Proc. Natl. Acad. Sci. USA
Vol. 80, pp. 4569–4571, July 1983
Psychology

Discovery, invention, and development: Human creative thinking*

HERBERT A. SIMON

Department of Psychology, Carnegie-Mellon University, Pittsburgh, Pennsylvania 15213

Contributed by Herbert A. Simon, August 24, 1982

The subject of this paper is a singularly appropriate one for this journal. We are all experts on the topic of scientific discovery, invention, and development. Members are elected to the National Academy of Sciences because they have been adjudged to have done some creative thinking that led to one or more significant discoveries, inventions, or applications of scientific ideas.

But the fact that a person sometimes does creative things does not mean that he understands the creative process. To imagine that a scientist can give a full scientific account of his own thought processes is no more reasonable than putting a Geiger counter on the podium at an American Physical Society meeting and expecting it to deliver a lecture on the theory of radiation. In building a theory of scientific discovery or invention, a scientist's behaviors are the data—the phenomena we seek to explain—not the direct source of theory.

Much of what goes on when we are thinking is inaccessible to our conscious awareness. We use such terms as "judgement," "intuition," and "creative insight" to name and label those phenomena that occur without awareness. But labels are not explanations; contrary to the view of Moliere's physician, opium's sleep-inducing power is not explained by attributing to it a dormitive property.

Many scientists and philosophers have doubted whether a scientific explanation of discovery is possible or whether there is even something describable as a method of discovery. Einstein himself is quoted as saying, "There is no logical path leading to [scientific] laws. They can only be reached by intuition, based upon something like an intellectual love [*Einfuehlung*] of the objects of experience." If Einstein meant by "logical path" a deductive route to discovery, he was surely right. If he meant to say that there is not method in the madness of discovery, we might question—and investigate empirically—whether his pessimism was justified.

The fact that many books and articles have been written on the subjects of invention, discovery, and creativity, suggests that we believe that at least some significant aspects of these processes are amenable to investigation and description. Until recent times, most of our knowledge about them has derived from the experience and observations of thoughtful practitioners. In this case, the Geiger counter does indeed have something to say about the theory of radiation. But creative processes are also now the objects of direct scientific investigation. In the past dozen years, some psychologists have sought to build and test theories about the processes, both conscious and unconscious, that go on in the human brain when discovery and invention are taking place. Here, I am going to say a little about this psychological research, but mainly, I am going to consider what we know about these matters from the information that is available to us more directly in our practice of science and engineering.

Processes of scientific discovery

Much can be said about the differences in process at the two ends of the continuum from basic discoveries to development and application. In particular, the development of products from basic discoveries takes place in a complex social and economic environment, in which both motivations and definitions of the problem differ very much from those in the environments where basic discoveries typically occur. On the other side, I am going to emphasize the communalities of process—and to argue, in fact, that while discovery and development usually address different substantive *problems*, the psychological *processes* of problem solution are quite similar.

Let me begin by defining creativity in a pragmatic way. From time to time, human beings arrive at ideas that are judged by their fellows to be both novel and valuable. The values discerned in these new ideas may be intellectual, esthetic, practical, or what not. It does not matter. A psychological theory of creativity (or discovery or invention) would account for the processes that are *involved* in bringing about such novel and valuable products. Of course, we should not beg the question: whether creativity is a unitary thing or whether quite different processes are involved in creative production in different domains is itself an empirical question, to be decided on the basis of evidence.

I will not try to survey the whole literature on creativity. It draws partly on historical, biographical, and autobiographical accounts of discovery, on systematic surveys of creative (and uncreative) persons, on a limited number of laboratory studies aimed at eliciting creative behavior or comparing creative with noncreative styles of problem solving, and a number of recent attempts to simulate creative behavior with computer programs. At the phenomenological level, the level of observable events, there is great consistency among all the findings of this research, whatever its methods. I can sum up these findings in a few generalizations.

What chiefly distinguishes creative thinking from more mundane forms are (*i*) willingness to accept vaguely defined problem statements and gradually to structure them, (*ii*) continuing preoccupation with problems over considerable periods of time, and (*iii*) extensive background knowledge in relevant and potentially relevant areas. Not all of these conditions—tolerance of ambiguity, persistence, and knowledge—are satisfied in all cases of discovery, but their presence has been observed and commented on too many times to suppose that their association with success in discovery is accidental. None of these conditions is very surprising; they all have a strong motivational component, and they satisfy our sense of justice—the virtues of patience, persistence, and diligence are likely to be rewarded, even though there may be, as in gold mining, a large chance element in who wins the reward. The virtues only allow you to buy a ticket in the lottery, and some tickets, as in all lotteries, pay off and some do not.

Insight—based on recognition

However, something seems to be left out of this picture. It seems

* Presented at the annual meeting of the National Academy of Sciences, Apr. 26, 1982, Washington, D.C.

to deny the more romantic view of discovery that speaks of "insight" and "intuition" and the "creative moment." In an opinion he once wrote, Justice Learned Hand introduced into our patent law the requirement that evidence of a "flash of genius" was required to make an invention patentable—thereby leaving a generation of patent attorneys with the desperate problem of making such flashes visible.

There seems to be no doubt, however, that the final step of discovery is often a sudden event and sometimes an unexpected and surprising one to the inventor, who can give little account of the process that arrived at it. Often, the final step was preceded by a period of "incubation" during which the discoverer was ostensibly preoccupied with other matters. Such events must also receive an explanation in a theory of creativity. They seem somewhat less inscrutable when we reflect that (*i*) conscious awareness is not an invariable property of mental processes and (*ii*) the theory is to account for the entire process and not simply the final step. In particular, it is typical of the act of recognition in humans that, when we recognize a familiar pattern (a human face or a beetle that is familiar from our scientific studies), we are unable to report reliably what cues in the stimulus were essential or instrumental to the recognition. It is notorious, for example, that the characteristics recorded in taxonomic keys (many determinable only under the microscope) are not at all the characteristics that an experienced taxonomist uses for immediate (and often quite reliable) recognition of species with which he is highly familiar.

Knowledge and persistence

The ability to achieve sudden insights into situations by recognizing familiar features in them depends on having stored a great deal of knowledge—knowledge about the familiar patterns that can be recognized and knowledge of the cues for recognition. We know a good deal today, from research on creativity in a few disciplines, about the amount of knowledge that a world-class expert must have, and about the length of time it takes to acquire it.

With respect to amount of knowledge, careful study of the ability of chess grand masters to recognize familiar configurations of chess pieces leads to the estimate that they have at least 50,000 familiar "friends"—patterns that they will note and recognize immediately in a position on a chess board. That number seems not unreasonable when we recall that it is of the same order of magnitude as the word recognition vocabulary of a college graduate for words in his native language.

With respect to the time required for a world class expert to acquire the recognition capability and other knowledge he requires, studies of biographical data in the domains of chess playing, musical composition, and painting arrive at a rather consistent minimum estimate of 10 years. That is, even the most talented persons do not reach world-class performance in their fields until they have devoted about 10 years (or more) of rather single-minded attention to becoming expert. (The time may be somewhat, but not a great deal, less for athletes). The 10-year figure holds for child prodigies (Bobbie Fischer, Mozart, Picasso) as well as for persons who begin the ascent to expertness later. Knowledge and persistence do seem, indeed, to be prerequisites to high-level performance.

Weak methods

To construct an adequate theory of discovery, one must attend to the conditions under which novelty is likely to emerge. Problems that call for creativity are precisely problems from domains that have not already been well worked over and in which sophisticated, systematic algorithms for solution do not exist. In such ill-structured problem domains, problem-solving methods cannot be closely attuned to the characteristics of the problem environment—such tuning requires a great deal to be known about that environment. If we are given a linear algebraic equation in one variable to solve, we simply apply a well-worn and overlearned algorithm to solve it. Ingenuity is required only when such an algorithm is not known to us.

Hence, we may predict that persons tackling problems whose solution will have marks of novelty and require creativity will use very general methods that do not rely on specific knowledge about the problem domain. Such methods, on the other hand, are likely to be highly inefficient; all that commends them is that no better ones are available.

Research using computer programming languages to design problem solvers has uncovered a number of such methods, which are nowadays usually called "weak methods." An example of a weak method is generate-and-test, which consists simply in devising possible solutions and then testing each one to see whether it satisfies the solution conditions of the problem—essentially trial-and-error search. A somewhat more powerful weak method, which requires for its application only a little more knowledge about the problem domain, is means-ends analysis. In this method, a present situation is compared with a goal situation and one or more differences between them are detected. When a difference is noted, it may (by recognition) call forth from memory an operator that is relevant to reducing differences of this kind. The operator is then applied, a new situation is created, and the whole cycle is repeated as often as desired.

Most weak methods require larger or smaller amounts of search before problem solutions are found, but the search need not be blind trial-and-error—in fact, usually cannot be, for the search spaces are generally far too vast to allow unselective trial and error to be effective. Weak methods generally incorporate Polya's idea of "heuristics"—rules of thumb that allow search generators to be highly selective, instead of searching the entire space.

A computer simulation of discovery

There is now a considerable body of evidence that problem solving of a kind that would be regarded as creative if exhibited by human beings can be produced by computer programs that use weak methods to search very selectively in unfamiliar problem domains and that are able, with the help of such methods, to detect interesting laws in data and to invent new interesting concepts. Now I must qualify that last statement a little bit. These programs have been tested mainly not by setting them out to search for genuinely new knowledge, previously beyond human ken, but by setting them the problem of rediscovering important scientific laws and concepts, starting from essentially the same situations that the original human discoverers did. They are like the young Gauss, you might say, who was (rightly) judged creative by his teacher when, at a tender age, he rediscovered the formula for the sum of the first N integers.

One of these programs, which I should like to describe briefly, has rediscovered Kepler's third law, Ohm's law, Black's law of heat, Snell's law of refraction, the Gay-Lussac law of gaseous chemical reactions, some aspects of Dalton's law of chemical reactions, and some others. In the course of doing these things, it reinvented the concepts of inertial mass, specific heat, refractive index, atomic weight, and molecular weight (distinguishing between the latter two). It also rediscovered conservation laws for momentum and heat.

The program whose exploits I have listed is called BACON, after Sir Francis Bacon, for it is a data-driven inductive machine. Its inputs are numerical and nominal data, it is capable of performing factorial experiments on these data, varying one independent variable at a time, but it is provided with no theoretical knowledge about the phenomena. The historical rediscovery tests we have posed to it were derived primarily from situations where the original discovery was, as far as we know, primarily data driven, and where little relevant theory existed before the discovery. The initial versions of BACON were produced by Patrick W. Langley as a doctoral project, and its extensions are joint products of Langley, Gary Bradshaw, and myself.

BACON carries out relatively little search—it certainly does not rely on brute-force trial-and-error exploration of the data. On the other hand, its methods are relatively few and simple, based on a handful of heuristics. What is interesting and remarkable about it is that discoveries (rediscoveries) of such magnitude could be produced by such a simple mechanism.

Let me give a very brief sketch of BACON's heuristics. First, the program tests for correlations among pairs of variables. When a substantial correlation is discovered (other variables being held constant), BACON tests the product or ratio of the variables (as appropriate) for invariance. The product or ratio, whether invariant or not, is introduced as a new variable. Thus, for example, BACON reaches Kepler's third law by examining P/D, P/D^2, and P^2/D^3. It engages in little extraneous search and finds the correct function in a few seconds.

Another rule of thumb used by BACON is to try first laws that are symmetric in appropriate ways. When examining calorimetric data, it discovers that the law depends on the product of the temperature, mass, and specific heat of one of the substances involved. It then hypothesizes directly that the final law will involve the product of the same properties of the other substance. A third rule of thumb is to pay attention to numbers that are simple multiples of each other. This heuristic leads BACON to the chemical laws of simple combinations and to atomic and molecular weight. When BACON finds that there is an invariant relation between properties of two different objects, it postulates a new property belonging to each object to account for the relational invariant. With this heuristic, it invents the concepts of inertial mass (from experiments involving mutual accelerations), specific heat (from the calorimetric experiments), and refractive index.

BACON is very far from a comprehensive theory of scientific discovery. In its present form, it cannot relate one level of phenomena to another (e.g., derive the laws of thermodynamics from the kinetics of gasses), cannot use theory to derive possible laws, does not invent observational instruments, and does not choose data for examination (although it has fair capacities for ignoring irrelevant data without much wasted motion). It has only a modest capacity for dealing with noisy data.

Even with these limitations, what BACON demonstrates is that scientific discoveries of sorts that have been highly regarded can be made with the use of a small armatorium of heuristics and weak methods. The processes it uses are not different from the processes that have been shown, in the psychological laboratory, to account for other, more mundane, human problem-solving performances. It should increase our optimism that scientific discovery can be understood as a natural phenomenon and that it follows laws already made familiar by other research in cognitive psychology. It is not a domain of strange, esoteric, human phenomena. Scientists need not fear for their membership in the human species.

Other computer discovery programs

BACON is not the sole swallow that makes this particular scientific summer. There are other programs that cast light on the discovery process by simulating it that I can only name. One is the AM program of Douglas Lenat that, given a set of basic concepts in set theory, went on to reinvent the integers, the concept of prime number, and the operations on the integers. It conjectured (but did not prove) the fundamental theorem of arithmetic and Goldbach's conjecture. Another is the MOLGEN program, of Stefik and Friedland, which is capable of designing simple experiments in molecular genetics. Through these and other modeling efforts, we are learning a good deal, today about the problem-solving processes that underlie scientific creativity, and we are discovering that these processes rely basically on weak methods and heuristic search.

Conclusion

All of this evidence bears mainly on the discovery process in basic research. I have said nothing about the development and application end of the spectrum. I would hypothesize that the major shift, as we move toward this latter end, is in the formulation of the criteria that define problem solutions. From a central concern with understanding phenomena in depth and a strategy of abstracting from all but the centrally relevant variables, we move to a concern for the whole technical and social context in which a principle is to be applied, including unanticipated consequences and side effects, as well as questions of efficacy and economy. Many of the difficulties in communication among the groups at different points along the research and development continuum are associated with these issues of problem formulation. I doubt whether they imply any differences in the basic psychological processes that are required for problem solution.

1. Hayes, J. R. (1978) *Cognitive Psychology* (Dorsey, Homewood, IL), pp. 215–244: This presents a brief general survey of psychological research on creativity.
2. Simon, H. A. (1977) *Models of Discovery* (D. Reidel, Dordrecht, Netherlands), Sect. 5, pp. 265–338: This presents an elaboration of the theory set forth here and contains numerous references to the literature.
3. Simon, H. A. (1981) *The Sciences of the Artificial* (MIT Press, Cambridge, MA), 2nd Ed., pp. 101–123: This discusses the knowledge of experts and its acquisition and cites relevant references.