

Is Experts' Knowledge Modular?

Fernand Gobet (frg@psyc.nott.ac.uk)
ESRC CREDIT
School of Psychology
University of Nottingham
Nottingham NG7 2RD, UK

Abstract

This paper explores, both with empirical data and with computer simulations, the extent to which modularity characterises experts' knowledge. We discuss a replication of Chase and Simon's (1973) classic method of identifying 'chunks', i.e., perceptual patterns stored in memory and used as units. This method uses data about the placement of pairs of items in a memory task and consists of comparing latencies between these items and the number and type of relations they share. We then compare the human data with simulations carried out with CHREST, a computer model of perception and memory. We show that the model, based upon the acquisition of a large number of chunks, accounts for the human data well. This is taken as evidence that human knowledge is organised in a modular fashion.

Introduction

An important goal of cognitive science is to understand the characteristics of knowledge, in particular the way it is acquired and used. To achieve this goal, research has employed a number of methods, including artificial laboratory experiments, such as nonsense syllable learning, and collection of naturalistic data, such as experts functioning in their natural environments. It is generally accepted that knowledge consists of different types (declarative, procedural, episodic) and that its acquisition follows a power law of learning. In addition, it has often been proposed that knowledge is modular, consisting, for example, of productions (e.g., Newell, 1990) or of perceptual chunks (e.g., Chase & Simon, 1973).

The goal of this paper is to explore, both with empirical data and with computer simulations, the extent to which modularity characterises human knowledge, and in particular experts' knowledge. We first describe the concept of modularity, and then show how it has been used in expertise research. This leads us to describe the CHREST architecture, which acquires knowledge by growing a discrimination net encoding chunks. Next, we present data aimed at characterising the properties of experts' chunks, and compare them with those acquired by CHREST. The comparison results in an excellent fit between the model and the human data. In the conclusion, implications for the modularity of human knowledge in general are drawn.

Modularity of Knowledge

Several formalisms, both modular and non-modular, have been developed in cognitive science to explain how humans represent and implement knowledge. Examples of modular representations are production systems, semantic networks, and discrimination nets. Examples of non-modular representations are distributed neural networks, holograms, and various mathematical representations based on matrix algebra. This classification should be considered with caution, however. On the one hand, production rules, for example, are typically organised in problem spaces (e.g., Newell, 1990), and their interdependence can be considerable, which counts against strict modularity. On the other hand, it could be argued that, in non-modular representations, modules emerge as the system develops or learns (e.g., Rumelhart & McClelland, 1986).

Modular knowledge organisation has attracted much interest in computer science and artificial intelligence, given the importance of how knowledge is indexed, structured, organised, and retrieved (e.g., Lane et al., 2000). In artificial intelligence, modularity has been defined as "the ability to add, modify, or delete individual data structures more or less independently of the remainder of the database, that is, with clearly circumscribed effects on what the system 'knows'" (Barr & Feigenbaum, 1989, p. 149). While a strong argument can be made that it is easier to understand modular and decomposable systems than systems that do not share these properties (e.g., Simon, 1969), and that the value of these properties has been demonstrated in fields such as software engineering, it is an empirical question whether human knowledge is modular or not. A rich source of data about this question has been gained from research into expert behaviour, to which we now turn our attention.

Chess Experts' Knowledge

In his seminal study, De Groot (1946/1965) subjected chess players to a number of problem-solving and memory experiments. The surprising result was that, in a choice-of-a-move task, there was no large skill difference in variables such as depth of search, number of moves considered, or search heuristics employed. However, a clear difference was found in a memory task where a chess position was presented for a few seconds. Masters could recall the entire position almost perfectly,

while weaker players could recall only a handful of pieces. De Groot concluded that expertise does not reside in any superior abilities but in knowledge.

Continuing de Groot's research, Chase and Simon (1973) carried out a study destined to have a huge impact in cognitive science. They used two tasks. In the *recall task*, based on de Groot's (1965) method, a chess position was presented for five seconds, and players had to reconstruct as many pieces as possible. In the *copy task*, the stimulus board remained in view, and the goal was to reconstruct it onto a second, empty board. As the stimulus and the reconstruction boards could not be fixated simultaneously, Chase and Simon used the glances between the boards to detect memory chunks. Comparing the latencies between successive pieces in the copy and recall tasks, they inferred that pieces replaced with less than 2 seconds' interval belonged to the same chunk, and that pieces placed with an interval of more than 2 seconds belonged to different chunks. Finally, they showed that the chunk definition based upon the latencies between two successive pieces was consistent with a definition based upon the pattern of semantic relations (attack, defence, proximity, colour, and type of piece) shared by these two pieces. This converging evidence was used to infer the chunks used to mediate superior performance, and to explore how they allowed masters to find good moves despite their highly selective search. A number of other experimental tasks (reviewed in Gobet & Simon, 1998) have brought converging evidence for the psychological reality of chunks, as defined either by latency in placement or by number of relations between pieces.

Simon and Gilmarin (1973) developed a computer program (MAPP; Memory-Aided Pattern Perceiver) implementing some of Chase and Simon's ideas. MAPP is based upon EPAM (Elementary Perceiver and Memorizer; Feigenbaum & Simon, 1984), a theory developed to account for empirical phenomena where chunking (i.e., acquisition of perceptual units of increasing size) is seen as essential. The basic idea in MAPP was that long-term memory (LTM) is accessed through a discrimination net, and that, once elicited, LTM chunks are stored in short-term memory (STM) through a pointer. MAPP's relatively low recall performance—slightly better than a good amateur, but inferior to an expert—was attributed to the small number of nodes, about two thousand, stored in its LTM. MAPP simulated several results successfully: increase in performance as a function of the number of chunks in LTM; kind of pieces replaced; and contents of chunks. However, in addition to its failure in simulating expert behaviour, the program had several limitations (De Groot & Gobet, 1996). In particular, the chunks were chosen by the programmers and not autonomously learnt, and the program made incorrect predictions for a number of experiments that were later carried out. These limitations were removed in the CHREST program discussed below.

CHREST

CHREST (Chunk Hierarchy and REtrieval STructures; De Groot & Gobet, 1996; Gobet & Simon, 2000) is a cognitive architecture similar to MAPP. CHREST originally addressed high-level perception, learning and memory, but various problem-solving mechanisms have been implemented recently. It is composed of processes for acquiring low-level perceptual information, an STM, attentional mechanisms, a discrimination net for indexing items in LTM, and mechanisms for making associations in LTM such as production rules or schemas. STM mediates the flow of information processing between the model's components. The central processing of CHREST revolves around the acquisition of a discrimination net based on high-level perceptual features picked up by attentional mechanisms and on the creation of links connecting nodes of this net together.

After the simulated eye has fixated on an object, features are extracted and processed in the discrimination net, and then, based upon the output of the discrimination, a further eye fixation is made, and so on. STM operates as a queue; that is, the first elements to enter are also the first to leave. STM has a limited capacity, which consists of four chunks (Cowan, 2001; Gobet & Simon, 2000). Processing is constrained by a number of restrictions, including time parameters such as the time to fixate a chunk in LTM (8 s) and capacity parameters such as the four-chunk limit of STM.

The discrimination net consists of *nodes*, which contain *images* (i.e., the internal representation of the external objects; images correspond to Chase and Simon's *chunks*); the nodes are interconnected by *links*, which contain *tests* allowing items to be sorted through the net. Learning happens as follows: once an item has been sorted through the net, it is compared to the image in the node reached. If the item and image agree but there is more information in the item than the image, then *familiarisation* occurs, in which further information from the item is added to the image. If the item and image disagree in some feature, then *discrimination* occurs, in which a new node and a new link are added to the net. Based on empirical data, it has been estimated that discrimination requires about 8 s and familiarisation about 2 s.

In addition to these learning mechanisms, CHREST has mechanisms for augmenting semantic memory by the creation of schemas (known as *templates*) and of *lateral* links connecting nodes together (Gobet, 1996); for example, these links can be created when nodes are sufficiently similar ('similarity links'), or when one node can act as the condition of another node ('condition links'). The creation of these links is consistent with the emphasis on processing limits present in both EPAM and CHREST, in that all nodes used for creating new links must be in STM.

Table 1: Copy, recall and *a priori* chess relations probabilities, for combinations of the five chess relations: Attack (A), Defence (D), Spatial Proximity (P), Same Colour (C), and Same Piece (S).

Relations	COPY				RECALL				<i>A priori</i> Probabilities	
	GAME		RANDOM		GAME		RANDOM		GAME	RANDOM
	WITHIN	BETWEEN	WITHIN	BETWEEN	≤ 2 sec	> 2 sec	≤ 2 sec	> 2 sec		
-	.037**	.172**	.086**	.129**	.052**	.190**	.051**	.284	.335	.297
A	.005**	.006	.031	.054**	.004**	.024	.000*	.054	.016	.024
P	.000	.006	.037**	.059**	.001	.006	.033**	.041*	.004	.010
C	.148**	.278	.152**	.203**	.132**	.247	.136**	.189	.255	.297
S	.016**	.056**	.040**	.049**	.040**	.102*	.059**	.054	.154	.144
AP	.000*	.000	.056**	.069**	.001	.003	.015	.027	.005	.028
AS	.000	.000	.003	.005*	.004**	.003	.000	.000	.001	.001
DC	.104**	.133**	.072**	.077**	.059**	.084**	.044	.068*	.035	.024
PC	.084**	.067**	.059**	.046**	.049**	.060**	.066**	.081**	.019	.009
PS	.002	.006	.044**	.064**	.006	.012	.018	.027	.006	.010
CS	.115	.094	.135*	.105	.111	.057*	.059*	.041	.096	.108
APS	.000	.000	.013*	.013	.001	.000	.018	.014	.001	.007
DPC	.109**	.078	.123**	.064**	.093**	.084*	.118**	.081*	.048	.028
DCS	.048**	.017**	.000	.000	.033**	.012**	.015**	.000	.002	.001
PCS	.196**	.039**	.127**	.039**	.202**	.060**	.232**	.041**	.011	.007
DPCS	.137**	.050**	.023**	.023**	.213**	.054**	.136**	.000	.013	.007
#observations	1283	180	1114	389	1563	332	272	74		

Note: * means $p < .01$, ** means $p < .001$ (both two-tailed). The statistical significance levels are based on the z -values that were computed using the following formula (assuming the normal approximation to the binomial distribution):

$$z = \frac{p_0 - p_e}{s. e.}, \quad \text{where } s. e. = \sqrt{\frac{p_e (1 - p_e)}{\text{sample size}}},$$

and where p_0 is the observed probability and p_e the *a priori* (expected) probability.

CHREST can reproduce a number of features of the behaviour of skilled and unskilled chess players in memory experiments, such as their eye movements, the size and number of chunks, the number and type of errors, and the differential recall of game and random positions (De Groot & Gobet, 1996; Gobet & Simon, 2000). As a psychological theory, CHREST has several strengths. It is parsimonious, with few free parameters. It provides absolute quantitative predictions, for example about the number of errors committed or the time taken by a subject to carry out a task. Together with EPAM, it simulates in detail a number of empirical phenomena from various domains, such as verbal learning, context effects in letter perception, concept formation, expert behaviour, acquisition of first language by children, and use of multiple representations in physics (see Gobet et al., in press, for a review).

A Replication of Chase and Simon (1973)

As noted above, Chase and Simon (1973) operationalised the concept of chunk using both the latencies between successive piece placements and the semantic

relations between them. Their experiment has recently been replicated and extended by Gobet and Simon (1998). The main difference between the two studies is that Gobet and Simon used a computer display to present the tasks instead of physical chessboards. In spite of this difference, there is an important overlap between the results of the two studies.

Gobet and Simon analysed 26 players (Chase and Simon had only 3) ranging from good amateurs to professional grandmasters, who were divided into three skill levels (Masters, Experts and Class A players). The results were in line with previous experiments, showing a massive skill effect with game position, and a small but reliable skill effect even with meaningless positions. Here, we focus upon the operationalisation of chunks, relying both upon Gobet and Simon's published data and upon additional analyses.

Latencies Predict Chunk Boundaries

Gobet and Simon essentially followed Chase and Simon's approach. They first estimated a time threshold (2 s) as a means to decide whether two pieces placed in succession belonged to the same chunk, and then

validated this threshold by showing that it led, on average, to similar chunks as those obtained by using semantic relations. If they are modular, chunks should be characterised by a high density of relations between the elements that constitute it, and by a low density of relations with elements from other chunks (Chase & Simon, 1973; Cowan, 2001). That is, there should be many more relations between successive pieces within the same chunk than between successive pieces on opposite sides of a chunk boundary. Thus, the relations between successively replaced pieces should be different depending on whether they are separated by short or long latencies. In addition, assuming that the same cognitive mechanisms mediate the latencies in the copy and recall experiments, the two experiments should show the same pattern of interaction between latencies and number of relations. In other words, the relations for the within-glance placements in the copy task should correlate with those for rapid placements (≤ 2 s) in the recall task and the relations for between-glance placements in the former should correlate with those for slow placements (> 2 s), in the latter.

These predictions are met in both the copy and the recall tasks, whose results correlate highly. Within chunks, small latencies correlate with a large number of relations, while large latencies occur when there are few relations between successive pieces. No such relationship is observed for successive pieces belonging to different chunks. The shortest latencies are found with four relations (Defence, Proximity, Colour, and Kind), which mainly occur with pawn formations.

Relations Predict Chunk Boundaries

The next step consists in showing that the pattern of relation probabilities for within-chunk, but not for between-chunk placements, differs from what could be expected by chance. Table 1 gives the probabilities of the presence of different combinations of relations in the various experimental conditions, with the three skill levels pooled. The last two columns give the *a priori* probabilities (for game and random positions, respectively) that were calculated by recording, for each position, all relations that exist between all possible pairs of pieces; the *a priori* probability for a relation is obtained by dividing the total number of occurrences of a relation by the total number of possible pairs. These *a priori* probabilities were based on 100 positions and 26,801 pairs. Finally, the *z*-values indicate whether the observed probabilities reliably differ from the *a priori* probabilities.

In the copy task, with game positions but not with random positions,¹ the between-glance probabilities are much closer to chance than the within-glance probabilities. This pattern holds also in the recall of both ran-

¹That this pattern does not hold with the copy of random positions may be due to the strategy used by subjects to replace these positions. Several subjects copied the positions line by line or column by column.

dom and game positions when slow placements (> 2 s) are compared with fast placements (≤ 2 s). The probabilities for pieces with three and four relations are high in the within-glance and fast (≤ 2 s) conditions compared with the between-glance and slow (> 2 s) conditions; the opposite is true for pieces with one relation or none. Note also that the probabilities for combinations of relations that include an attack (A) are conspicuously low, compared with chance, for game positions but not for random positions.

One way to make sense of Table 1 is to analyse the correspondence between the number of chess relations and the deviations from *a priori* probabilities, computed by subtracting the *a priori* probabilities from the observed frequencies of a given condition. Based on the notion of modularity, it should be expected that the within-chunk deviations from *a priori* probabilities would be highly correlated with the number of relations, while this would not be the case for the between-chunk deviations. This is exactly what was found. The correlations with number of relations are high for the within-chunk conditions (copy game within-glance: 0.81; copy random within-glance: 0.68; recall game short latencies: 0.86; recall random short latencies: 0.79; all the correlations are statistically significant at $p = .005$). The correlations are smaller with the between-chunk conditions (copy game between-glance: 0.61; copy random between-glance: 0.56; recall game long latencies: 0.58; recall random long latencies: -0.15; none of the correlations are significant at the .01 level). These results are illustrated graphically in Figure 1, which shows the results for game and random positions as a function of whether the placements were within-chunk or between-chunk. From the Figure, it is clear that, for within-chunk conditions, the placements having few relations are below chance, while the placements having several relations are above chance. There is no such clear relation for the between-chunks placements.

Computer Simulations

We now show that CHREST captures the composition of chunks and the pattern of relations of within- and between-chunk placements. Simulations of similar phenomena, carried out by Simon and Gilmartin (1973) using MAPP, were limited to a single subject and matched the data only approximately.

Methods

In the *learning phase*, the program scanned a large database of master-game positions, fixating squares with simulated eye movements, and learning chunks using discrimination and familiarisation. Three nets were created, estimated to correspond roughly to the recall percentages of Class A players, experts, and masters with a five-second presentation time. These nets had respectively 1,000 nodes, 10,000 nodes, and 100,000 nodes.

For the simulations of the *performance phase*, the program was tested with 100 game positions and 100

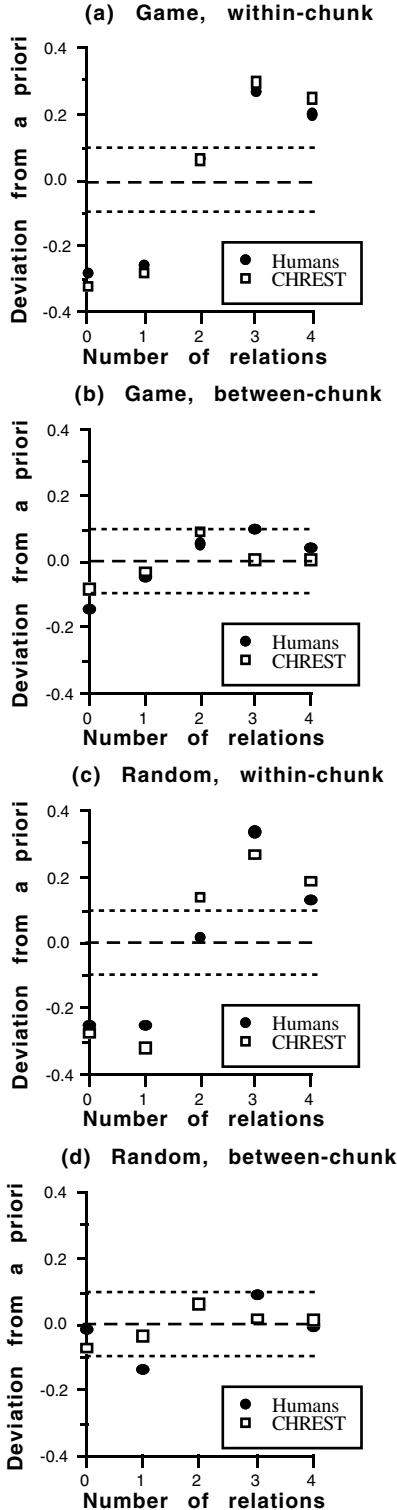


Figure 1: Relation between chess relation probabilities and the number of relations shared by two pieces successively placed. The long-dash line indicates zero deviation, and the short-dash lines indicate deviations of 0.1 above or below zero.

random positions. Learning was turned off. During the five-second presentation of a position, CHREST moved its simulated eye around the board. Each eye fixation defined a visual field (all squares within two squares from the square fixated); the pieces within the visual field are treated as a single pattern and sorted through the discrimination net. Other patterns are defined by the pieces focused upon in two successive eye fixations. If a chunk is found in the discrimination net, a pointer to it is placed in STM.

During the reconstruction of a position, CHREST used the information stored in STM. When a piece belonged to several chunks, it was replaced only once. In case of conflicts (e.g., a square is proposed to contain several pieces), CHREST resolved them sequentially, based on the frequency with which each placement is proposed. Like humans, it sometimes made several different proposals about the location of a piece or about the contents of a square. Finally, some weak heuristics were used, such as the fact that only one white king can be replaced in a position. (See Gobet & Simon, 2000, for more detail.)

A chunk refers to the image of a node in the discrimination net. It is therefore straightforward to decide whether two pieces do or do not belong to the same chunk. The relations between pieces were extracted using the same program as that used with the human data.

Results

Table 2 gives the probabilities of observing a pattern of relations, as a function of the type of position and the kind of placement. Although the fit with the corresponding human data shown in Table 1 is reasonable

Table 2. Recall and *a priori* chess relations probabilities, for combinations of the five chess relations: Attack (A), Defence (D), Spatial Proximity (P), Same Colour (C), and Same Piece (S).

Relations	Game positions			Random positions		
	With in	Between	A priori	With in	Between	A priori
-	.009	.254	.335	.018	.231	.297
A	.005	.034	.016	.021	.061	.024
P	.013	.011	.004	.050	.026	.010
C	.104	.208	.255	.040	.216	.297
S	.021	.148	.154	.050	.136	.144
AP	.004	.013	.005	.030	.027	.028
AS	.000	.001	.001	.001	.005	.001
DC	.042	.059	.035	.038	.042	.024
PC	.097	.050	.019	.092	.039	.009
PS	.020	.019	.006	.061	.018	.010
CS	.064	.113	.096	.094	.111	.108
APS	.004	.005	.001	.008	.017	.007
DPC	.162	.031	.048	.148	.033	.028
DCS	.007	.000	.002	.009	.001	.001
PCS	.186	.032	.011	.147	.015	.007
DPCS	.259	.021	.013	.193	.023	.007

(the r^2 are: game within-chunk: .83; game between-chunk: .82; random within-chunk: .58; random between-chunk: .75), not too much weight should be given to them, because they are sensitive to a few large values, and because they may in part reflect the statistics of the chess environment (i.e., the *a priori* probabilities). As with the human data, we subtracted the *a priori* probabilities from the recall probabilities, and took the sum for each number of relations. Figure 1 shows the results for both the humans and CHREST. The model fits the human data quite well. In particular, the between-chunk placements show little deviation from the *a priori* probabilities, in contrast to the within-chunk placements, which are clearly below chance with zero and one relation, and above chance with three and four relations. All conditions pooled, CHREST accounts for 90% of the variance of the human data.

Conclusion

EPAM and CHREST's learning mechanisms, based upon the construction of a discrimination net of chunks, offer a crisp and computational definition of the concept of knowledge module. Using this definition, Chase and Simon (1973) have found, and Gobet and Simon (1998) have confirmed, that relations and latencies between pieces offer converging evidence for validating the psychological reality of chunks. This paper has shown that, with the same mechanisms used to account for a variety of chess data, CHREST acquires chunks that have the same relational properties as humans'.

The acquisition mechanisms consisting in learning pieces within the visual field and between two eye fixations largely explain the high number of relations within chunks. It is important to note that this phenomenon is not trivial to simulate, however. For example, learning mechanisms such as Saariluoma and Laine's (2001) frequency-based heuristic, where chunk construction is not constrained by spatial contiguity, would fail to account for the data, because they do not capture the relation of proximity which is essential in the chunks acquired by humans (cf. Table 1).

These results, as well as others, indicate that the modular structure of the type of discrimination net used by EPAM and CHREST captures essential aspects of human cognition. Chunks, whose elements share a number of relations, are built up gradually and recursively, with later chunks being built from smaller 'sub-chunks'. Some of these chunks evolve into schema-like structures, and some get later connected by lateral links, thereby constructing both a net of productions and a semantic network. It is not only the presence of a node storing a piece of knowledge which matters, but also the richness with which this node is perceptually indexed and the density with which this node is connected to other nodes. These two aspects give some computational meaning to "conceptual understanding": a richly-connected network of links connecting productions and schemas, that is accessible through perceptual

chunks. In addition to expert behaviour, CHREST, which incorporates mechanisms for all these kinds of learning, including the acquisition of modular structures, accounts for empirical phenomena in a variety of domains.

Acknowledgements

I am grateful to Herb Simon for his involvement in many aspects of this research and to the members of the CHREST group for comments on this paper.

Reference List

Barr, A., & Feigenbaum, E. A. (1989). *The handbook of artificial intelligence*. (Vol. 1). New York: Addison-Wesley.

Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, 4, 55-81.

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24.

de Groot, A. D. (1965). *Thought and choice in chess*. The Hague: Mouton. (First edition in Dutch, 1946).

de Groot, A. D., & Gobet, F. (1996). *Perception and memory in chess*. Assen: Van Gorcum.

Feigenbaum, E. A., & Simon, H. A. (1984). EPAM-like models of recognition and learning. *Cognitive Science*, 8, 305-336.

Gobet, F. (1996). Discrimination nets, production systems and semantic networks: Elements of a unified framework. *Proceedings of the 2nd International Conference on the Learning Sciences* (pp. 398-403). Evanston IL: Northwestern University.

Gobet, F., Lane, P. C. R., Croker, S., Cheng, P.C-H., Jones, G., Oliver, I., & Pine, J. M. (in press). Chunking mechanisms in human learning. *Trends in Cognitive Science*.

Gobet, F., & Simon, H. A. (1998). Expert chess memory: Revisiting the chunking hypothesis. *Memory*, 6, 225-255.

Gobet, F., & Simon, H. A. (2000). Five seconds or sixty? Presentation time in expert memory. *Cognitive Science*, 24, 651-682.

Lane, P. C. R., Gobet, F. , & Cheng, P. C-H. (2000). Learning-based constraints on schemata. *Proceedings of the 22nd Meeting of the Cognitive Science Society*. (pp. 776-782). Mahwah, NJ: Erlbaum.

Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.

Rumelhart, D. E., & McClelland, J. L. (1986). *Parallel distributed processing*. Cambridge, MA: MIT Press.

Saariluoma, P. & Laine, T. (2001). Novice construction of chess memory. *Scandinavian Journal of Psychology*, 42, 137-147.

Simon, H. A. (1969). *The sciences of the artificial*. Cambridge: MIT Press.

Simon, H. A., & Gilmarin, K. (1973). A simulation of memory for chess positions. *Cognitive Psychology*, 5, 29-46.