9
From Bounded Rationality to Expertise

Fernand Gobet

Introduction

Historically, a pervasive assumption in the social sciences, in particular economics, is that humans are perfect rational agents. Having full access to information and enjoying unlimited computational resources, they maximize utility when making decisions. As is well known, Herbert A. Simon rejected this assumption, calling it a fantasy for two main reasons. First, the complexity of the environment makes it impossible for humans to have full access to information. Second, a number of important restrictions impede the human cognitive system, such as limited attention and slow learning rates. Therefore, humans display only a bounded rationality and must satisfice, i.e., make decisions that are good enough, but not necessarily optimal.

Research into expertise has contributed to the question of rationality in two important ways. First, to what extent can some of the very best among us – super experts – approximate full rationality? Second, by what means do experts, at least in part, circumvent the constraints imposed by bounded rationality?

This chapter takes the shape of a fugue, with the themes of bounded rationality and expertise first played in the background of personal recollections, and then elaborated with a more formal survey of Simon's research into expertise. The themes are played a third and final time with a discussion of the heuristics (rules of thumb) proposed by Simon for achieving a successful career in science.

R. Frantz et al. (eds.), *Minds, Models and Milieux*

Becoming an Expert: A Personal Recollection

My collaboration with Herbert A. Simon lasted over ten years, including six years spent at Carnegie Mellon. While I was working on my PhD thesis on chess players' memory, I secured a research fellowship from the Swiss National Science Foundation to work with him. The qualifications I listed in my introductory letter to Simon were rather limited, a first degree in psychology and the title of International Chess Master. Simon, who probably saw an opportunity to reactivate research he carried out on chess expertise in the 1960s and 1970s (see below), but which had been dormant since, accepted to host me.

Meeting the Man

I can still recall our first meeting on a beautiful morning in January 1990. His office was welcoming, but also rather disorganized, with stacks of papers and books hiding his desk. The meeting was short but cordial, and Simon gave me advice about life and housing in Pittsburgh, and briefly talked about the projects he was currently involved in.

The second meeting was my first real scientific discussion with Simon. It was actually a shock. In a polite and friendly way, Simon demolished the research line I had in mind for my PhD, which was to elicit a chess grandmaster's knowledge of a small and specific domain (rook and pawn endgames), and to build a program implementing this knowledge. The aim, inspired by research on expert systems, was to compare the amount of procedural (knowing how) and declarative (knowing that) knowledge. Simon found that the project was not realistic enough ('A player like Kasparov will give you lectures on rook endgames for several days; what are you going to do with all these data?'). In addition, he thought that the project would dovetail better with the research of his colleague John Anderson. I can still feel the panic that invaded me when he told me this, as it was an invitation to sever collaboration before the end of the first meeting! In the discussion that followed, he made it clear that he would prefer a project directly linked to the chunking theory he had developed with Chase in the 1970s to account for chess expertise (Chase and Simon, 1973a, 1973b). This influential theory, and in particular the computer model MAPP (Memory-Aided Pattern Perceiver; Simon and Gilmartin, 1973) that implemented it in part, had been severely criticized, and Simon wanted to improve on it. Thus, my first lesson was that Simon, while open to other ideas, was very selective about the research lines he invested time in, and made sure that they addressed his central interests.

Luckily, my back-up project precisely addressed the issues that Simon mentioned. However, it was less developed than my first choice, and I spent a rather anxious week trying to improve on it. The idea was to carry out a series of experiments on chess players' memory, focusing on the amount of information that they could memorize after a brief presentation, and the extent to which this information was encoded with specific cues about spatial localization. The experimental results would be simulated by a program based on MAPP and the idea of chunking. With hindsight, it is obvious that this second project, which combined experiments with computer modeling, fitted Simon's scientific approach much better. In addition, MAPP was a variation of EPAM (Elementary Perceiver and Memorizer; Simon and Feigenbaum, 1964), a general theory of learning embodied as a computer program, which played a central role in Simon's approach to expertise and bounded rationality more generally.

My visit was supposed to last for one year. However, the privilege of working with Simon and the exciting, interdisciplinary CMU environment convinced me to extend it, and I eventually stayed for six years. I would meet Simon face-to-face for one hour a week on average. In general, I would decide on the agenda, which typically included our research on chess memory and a discussion of some of his other research projects, not only in psychology but also in other fields, mostly artificial intelligence and philosophy.

'What new data do we have about chess today?' Simon would often ask at the beginning of our meetings. Sitting on a simple chair, with no desk between him and his guest, he had an informal and welcoming style. He was immensely curious about new phenomena, and took obvious pleasure in analysing data, always looking for hidden patterns. He was outstanding at making sense out of complex data sets, but would also often ask me to re-analyse data, by using better representations for example.

Meetings with him were alternatively dense and relaxed, focused and wide-ranging. His informal style made me sometimes forget that I was talking to one of the greatest minds and one of the last true humanists of the twentieth century. While open to new ideas, he also would immediately, albeit elegantly, rebut bad arguments. However, even then, he managed to refute them in a way that did not make you feel too dejected – even though more often than not, your idea had been irremediably torn apart.

In spite of his age, Simon was extremely energetic. Not only did he have an active program of research, but he also had a normal teaching load and was deeply involved with the politics of CMU. (Imagining his energy when he was in the prime of his life was a rather depressing exercise.) For him, research was a hobby, and working 80 hours a week seemed reasonable. Of course, he expected his collaborators to show a similar dedication to research, including their working hours. This expectation was a logical application of his work on expertise, where he suggested that, in most domains, a minimum of 10,000 hours of practice and study was necessary to become an expert (Simon and Chase, 1973).

Simon was exigent and expected his collaborators to share his passion for work. Working with him was extremely rewarding, and most of his students and collaborators did show strong commitment, although few managed to work for 80 hours a week consistently. In a very elegant and efficient way he taught us much about science, mainly that scientific research can be an exciting voyage of discovery.

Another striking characteristic of Simon was his curiosity; the title of one of his talks was 'The cat that curiosity couldn't kill.' He had mastered an astonishing range of domains, from piano to chess to entomology, and was interested in all sorts of topics, including non-scientific questions. Once he had decided to find an answer to a question – sometimes a trivial question – he would work on it obsessively; it was as if heuristic search was a goal in itself. This, of course, is a powerful personality trait for a researcher to have.

Simon was also very generous with his students and colleagues, and always ready to support us in difficult moments. When some arcane immigration regulation allowed me to be employed by Carnegie Mellon University but not to be paid, he did not hesitate to support me financially for a couple of months. His generosity went beyond the academic world and he gave the money associated with his Nobel Prize to charity.

Collecting Data

As noted above, our collaborative work focused on chunking theory and on the empirical difficulties it faced. In their seminal work, Chase and Simon (1973a, 1973b) had analysed in great detail the way players of different skill levels recalled briefly-presented chess positions. Their hypothesis was that experts in chess – and other domains – encode domain-specific material as *chunks*: groups of items that are related perceptually and semantically. In line with bounded rationality, a chunk takes a relatively long time to store in long-term memory, but it can be accessed rapidly afterwards, in a few hundred milliseconds.

However, Simon had been tormented by an anomaly in the data ever since the publication of his work with Chase. Although the master used larger chunks than the weaker players, as predicted by the theory, he also used more chunks, which was inconsistent with the hypothesis that all players had the same limited short-term memory capacity. In addition, several experiments with individuals trained to remember long lists of digits presented for one second each (Chase and Ericsson, 1982), clearly indicated that information could be encoded more rapidly into long-term memory than postulated by theory. Simon hypothesized that experts in chess and in the digit-span task used *retrieval structures*, that is, long-term memory schemas allowing them to encode information rapidly. What eluded him was the exact nature of these structures.

Our replication of Chase and Simon's experiment improved on it in two important ways: we had a much larger sample; and we collected more reliable data by using a computer program rather than physical boards and pieces. The results clearly showed that masters possessed and used much larger chunks than had been found in the original study (Gobet and Simon, 1998). In addition, critically, the number of chunks replaced did not differ between skill levels, all players used three chunks on average. Simon was very pleased with these results, not only because they corrected the anomaly found in Chase and Simon's study, but also because they were consistent with data he had collected a few years earlier on the memory for Chinese ideograms (Zhang and Simon, 1985). This interest in numerical parameters of cognition (another is eight seconds to learn a chunk) was a signature of Simon's scientific style and a logical consequence of the hypothesis of bounded rationality.

Support for the three-chunk capacity of visual short-term memory was also found in a series of experiments where we tried to establish the limit of chess experts' memory. In these experiments, chess players had to memorize not only one chess position presented for a short amount of time (typically five seconds), but also several of them, each containing around 25 pieces (Gobet and Simon, 1996d). This task was very difficult, and only masters could cope with it. Interestingly, there seemed to be a limit of around three or four boards – again, the magical number three!

Simon was intrigued by this task, and convinced me to conduct an experiment with his favorite research method, to study a single subject in great detail. He also persuaded me to volunteer as the subject, I was an International Chess Master after all! In a longitudinal experiment that lasted nearly two years (with a few interruptions), I practiced memorizing as many positions as possible, several times a week. The software I had written, together with a large database of chess positions, allowed me to carry out a well-controlled experiment, despite being both the experimenter and the subject. This experiment produced a huge amount of data (every single action and its timing was automatically recorded), which, unfortunately, we were slow to analyse (For preliminary analyses, see Gobet and Simon, 1996d, Gobet, 2013). I improved from four boards to a personal record of ten boards with a minimum of 80% accuracy on each board, but was never able to go beyond. At the beginning of our weekly meetings, Simon would ask me about my progress – and often my lack of progress – and we used these data to develop mechanisms for retrieval structures in chess in what became the *template theory* (Gobet and Simon, 1996d).

In another experiment, we studied the effect of modifying chess boards by mirror image reflections (Gobet and Simon, 1996a). The results provided useful information on the way chunks are encoded by chess players. In particular, they showed that information about location is encoded in chunks. Another experiment systematically varied the presentation time from one second to 60 seconds (Gobet and Simon, 2000). The results confirmed, with a visuo-spatial task, a key parameter that Feigenbaum and Simon (1984) had estimated with a verbal task, that it takes around eight seconds to create a new chunk in long-term memory. While these results were important, it was another result, which I had mentioned in passing and rather anecdotally, that captured Simon's imagination. World champion Garry Kasparov had played several matches against national teams, facing up to eight grandmasters or masters simultaneously. In most cases, he had won these matches. Crucially, computing Kasparov's performance in these matches showed that he played, on average, at a level that still placed him in the six best players in the world at the time (Gobet and Simon, 1996c). For Simon, this was a spectacular illustration of the role of pattern recognition and selective search in expert decision making, of how experts can (partly) overcome the limits imposed by bounded rationality on their cognitive abilities.

Building Computer Models

One exciting aspect of my collaboration with Simon was the development of several computer models. It was also a very challenging experience, not only for technical reasons, but also because, whatever the beauty of one's model, the moment of truth is whether the model accounts for the experimental data – and experimental data are ruthless. The first version of CHREST (Chunk Hierarchy & REtrieval STructures) combined MAPP with PERCEIVER, a model simulating chess players' eye movements (Simon and Barenfeld, 1969). One improvement over the earlier models was that chunks were not input by the modeler, but were learned autonomously as a function of the positions that the model had seen. Eye movements played an important role in learning and, in turn, the chunks that had been learned were essential in directing future eye movements.

The key contribution of CHREST was closely to link mechanisms for perception, learning, and memory, and to provide mechanisms for the concept of a retrieval structure. Our first hypothesis was that chess players' retrieval structures are similar to the structures used by individuals specialized in memorizing digits; such structures are generic and can be used with any kind of material as long as it is taken from the domain of expertise. This version of CHREST was able to simulate several empirical data successfully, but also suffered from some serious weaknesses. In particular, it overestimated recall performance with random positions, and could not replicate the experimental results obtained with modifying chess boards by taking their mirror image.

After much trial and error, we reached the conclusion that chess players' retrieval structures were specific; that is, they could be used only when the board contained some specific patterns. In a sense, these templates were more similar to the schemas discussed in psychology and artificial intelligence than to the structures identified in the digit-span task. The modified version of CHREST accounted for a wide range of data concerning eye movements, recall performance with diverse types of positions (game positions, random positions, positions modified by mirror image), number and type of errors made, and type of chunks used. An important contribution of the model, which goes beyond chess, was that it provided mechanisms explaining how schemas are constructed automatically, including the way variables and default values are built. CHREST was later applied to other domains of expertise and to the simulation of first language acquisition, which can be considered as a kind of expertise (Gobet *et al.*, 2001).

An interesting episode in the development of CHREST is worth mentioning, since proving Simon wrong was extremely difficult. When simulating the recall of random positions, the second version of CHREST systematically predicted that there should be a skill effect, although it was much smaller than with game positions. This prediction was contrary to Chase and Simon's (1973a) result, where the master performed as badly as the weaker players with random positions. In fact, this lack of skill difference had become a standard result in psychology, found in most textbooks. Simon first thought that there were bugs in my program. However, after much double-checking and many replications, it was clear that the effect was genuine; even in random positions, a version of the model knowing many chunks is more likely than a version knowing fewer chunks to recognize, by chance, a pattern of pieces on the board. When we ran new experiments and collated all studies in the literature that had used random positions as a control condition, it became clear that the effect was also present with humans. In most experiments, the effect was small and statistically non-significant, but it was reliable when all experiments were combined (Gobet and Simon, 1996b). Ironically, these results provided strong support for chunking theory, since they are difficult to explain with other theories of expertise (Gobet, 1998).

CHREST was developed at the same time as two other models based on EPAM. The first one accounted for results in categorization, and in particular highlighted the role of strategies in that task (Gobet *et al.*, 1997). The second explained how an individual with a normal short-term memory capacity managed, after intense practice and training, to memorize 106 digits, each presented for one second only (Richman *et al.*, 1995, Richman *et al.*, 1996). In the second part of my stay at CMU, these models were discussed during near-weekly meetings of the EPAM group, which also included Howard Richman, Jim Staszewski, and Shmuel Ur. These meetings were very lively and included a considerable amount of brainstorming, and while sometimes lacking structure, they offered a productive environment for exploring various aspects of EPAM. Simon was active in these discussions, but was non-directive.

Bounded Rationality and Expertise

The kind of simple tasks typically studied in psychology can only go so far to identify the properties of human cognition, including strategies and cognitive invariants. By studying much more demanding tasks, research into expertise offers a unique window into cognition (Gobet, 2015), and, in particular, how humans cope with complex environments. As a first approximation, it is possible to divide Simon's research on expertise into three periods: problem solving (until the mid-1960s); perception and memory in chess (late 1960s to mid-1970s); and broadening the horizon (from the late 1970s to Simon's death).

Problem Solving

Simon's interest in expertise is apparent in his early books, such as *Administrative Behavior* (Simon, 1947a) and *The Technique of Municipal*

Administration (Simon, 1947b), where the issue is how organizations can best use expert skills, particularly with respect to decision making. Simon argues that hierarchies offer the best structure to do so, as they allow decisions to be made in the part of the organization where they are most useful.

During the 1950s, Simon started to study expertise empirically through his work on chess problem solving. Dutch chess master and psychologist Adriaan de Groot had shown that chess players' search is highly selective; among the numerous moves possible in a position, they look only at a handful (De Groot, 1946). Indeed, the stronger the player, the narrower the search behavior, inferior options being rarely considered. To flesh out the mechanisms that allowed moves to be generated so selectively, Newell and Simon (1965, 1972) carefully analysed the verbal protocol of a single player trying to find the best move in a given position. Several characteristics of move generation were identified. For example, if the analysis of a move leads to a positive evaluation, then this move is further investigated. If the evaluation is negative, then a different move is examined.

The research led to several simulation models, which are of considerable interest since they represent a direct attempt to model concepts borrowed from Simon's theory of bounded rationality. As is well known, Simon strongly advocated the use of formal models in the social sciences. While he originally used mathematical methods, such as differential calculus, he had noted essential limitations with them and concluded that other techniques were necessary. This, of course, led to the development of artificial intelligence and computer modeling, tools that were not, for many years, distinct in his mind. For the study of expertise, computer modeling has clear advantages: theories are precisely stated; clear predictions can be made, both quantitatively and qualitatively; and simulations can examine the structure of the task environment.

NSS, a chess program developed by Newell, Shaw and Simon (1958) uses goals such as maintenance of material balance and control of the center. Based on these goals, two move generators are engaged: the first generates possible moves in the problem situation; the second generates moves that are possible during look-ahead search (some search is carried out to evaluate the suitability of the proposed moves). NSS directly implements the concept of *satisficing* by playing the first move that is evaluated above an aspiration threshold. The program demonstrated that it is possible to choose reasonable moves with very small search trees (less than 100 positions). However, the quality of its play was low.

Selective search is also demonstrated by MATER (Baylor and Simon, 1966). This was achieved by limiting search to forced moves and moves that minimized the number of options available to the opponent. The program played good chess in positions that contained a forced checkmate combination; however, it was limited to this sub-domain of chess.

Science is, of course, a kind of expertise, and Simon devoted considerable attention to this topic. In his early writings about scientific discovery and creativity, he was mostly interested in technical and philosophical aspects (e.g., most of the essays collected in Simon, 1977). However, following their work on problem solving, Newell, Shaw and Simon (1962) speculated in the late 1950s and early 1960s about how the human mind can be creative, and whether this could be described objectively and explained scientifically. Their answer was that creativity is a special case of problem solving, and thus can be studied with the same conceptual tools.

Perception and Memory in Chess

Perhaps Simon's central contribution to the study of expertise was made in trying to answer another question studied by De Groot (1946), how do perception and memory mechanisms allow masters to understand the gist of a position rapidly, in a matter of seconds? As we saw earlier, Chase and Simon (1973a, 1973b) developed a means to identify chunks and conducted a series of clever experiments on chess players' perception and memory. In addition, they proposed mechanisms not only accounting for these experiments but also explaining selective search. The central ideas of their chunking theory are clearly linked to Simon's concept of bounded rationality. Experts' cognition suffers from a number of limitations (e.g., limited-capacity short-term memory and slow learning rates), and these limitations are assuaged by knowledge. Building knowledge predominantly consists of acquiring a large number of chunks, which are both perceptual and semantic units. These units are linked to possible actions, forming *productions*; for example, in chess, if a file is open, occupy it with a rook. Thus, pattern recognition makes it possible to demonstrate expertise despite strict limits on computational capabilities. A strong implication of the theory is that expert intuition is essentially pattern recognition. Another implication is that the best way to explain expertise in chess, and in other domains, is to use the formalism of a production system (i.e., a system specifying how productions are used for solving problems).

As noted earlier, several computational models were developed by Simon and his colleagues to account for perception and memory in chess. PERCEIVER (Simon and Barenfeld, 1969) simulated the eye movements of a chess player, using some of the mechanisms present in MATER, in particular, to compute piece movement. The program was built to take issue with the claim of gestalt psychology that perception is holistic and complex. Using local and simple mechanisms, PERCEIVER was able to convincingly reproduce the eye movements of one player in a given position. The key assumption was that perceptual information relates to attack/defense relations between pairs of pieces or between a piece and a square. One limit of the study was that only one position was presented to test the validity of the program.

MAPP (Simon and Gilmartin, 1973) incorporated some of the mechanisms postulated by chunking theory and, like EPAM, uses a discrimination net to organize information. A discrimination net is a hierarchical network of nodes where features of the objects to learn are tested to determine what new information should be added to the existing hierarchy. MAPP uses two parameters that strongly limit its cognitive abilities: short-term memory can store only seven chunks; and learning a new chunk takes eight seconds. During the presentation of a chess position, MAPP tries to recognize chunks in long-term memory and, when information is successfully identified, pointers to those chunks are placed in short-term memory. During the reconstruction phase, MAPP simply unpacks the information provided by these chunks.

MAPP was able to simulate the recall performance of a strong amateur, but not of a master. Using mathematical extrapolations from the computer simulations, Simon and Gilmartin concluded that from 10,000 to 100,000 chunks (50,000 as a first approximation) are necessary to reach expert level in chess and in other domains. Later simulations suggested that the number might be as large as 300,000 (Gobet and Simon, 2000).

Broadening the Horizon

During this final period, Simon both revisited old research topics and studied new domains of expertise. One new topic was novice and expert differences in solving physics problems, which Simon studied using verbal protocols (Bhaskar and Simon, 1977, Larkin *et al.*, 1980). The results showed that novices tend to search backward, from the goal to the givens of the problem, while experts tend to search forward, from the givens to the goal. However, when problems are difficult, experts revert to backward search. Regardless of difficulty, experts use heuristics to reduce the amount of search and draw on more efficient representations of the problem. The importance of representations is also clear in economics, where experts are better at developing multiple representations, for example verbal and diagrammatic representations (Tabachnek-Schijf *et al.*, 1997).

To account for these empirical results, Simon built several computational models, implemented as production systems. One model accounted for problem solving in thermodynamics, a semantically rich domain (Bhaskar and Simon, 1977). An important offshoot of this work was SAPA, a program that semi-automatically coded verbal protocols. Another production system, ABLE, provided mechanisms accounting for how novices become experts in solving physics problems (Larkin and Simon, 1981), including the change from backward search to forward search. An interesting aspect of ABLE is that it can use declarative statements to derive new results; in turn, these results can be used to solve new problems. Understanding task instructions and problem descriptions is, of course, crucial for building internal representations of problems and solving them, not least because different instructions lead to different representations. This process had been modeled by Hayes and Simon (1974) with respect to several variants of the puzzle known as the Tower of Hanoi.

In economics, Simon and colleagues developed CaMeRa, a model simulating visual reasoning and the way experts combine different kinds of representation (Tabachnek-Schijf *et al.*, 1997). When solving problems, CaMeRa can interact with external representations, such as diagrams. It uses several formalisms: a parallel system accounting for low-level vision; a semantic network storing semantic knowledge; and a production system used for problem solving.

This period also saw a return to the study of scientific discovery. Together with several collaborators, Simon developed a number of computer programs able to simulate famous discoveries in the history of science (Langley *et al.*, 1987, Bradshaw *et al.*, 1983). For example, a computer program called BACON re-discovered several scientific laws, including Kepler's third law of planetary motion, using the same data as those available in the original discoveries. Heuristics made it possible for the program to search selectively through the space of possible equations; interestingly, experiments with students confirmed that they use the same heuristics (Qin and Simon, 1990). Another program, KEKADA, was able to design experiments, change theory as a function of the results, and then design new experiments; it was able to simulate Krebs' discovery of the urea cycle in 1932. In his autobiography, Simon (1991) suggests that these programs would be a good start for somebody trying to simulate his own scientific creativity!

In the first part of this chapter, I described the work that Simon undertook with chess during this period. At the same time we developed CHREST to account for chess expertise, Simon worked on EPAM-IV (Richman *et al.*, 1995), a model accounting for superior digit memory. This model is particularly relevant in the context of bounded rationality, since it shows how experts can strategically compensate for structural limits in their cognitive architecture, in this case, short-term memory. The task under study was the digit-span task, where one has to memorize a sequence of digits that are dictated rapidly (typically, one second each). While most of us can remember only about seven items, some individuals were trained to recall many more. For example, DD, the human subject simulated by Richman and colleagues, recalled up to 106 digits. DD used different kinds of mnemonics (techniques aimed at increasing one's memory), and a sophisticated semantic knowledge of numbers (historical dates, typical running times), to produce such an incredible feat. In line with previous research, he also used retrieval structures (Chase and Ericsson, 1982), structures that enable a rapid storage in long-term memory. DD's behavior and performance are obviously hard to explain with chunking theory, and consequently spurred Simon on to develop a model accounting for these results.

Just like CHREST, EPAM-IV combines chunking mechanisms with the notion of a retrieval structure. The difference, however, is that the retrieval structures postulated in the digit-span are acquired deliberately and consciously. The model specifies, in great detail, structures and mechanisms for short-term memory and long-term memory, and the way chunks, semantic knowledge, and retrieval structures are acquired through learning. Each cognitive process has a time cost, which makes it possible to simulate DD's performance with great precision. The simulations showed that the model successfully accounts for how DD acquired expertise in this domain; indeed, the model was able to capture his development both qualitatively and quantitatively.

Bounded Rationality and Heuristics

During his career, Simon had amassed considerable evidence that human rationality is bounded. Bounded rationality does not mean that humans are irrational, but that humans are rational within the confines of their computational capabilities. An important means to reach this rationality is by taking advantage of the statistical structure of the task environment and extracting regularities through learning. The importance of learning is amply supported by data from expertise research and is perhaps most apparent in the central role played by pattern recognition in expertise. However, using such regularities is not enough, it is still important to carry out search through the space of promising options. As systematic search is not possible due to the limits of human cognition (in particular the limits of short-term memory), searching through the problem space must be selective and guided by heuristics. Again, research into expertise supports this hypothesis and selective search is a constant theme, as has been illustrated several times in this chapter.

In order to provide a complete theory of bounded rationality, including the way experts manage to circumvent their limited computational capabilities, one needs to provide mechanisms both for pattern recognition and search. The models developed by Simon to account for expertise provide examples of such mechanisms. Selective search is sometimes guided by pattern recognition, and it is sometimes guided by heuristics. Thus, to become an expert it is essential to acquire powerful heuristics, and experts in science are no exception.

In writings and talks he would give around the CMU campus, Simon provided heuristics to help students and colleagues succeed in science. Many showed how selective attention can be a powerful tool to deal with the limits imposed by bounded rationality. Here is a small sample of Simon's heuristics that have impressed me (for additional examples, see Valdes-Perez, 2002, Langley, 2002). 'What is worth doing is worth doing badly. Carry out your research diligently, but not more so than necessary.' This is a direct application of the notion of satisficing, optimal solutions are out of reach and one has to content oneself with good-enough solutions. To Simon, 'a PhD thesis is only a progress report.' This is a particularly useful heuristic since many students are paralyzed by the myth that a PhD thesis has to make a major contribution to its field of research. In some cases, this paralysis is so severe that the PhD is never finished. How did Simon find the time to master numerous scientific domains? 'Your time is precious. Don't waste it by reading newspapers and watching TV. If something really important happens in the world, you'll know it through your friends.' Some heuristics dealt with the content of research: 'Choose important but also realistic research questions;' and 'Play with your knowledge, explore unexpected connections.' Perhaps his most powerful heuristic was to be surrounded by collaborators and friends. I was very fortunate that Simon used this heuristic with me.

References

Baylor, G. W. and Simon, H. A. (1966). A Chess Mating Combinations Program. *1966 Spring Joint Computer Conference.* Boston.

Bhaskar, R. and Simon, H. A. (1977). Problem Solving in Semantically Rich Domains: An Example from Engineering Thermodynamics. *Cognitive Science,* 1, 193–215.

Bradshaw, G., Langley, P. W. and Simon, H. A. (1983). Studying Scientific Discovery by Computer Simulation. *Science,* 222, 971–75.

Chase, W. G. and Ericsson, K. A. (1982). Skill and Working Memory. *The Psychology of Learning and Motivation,* 16, 1–58.

Chase, W. G. and Simon, H. A. (1973a). The Mind's Eye in Chess. In *Visual Information Processing,* ed. W.G. Chase. New York: Academic Press.

Chase, W. G. and Simon, H. A. (1973b). Perception in Chess. *Cognitive Psychology,* 4, 55–81.

De Groot, A. D. (1946). *Het Denken van den Schaker,* Amsterdam, Noord Hollandsche.

Feigenbaum, E. A. and Simon, H. A. (1984). EPAM-like Models of Recognition and Learning. *Cognitive Science,* 8, 305–36.

Gobet, F. (1998). Expert Memory: A Comparison of Four Theories. *Cognition,* 66, 115–52.

Gobet, F. (2013). Chunks and Templates in Semantic Long-term Memory: The Importance of Specialization. In *Expertise and Skill Acquisition: The Impact of William G. Chase,* ed. J. J. Staszewski. New York: Psychology Press.

Gobet, F. (2015). *Understanding Expertise,* London, Palgrave Macmillan.

Gobet, F., Lane, P. C. R., Croker, S., Cheng, P. C. H., Jones, G., Oliver, I. and Pine, J. M. (2001). Chunking Mechanisms in Human Learning. *Trends in Cognitive Sciences,* 5, 236–43.

Gobet, F., Richman, H. B., Staszewski, J. J. and Simon, H. A. (1997). Goals, Representations, and Strategies in a Concept Attainment Task: The EPAM Model. *The Psychology of Learning and Motivation,* 37, 265–90.

Gobet, F. and Simon, H. A. (1996a). Recall of Random and Distorted Positions. Implications for the Theory of Expertise. *Memory and Cognition,* 24, 493–503.

Gobet, F. and Simon, H. A. (1996b). Recall of Rapidly Presented Random Chess Positions is a Function of Skill. *Psychonomic Bulletin and Review,* 3, 159–63.

Gobet, F. and Simon, H. A. (1996c). The Roles of Recognition Processes and Look-ahead Search in Time-constrained Expert Problem Solving: Evidence from Grandmaster Level Chess. *Psychological Science,* 7, 52–5.

Gobet, F. and Simon, H. A. (1996d). Templates in Chess Memory: A Mechanism for Recalling Several Boards. *Cognitive Psychology,* 31, 1–40.

Gobet, F. and Simon, H. A. (1998). Expert Chess Memory: Revisiting the Chunking Hypothesis. *Memory,* 6, 225–55.

Gobet, F. and Simon, H. A. (2000). Five Seconds or Sixty? Presentation Time in Expert Memory. *Cognitive Science,* 24, 651–82.

Hayes, J. R. and Simon, H. A. (1974). Understanding Written Instruction. In *Knowledge and Cognition,* ed. L. W. Gregg. Hillsdale, NJ: Erlbaum.

Langley, P. (2002). Heuristics for Scientific Discovery: The Legacy of Herbert Simon. In *Models of a Man: Essays in Memory of Herbert A. Simon*, ed. E. Augier and J.G. March, J. G. Boston: MIT Press.

Langley, P., Simon, H. A., Bradshaw, G. L. and Zytkow, J. M. (1987). *Scientific Discovery,* Cambridge, MA, MIT Press.
Larkin, J., Mcdermott, J., Simon, D. P. and Simon, H. A. (1980). Expert and Novice Performance in Solving Physics Problems. *Science,* 208, 1335–42.
Larkin, J. H. and Simon, H. A. (1981). Learning through Growth of Skill in Mental Modeling. *Proceedings of the Third Annual Conference of the Cognitive Science Society.* Berkeley: Cognitive Science Society.
Newell, A., Shaw, J. C. and Simon, H. A. (1958). Chess-playing Programs and the Problem of Complexity. *IBM Journal of Research and Development,* 2, 320–35.
Newell, A., Shaw, J. C. and Simon, H. A. (1962). The Process of Creative Thinking. In *Contemporary Approaches to Creative Thinking, ed.* Gruber, H. E., Terrell, G. and Werheimer. New York: Atherton Press.
Newell, A. and Simon, H. A. (1965). An Example of Human Chess Play in the Light of Chess-playing Programs. In *Progress in Biocybernetics, ed.* Weiner, N. and Schade, J. P. Amsterdam: Elsevier.
Newell, A. and Simon, H. A. (1972). *Human Problem Solving,* Englewood Cliffs, NJ, Prentice-Hall.
Qin, Y. and Simon, H. A. (1990). Laboratory Replication of Scientific Discovery Processes. *Cognitive Science,* 14, 281–312.
Richman, H. B., Gobet, F., Staszewski, J. J. and Simon, H. A. (1996). Perceptual and Memory Processes in the Acquisition of Expert Performance: The EPAM Model. In *The Road to Excellence, ed.* Ericsson, K. A. Mahwah, NJ: Erlbaum.
Richman, H. B., Staszewski, J. J. and Simon, H. A. (1995). Simulation of Expert Memory with EPAM IV. *Psychological Review,* 102, 305–30.
Simon, H. A. (1947a). *Administrative Behavior,* New York, NY, Macmillan.
Simon, H. A. (ed.) (1947b). *The Technique of Municipal Administration* (3rd edn), Chicago: International City Managers' Association.
Simon, H. A. (1977). *Models of Discovery and Other Topics in the Methods of Science,* Dordrecht, Reidel.
Simon, H. A. (1991). *Models of My Life,* New York, Basic Books.
Simon, H. A. and Barenfeld, M. (1969). Information Processing Analysis of Perceptual Processes in Problem Solving. *Psychological Review,* 7, 473–83.
Simon, H. A. and Chase, W. G. (1973). Skill in Chess. *American Scientist,* 61, 393–403.
Simon, H. A. and Feigenbaum, E. A. (1964). An Information Processing Theory of Some Effects of Similarity, Familiarity, and Meaningfulness in Verbal Learning. *Journal of Verbal Learning and Verbal Behavior,* 3, 385–96.
Simon, H. A. and Gilmartin, K. J. (1973). A Simulation of Memory for Chess Positions. *Cognitive Psychology,* 5, 29–46.
Tabachnek-Schijf, H. J. M., Leonardo, A. M. and Simon, H. A. (1997). CaMeRa: A Computational Model of Multiple Representations. *Cognitive Science,* 21, 305–50.
Valdes-Perez, R. E. (2002). Personal Recollections from 15 Years of Monthly Meetings. In *Models of a Man: Essays in Memory of Herbert A. Simon, ed.* Augier, E. and March, J. G. Boston: MIT Press.
Zhang, G. and Simon, H. A. (1985). STM Capacity for Chinese Words and Idioms: Chunking and Acoustical Loop Hypothesis. *Memory and Cognition,* 13, 193–201.