

Heuristics for Scientific Discovery: The Legacy of Herbert Simon

Pat Langley

Institute for the Study of Learning and Expertise
2164 Staunton Court, Palo Alto, CA 94306

Introduction

I had the good fortune to interact with Herb Simon on a regular basis between 1975 and 1984, initially as a graduate student in the psychology department at Carnegie Mellon University and later as a collaborator in extensions to my dissertation research. During the latter period, we were joined by Gary Bradshaw, another Carnegie Mellon PhD student, and Jan Żytkow, a visiting researcher with background in the philosophy of science.

My presence at CMU was not accidental, as I had come there specifically to work with Herb Simon. During my undergraduate days at Texas Christian University, I had an odd, recurring experience. For every topic that drew my interest – human cognition, artificial intelligence, philosophy of science, econometric models – it seemed that Simon had published one or more seminal papers, and I decided that I needed to study with this many-faceted scientist. Of course, after arriving in Pittsburgh, I learned that I had seen only the tip of the iceberg, since he had contributed significantly to other fields as well.

Clearly, Herb Simon was fascinated by many phenomena, but in this essay I will focus on only two of them. One concerns the heuristic nature of human decision making and problem solving. Heuristics are rules of thumb that let one simplify a decision task or problem to make it more tractable. Their use typically does not guarantee that one will find the optimal solution, or even any solution at all, but they can make practical the achievement of many tasks that would otherwise fall beyond the reach of human memory and cognition.

On another front, Herb Simon was intrigued with science itself as an instance of complex cognitive activity. He showed a special interest in the creative aspects of science that lead to the discovery of laws, the formation of explanatory theories, and the design of experiments. He was committed to moving beyond vague accounts of discovery processes to characterize them in terms of cognitive mechanisms. Naturally, the notion of heuristics figured centrally in his work on scientific discovery, as it did in his studies of problem solving and other aspects of human thought (e.g., Newell & Simon, 1972).

Thus, it seems appropriate to review Herb Simon's career in terms of his personal heuristics for scientific research, both because it will illuminate many aspects of his scientific style and because it may hold useful ideas that others can incorporate into their own work.¹ Moreover, it makes sense to illustrate these rules of thumb with examples taken from Simon's work on discovery processes. Fortunately, I was in a good position to collect these examples, since our collaborations over the years centered on just this topic.

1. For Simon's own thoughts about his research heuristics, see the afterword, "The Scientist as Problem Solver", in Simon (1991, pp. 368-387).

Be Audacious

Herb Simon was audacious in that he was willing to tackle imposing problems which others had been reluctant to face or even admit might be solvable. Not only did this heuristic increase his chances of making large-scale advances, but it took him regularly into more interesting, less explored territories than other researchers. One early problem, which he pursued well before my time, focused on building the first computer to prove theorems in logic and, arguably, the first to emulate human thought processes (Newell & Simon, 1956).

A second important pursuit, which I was lucky enough to partly witness, involved uncovering the cognitive and computational mechanisms that support the processes of scientific discovery. In 1966, Herb Simon published a chapter, “Scientific Discovery and the Psychology of Problem Solving”, in which he had the audacity to suggest that one could explain and model scientific discovery, which often involved some creative insight, in terms of normal cognitive processes, using the same mechanisms that underlie everyday human problem solving and decision making.

We can only understand the radical nature of this claim by reviewing the intellectual climate of the time. Philosophers of science were focused almost entirely on the logic of verification, by which one evaluates models and theories in terms of evidence. Indeed, some prominent philosophers like Popper (1961) rejected even the possibility of a logic of discovery. In other circles, there was a general air of mysticism about scientific discovery, which supposedly relied on unfathomable intuitions or creative sparks. Herb Simon’s position stood in stark contrast to these neovitalists, and his radical chapter set both the agenda and the tone for research on cognitive and computational studies of scientific discovery for the next 35 years. Challenging problems of this sort do not fall to a single blow of trumpets, however piercing, but we will see later that Simon responded to the complexity of science by drawing on additional heuristics.

However, I should note the contrast between Herb Simon’s audacity as a scientist and his humility as an individual. Meetings in his Baker Hall office, despite the books lining its walls, the papers stacked on its tables, and the expressions on its slate blackboard, felt less like visiting an academic’s office than coming to the living room in the house of an older, valued friend. Herb sat in an easy chair, with no intervening desk, and he treated even the least experienced student as a colleague. He was always interested in your ideas and, though he might disagree with you, he never seemed to judge. As a result, the physical space that he engendered was friendly but fully professional, challenging but always respectful, and comforting but genuinely creative. Despite his many scientific breakthroughs, Herb maintained that he was nothing special, just a normal human engaging in everyday problem solving, and he reflected that attitude in his style of interaction. I count my visits to Herb Simon’s office, and my meetings with him and colleagues there, as the most memorable, enjoyable, and important times of my life.

Ignore Discipline Boundaries

Another heuristic that Herb Simon applied repeatedly was to ignore traditional boundaries between disciplines. For any given research problem, he became familiar with the work done in every relevant field and incorporated the best ideas from each one into his own approach to the task. This strategy provided him not only with a rich source of metaphors and methods; it also imposed constraints that narrowed his search for a solution.

In his research on scientific discovery, Herb Simon borrowed concepts from a variety of disciplines. Naturally, he took advantage of his earlier ideas, and those of his collaborators, on the nature of human problem solving, especially the notion that this activity involves constrained search through a problem space. He also incorporated techniques from artificial intelligence, many of which he co-developed, to implement the search metaphor in running computer programs. In addition, he read widely in the philosophy of science, borrowing ideas about what constitutes a legitimate theory, and in the history of science, from which he collected phenomena to explain with his computational models.

During our research meetings on scientific discovery, Herb frequently suggested papers or books that were relevant to our current discussion. These might come from the literatures on cognitive psychology, artificial intelligence, philosophy, or history, and we could never predict in advance from which discipline he would draw. Herb Simon had an incredibly dense semantic network that indexed ideas not by their field of origin, but rather by their relevance to problems that interested him. Of course, we were able to read only a fraction of the items that he suggested, as we found it difficult to match his pace and scope.

Moreover, Herb Simon's Renaissance scholarship was bidirectional, in that he made his research results accessible to members of the communities on whose earlier ideas he built. Thus, he published his contributions on scientific discovery in the literature on cognitive science (e.g., Qin & Simon, 1990), artificial intelligence (e.g., Valdés-Pérez, Żytkow, & Simon, 1993), and philosophy of science (e.g., Simon, Langley, & Bradshaw, 1981), and he kept in touch with historians whose analyses he had used. He also authored works designed to make contact simultaneously with all of these communities (e.g., Langley, Simon, Bradshaw, & Żytkow, 1987). In each field, Simon's papers gradually convinced many that discovery was a topic suitable for scientific study, and his work gained a respectable following therein.

Use a Secret Weapon

A third heuristic that appeared in Herb Simon's research was his repeated use of "secret weapons" – methods and metaphors that he had mastered but that were not yet widely available to the broader community. These intellectual tools let him solve problems more rapidly, and with less effort, than most other scientists, and their use accounts, at least partly, for his great research productivity. Herb Simon was quite aware of this strategy, and I have borrowed the epigram "use a secret weapon" directly from his advice to students.

In his research on scientific discovery, Simon drew repeatedly on two frameworks that had been developed at Carnegie Mellon to model cognitive processes in humans. The first was the idea, mentioned earlier, that problem solving involves heuristic search through a problem space. Adopting this view requires one to specify the initial state, operators for generating new states, heuristics for selecting among alternatives, and some halting criterion. The second innovation concerned the use of production systems to model cognition. A production system consists of a long-term memory, stated as a set of condition-action rules, and a short-term memory that contains a set of goals or beliefs. On each cycle, the system matches the conditions of each rule against the contents of short-term memory, selects one or more matching rules, and executes their actions, which change these contents.

Both frameworks suggested mechanisms that we used in our models of scientific discovery, but they also supplied constraints on the systems we developed. They constituted secret weapons because they provided powerful metaphors, along with associated computational techniques, that had not yet been adopted by the research communities in either cognitive psychology or artificial intelligence, although Simon and his colleagues had been using both for some time. Ironically, Herb's secret weapons were never truly secret, in that he was willing to share his metaphors and methods with whoever was willing to adopt them. Over the years, he did much to popularize the notions of problem space search and production systems across different research communities, both in his papers and in his talks, and both in the context of specific models and in more general terms.

Balance Theory and Data

Yet another aspect of Herb Simon's research that we can characterize as a heuristic was his balance between theory and data. He realized that scientific models must explain observations, as this is a major criterion by which one judges their success. But he also understood the need for these models to remain connected with existing knowledge that had previously explained other, equally important, phenomena. When constructing models of cognition, he utilized effectively the constraints imposed by both data and knowledge to guide his efforts.

Herb Simon's work on computational models of scientific discovery exhibited this balance between theory and data. In this case, the phenomena involved episodes from the history of science, ranging from the discovery of descriptive relations, such as Kepler's third law of planetary motion and the chemical behavior of acids and alkalis, to the formation of explanatory models, such as structural models of chemical compounds and biochemical pathways in organisms. These were important events that should be handled by any complete account of the scientific process. At the same time, such an account should be constrained by established knowledge about human cognition, including the heuristic nature of problem solving, and connect to useful theoretical formalisms, such as production systems.

However, although Simon always operated within a clear theoretical framework, he was willing to expand upon and revise it when the need arose. For example, many ideas from

our early work on the discovery of numeric equations (like Kepler's law) proved transferrable to modeling the construction of qualitative laws (like those about acids and alkalis), but they were less relevant to modeling the formation of structural or process models (like the molecular model of water). In this case, he introduced new state representations and new operators to handle the novel aspects of these problems, though the models were still cast within the broader frameworks of heuristic search and production systems, which were still consistent with these discovery phenomena.

Satisfice

One of Herb Simon's central contributions, made well before he delved into either cognitive simulation or artificial intelligence, was the observation that humans seldom find optimal solutions to the problems that confront them. Rather, they *satisfice* by finding solutions that are good enough for their purposes but that they can obtain with the limited cognitive resources they have available. Although developed originally to explain aspects of human decision making, Simon and his colleagues also applied this idea profitably to their development of artificial cognitive systems.

However, Herb Simon treated the notion of satisficing as more than a useful theoretical construct; he also invoked it regularly as a heuristic in his own research. Our initial work on scientific discovery addressed a number of challenging problems, but we also idealized these tasks enough to make them tractable. For instance, we quite consciously focused on episodes from the early history of science, where relatively simple and general methods had sufficed to make important discoveries. Moreover, we deliberately ignored issues of problem formulation, variable selection, design of measuring instruments, and other aspects of scientific reasoning.

These idealizations were a source of continuing criticism of our discovery work for many years, but Herb's response reflected his approach to research. He pointed out that science operates in incremental steps, not by producing completed theories that emerge fully grown from the scientist's brow. Our work, he claimed, constituted progress in that it provided partial computational accounts of scientific phenomena where none had existed before. Nevertheless, he also acknowledged the limits of our idealizations and the need for additional research on this important topic.

On a related note, I have always been something of a perfectionist, which has often led me to take longer in finishing papers than might be desirable. One of Herb's favorite lines was "Anything worth doing is worth doing badly", which he used to encourage me and, I suspect, many others. It took me years to realize that he meant this as more than just practical advice, and that it followed directly from the theory of satisficing and bounded rationality that was the cornerstone of his intellectual career. However, to my knowledge, Herb Simon never wrote a bad paper in his long and productive life. But then many of us have long suspected that his rationality was somewhat less bounded than our own.

Persevere

The gradual nature of science means that major advances seldom occur overnight. For this reason, Herb Simon combined his satisficing strategy with another, complementary, one. His research program involved *persevering* over extended periods of time, letting him build incrementally on his previous results and extend his models to cover ever more phenomena. He was realistic about the time and energy that true progress in science demands, even when taking advantage of his other heuristics, and he was willing to devote the effort required.

Herb Simon's research on computational models of scientific discovery provides a clear example of such perseverance. Not even counting his early essays on the topic, he worked steadily in this area for well over 20 years. Most of this research involved collaboration with others, but none except Herb had the sheer staying power to continue working actively on the topic throughout this period. The result was a substantial body of results on the nature of scientific discovery, cast in cognitive and computational terms that made clear contact with other knowledge about human behavior.

Moreover, Herb Simon and his colleagues incorporated the notion of perseverance in research to their work on the scientific process itself. Our early efforts had focused on isolated discoveries from the history of science, despite the clear evidence that research often involves a lengthy sequence of interconnected steps. Later, in joint work with Deepak Kulkarni, Simon developed a computational model of extended episodes in science, such as Krebs' discovery of the urea cycle. Their account (Kulkarni & Simon, 1988) combined heuristics for designing experiments, formulating qualitative relations, and specifying new problems based on surprising results, all of which are evident in Simon's own career.

One additional benefit of perseverance, less obvious at first glance, is the accumulation of knowledge in the form of domain-specific *chunks*. This knowledge not only makes the research process itself more effective, but can aid in the communication of results. Herb Simon told a story about sitting down to write a paper and, with little effort, completing the introductory paragraph. But then he noticed the text seemed very familiar. He turned to his filing cabinets and, after some search, uncovered another paper, written over a decade earlier, with nearly the same introduction, word for word. We all acquire chunks from experience but, as usual, Herb did things in a bigger way than most of us.

A Research Challenge

Although Herb Simon is no longer with us, he would certainly have wanted us to carry on the research program he started so many years ago. And despite his many achievements, he would acknowledge that our understanding of scientific discovery remains incomplete, and that we should continue to study this diverse activity and extend our computational models of its operation. In our efforts, we should take advantage of the heuristics he utilized in his own research career, along with any others that aid us in addressing this complex task.

In this context, I will suggest a research problem that would stretch our abilities and shed further light on the nature of science. I propose that we develop a computational model for the behavior of a scientist who, over the course of his career:

- formulated the notion of satisficing in human decision making;
- co-invented list processing and heuristic search on computers;
- co-developed computational theories of human memory and problem solving;
- used these theories to model scientific discovery and other key phenomena; and
- fostered a new field, cognitive science, that knows no discipline boundaries.

We have already seen some of this scientist's strategies, and we have detailed records of his many accomplishments. The task of modeling his wide-ranging behavior seems challenging, but that would not have daunted Herb Simon, and, armed with the legacy of his heuristics, we know how to proceed.

References

- Kulkarni, D., & Simon, H. A. (1988). The processes of scientific discovery: The strategy of experimentation. *Cognitive Science*, *12*, 139–176.
- Langley, P., Simon, H. A., Bradshaw, G. L., & Żytkow, J. M. (1987). *Scientific discovery: Computational explorations of the creative processes*. Cambridge, MA: MIT Press.
- Newell, A., & Simon, H. A. (1956). The logic theory machine. *IRE Transactions on Information Theory*, *IT-2*, 61–79.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall, Inc.
- Popper, K. R. (1961). *The logic of scientific discovery*. New York: Science Editions.
- Qin, Y., & Simon, H. A. (1990). Laboratory replication of scientific discovery processes. *Cognitive Science*, *14*, 281–312.
- Simon, H. A. (1966). Scientific discovery and the psychology of problem solving. In R. G. Colodny (Ed.), *Mind and cosmos: Essays in contemporary science and philosophy* (pp. 22–40). Pittsburgh: University of Pittsburgh Press.
- Simon, H. A. (1991). *Models of my life*. New York: Basic Books.
- Simon, H. A., Langley, P., & Bradshaw, G. (1981). Scientific discovery as problem solving. *Synthese*, *47*, 1–27.
- Valdés-Pérez, R., Żytkow, J. M., & Simon, H. A. (1993). Scientific model-building as search in matrix spaces. *Proceedings of the Eleventh National Conference on Artificial Intelligence* (pp. 472–478). Menlo Park, CA: AAAI Press.