularly when genes get to cooperate with each other as well as compete (Forrest & Miller 1991). Again, structure emerges as the system evolves. Incidentally, Boden mistakenly attributes the power of GAs to the mutation operator, thinking that this provides an example of the effectiveness of randomness in creative processes (p. 16 of the target article). In fact, mutation plays only a background role in a GA and is generally kept to the minimum required to ensure that no "allele" disappears for too long. The real power of a GA is in reproduction, in which the better schemata are optimally favoured. The crossover and inversion operators explore the space of schemata without compromising this strategy (Holland 1975).

The concept of emergent systems ties in well with Minsky's "Society of Mind" idea, in which thinking arises through the interactions of large numbers of agents (Minsky 1985). A number of emergent models based on these ideas (particularly the "k-line" theory of memory) are described in (Partridge & Rowe, in press; Rowe 1991). One such model performs a best-first search in a bottom-up, stochastic manner. When a point is reached in the search where an alternative line of attack would be more fruitful than the current one, the search switches focus in a way reminiscent of "inspiration." The background stochastic activity that is not in the current focus may be considered "incubation."

In conclusion, I find Boden's presentation superficially accurate but theoretically misleading. An emergent system approach seems far more hopeful than the application of metarules.

Respecting the phenomenology of human creativity

Victor A. Shames and John F. Kihlstrom

Department of Psychology, University of Arizona, Tucson, AZ 85721; victor@ccit.arizona.edu; kihsltrm@ccit.arizona.edu

We can agree with Boden that creative products are the outputs of cognitive processes that, in principle, can be modelled in computer simulations. The success of such models, however, depends on the extent to which they incorporate the fundamental elements comprising creative processes. It seems that the particular programs favored by Boden exclude some essential elements because of specific limitations in the way researchers who design and advocate these programs construe human creative thought.

For example, Boden dismisses the accounts provided by such individuals as Mozart, Coleridge, and Poincaré as being unscientific, and she warns that "introspective accounts of creative episodes cannot be taken at face value" (p. 243). We do not wish to argue about the validity or the limitations of introspective data; both are widely recognized. What we wish to emphasize is that all introspective reports deserve equal scrutiny – it is not fair simply to criticize those that fail to conform to our own theories. Boden accepts the introspections of the pointillist painter Seurat as scientifically valid, because he writes, "I apply my method, and that is all there is to it" (p. 249). Is there any empirical reason to conclude that this account is more accurate than those of individuals who claim that their ideas appear "suddenly, unheralded by previous consciousness" (p. 243)?

Recent reviews of the creativity literature include a number of studies in which experimental evidence is obtained for the kinds of experiences that Boden dismisses as "romantic and inspirational" (p. 240). For example, Shames and P. G. Bowers (1992) have shown that the psychological constructs of effortlessness, absorption, and imaginative involvement, all of which correspond to experiences described in the anecdotal literature, are concomitants of creativity as measured in the laboratory. Dorfman et al. (in preparation) examine three fascinating components of the creative process that have escaped current computer simulations. These are intuition, incubation, and insight – the phenomenological hallmarks of human creative thought.

The creative thinker, faced with the problem of coming up with something new, begins with an intuition that this is possible: that an image or a phrase – even one like "The deckchairs are on the top of the mountain, three miles from the artificial flowers" (Boden précis, Ch. 3, para. 12) – can be turned into a line of poetry, or that a melody or rhythm can serve as the basis for a symphonic movement. The end point of the creative process seems to be marked by the experience of insight – the affect-laden recognition that the solution is in hand, that the intuition works, and that everything to follow is simply a technical matter of working out the details. Intuition and insight may be linked by incubation – the progress we make towards a solution even when not consciously working on the problem at hand. In our view, these are not mere epiphenomena of a mechanical process grinding its way from start to finish: they are important elements of mental life that must be addressed in any psychological theory of creativity.

All three of these phenomena have been studied in the laboratory. K. S. Bowers and his colleagues (Bowers et al. 1990) have shown that subjects are able to respond discriminatively to coherence on a word problem and a gestalt closure task even when they are unable to identify the coherence consciously. Metcalfe and Wiebe (1987) found that problem-solving processes that involve insight can be distinguished from those that do not based on the pattern of metacognitions associated with each process; subjects who correctly solve an insight problem are unable to detect an impending solution until it actually occurs. Smith and Blankenship (1991) observed incubation effects when misleading or inappropriate information interfered with subjects' initial attempts to solve a problem. Yaniv and Meyer (1987) observed incubation when it was assessed by implicit rather than explicit measures of a solution's accessibility.

In an integrated model of human creativity, all three of these phenomena play an important role in the production of something improbable or impossible. Intuition provides motivation for creative problem solving. The problem of coming up with something new may have been set by external forces – a commission from a patron, or the need to earn a decent living by one's craft. But the intuition that a solution is possible creates the intrinsic motivation for the act of creation: artists cannot respond to the extrinsic demand for creative products unless they believe they can follow their ideas through to the end. Since many creative individuals do their work in the absence of any prospect or promise of external reward, perhaps it is intuition that keeps them going.

While the study of intuition is relatively new, research on incubation has a long history. There are three basic points of view: one is that incubation reflects continued problem solving, outside of awareness, during the period in which the problem solver is ostensibly inactive. Another view is that a period of genuine inactivity allows the problem solver to begin anew at another entry point that is more likely to lead to an acceptable solution. A third view is that the mere passage of time permits exposure to a larger number of environmental cues, one of which leads to a solution. These are all information-processing theories that could be modelled in a computer simulation. But the success of the simulation may not tell us which of these hypotheses is the correct one. For that purpose, we need empirical data from the psychological laboratory.

If intuition provides the intrinsic motivation for creativity, insight provides the reward for following one's intuitions to the end: the affective experience of satisfaction (Aha! Eureka!) in a job well done. Insight initiates the verification process, and in divergent thinking situations where the question of creativity is relevant the intensity of the affective experience may be a clue to the quality or value of the solution. Thus, insight provides feedback about our intuitions: because we were able to follow our intuitions successfully this time, and get such a charge out of it, we may be more likely to trust them in the future.

We have seen that intuition, incubation, and insight are observable phenomena that, in all likelihood, serve an important purpose in creative thought. A meaningful psychological theory of human creativity must address these phenomena and explore the relationships among emotional, motivational, and cognitive aspects of the creative process. Our impression is that Boden too readily dismisses phenomena that she is unable to fully explain, and that the computational models she tends to favor have little to say about incubation, intuition, or insight. In evaluating the role of "felt experience" in creativity, Boden concludes, "It's not obvious that consciousness in *that* sense is essential to creativity" (p. 280). Although she may be right, it would be premature to exclude potentially relevant variables from our models of creativity. In fact, the exclusion of such variables may guarantee that our models will fail.

Individual differences, developmental changes, and social context

Dean Keith Simonton

Department of Psychology, University of California, Davis, Davis CA 95616-8686; dksimonton@ucdavis.edu

Once I had the rare opportunity to chat one-on-one with Herbert Simon. I was curious how his theories of human problem solving and artificial intelligence might explain the huge individual differences we observe in creative output (Simonton 1984b, Ch. 5). For example, why is the distribution of creative productivity so skewed that around 10% of the contributors to any field typically account for half of everything published, patented, exhibited, or performed? I must stress that this question was not as off-the-wall as might appear: earlier in his career Simon (1955) had offered a mathematical model to explicate the distinctive probability distribution of total productivity. Justified or not, I was frankly disappointed with his response. He could not make a direct connection between the cognitive processes that produce single discoveries and the processes that are responsible for a career of discoveries. Indeed, he only ventured the raw speculation that the highly skewed distribution of lifetime output may merely reflect underlying individual differences in *motivation*! In making such a remark, was Simon admitting that one of the most characteristic features of a creative life was not a proper subject for cognitive science?

This anecdote reflects the implicit Zeitgeist that guides the work of most cognitive scientists. As Gardner (1987, p. 6) expressed the prevailing attitude, cognitive science mostly ignores "the influence of affective factors or emotions, the contribution of historical and cultural factors, and the role of background context in which particular actions or thoughts occur." This neglect is immediately apparent in many of the recent books that treat cognitive approaches to creativity (e.g., Finke et al. 1992; Hayes 1989; Perkins 1981; Weisberg 1993). This same narrow emphasis on cognition to the exclusion of all else is equally apparent in Boden's brilliant exposition. Her book's main title is well chosen; she is only interested in the creative *mind*.

In my view, this constricted focus is unfortunate. We may never attain a comprehensive scientific account of creativity unless we try to handle the key features of the phenomenon. Besides describing the mental processes that produce creative ideas, the ideal theory must address three sets of questions:

1. Why are there substantial individual differences in manifest creative powers? Why do some creative workers come up with far more novel ideas than do their colleagues? How are these differences related to contrasts in intelligence, cognitive style, motivation, or values? What is the role of individual variation in genetically-based traits?

2. How does potential and actual creativity change over the life span? What developmental experiences in childhood or adolescence most strongly determine the quality and quantity of creativity observed during adulthood? Why does creative output exhibit such a distinctive trajectory over the course of a career? Why does productivity peak early in life for creators in some disciplines but peak late in life for those in other disciplines? Is there a reason why creative careers tend to consist of an interrelated "network of enterprises" (Gruber 1989) guided by some overriding thematic preoccupation?

3. What is the impact of the cultural, social, and even political environment on the origination and expression of new ideas? How do collaborators, colleagues, and rivals contribute to the emergence of creative output? In what ways does the ideological milieu shape the production and reception of discoveries, inventions, and other novelties? To what extent does the phenomenon of multiples – where two or more scientists come up with the same idea independently – imply that science is not produced by scientists at all, but rather by the scientific community acting in concert?

At this point, Boden and other students of artificial intelligence and problem solving might ask: Why should we pay any attention to these issues? Why can't we focus on the creative mind and leave these extraneous matters to personality, developmental, and social psychologists? There are actually multiple reasons, but here I have space to present just two (see also Csikszentmihalyi 1988a; Sternberg 1989; Tweney 1990).

First, a fuller appreciation of the phenomenon may impose