

Gobet, F. (2000). Long-term working memory: A computational implementation for chess expertise. *Proceedings of the 3rd International Conference on Cognitive Modelling*, pp. 142-149. Veenendaal, The Netherlands: Universal Press.

Long-Term Working Memory: A Computational Implementation for Chess Expertise

Fernand Gobet (frg@psyc.nott.ac.uk)

ESRC CREDIT

Department of Psychology

University of Nottingham

Nottingham NG7 2RD, U. K.

Abstract

Long-term working memory (Ericsson and Kintsch, 1995) is a theory covering empirical data from several domains, including expert behaviour. One difficulty in applying and evaluating this theory, however, is that it is framed in rather general terms, and that several mechanisms and parameters are left unspecified. This paper proposes a computer implementation of the theory for a domain that Ericsson and Kintsch cover in depth, namely chess memory. Simulations of Saariluoma's (1989) experiment where both game and random chess positions are presented auditorily make it possible to analyse two key ingredients of the theory: encoding through elaboration of LTM schemas and patterns, and encoding through retrieval structures. In the simulations, these mechanisms were modulated by two parameters. The results show that random positions, but not game positions, are sensitive to these parameters' values.

The study of expert behaviour is currently an important area of research in cognitive science. Several theories, including the chunking theory (Chase & Simon, 1973), the skilled memory theory (Chase & Ericsson, 1982), Soar (Newell, 1990), ACT (Anderson, 1983), and the template theory (Gobet & Simon, in press) have been advanced to explain how certain individuals excel in their domain of expertise. Recently, an important attempt to offer an integrative theory of cognition and expertise has been proposed by Ericsson and Kintsch (1995) with their long-term working memory (LT-WM) theory.

Long-Term Working Memory: Overview

After a detailed review of research on expertise and on text comprehension, Ericsson and Kintsch (1995) conclude that experts in various fields can encode information into long-term memory (LTM) more rapidly than had been postulated by traditional models of human memory, such as those of Anderson (1983) or Chase and Simon (1973). Based upon this analysis, Ericsson and Kintsch (1995) extend Chase and Ericsson's (1982) skilled memory theory into the LT-WM theory. The tenets of LT-WM are that "cognitive processes are viewed as a sequence of stable states representing end products of processing" and that "acquired memory skills allow these end products to be stored in long-term memory and kept directly accessible by means of retrieval cues in short-term memory [...]" (Ericsson & Kintsch, 1995, p. 211). Two intertwined mechanisms allow rapid storage into LTM: (a) encoding through a retrieval structure, and (b) encoding through knowledge-based associations connecting items either to other items or to LTM patterns and schemas, which allows for an integrated representation of the information in LTM. Ericsson and Kintsch also note that the demands that the task makes on memory will constrain which encoding mechanism will be employed, and that the relative roles played by encoding through retrieval structures or through LTM elaborations vary from task to task.

An important component of LT-WM is the concept of retrieval structure. Retrieval structures are "a set of retrieval cues [that] are organised in a stable structure" (Ericsson & Kintsch, 1995, p. 216). It is assumed that, through practice and study, experts develop such structures for their domain of expertise. Perhaps the strongest empirical support for retrieval structures comes from the extremely detailed analysis of how SF and DD, otherwise unremarkable college students, became world-class experts in the digit-span task (Chase & Ericsson, 1982; Richman et al., 1995). The LT-WM proposal that SF and DD store groups of digits and additional semantic information in a hierarchical retrieval structure is well supported by a large set of data including reaction times, verbal protocols, and direct experimental manipulations. However, Ericsson and Kintsch also argue that retrieval structures alone are not sufficient, and that they need to be supplemented by

LT-WM's second mechanism, elaboration of previously stored semantic memory; this elaboration facilitates the encoding of information by providing redundant cues (see Figure 1).

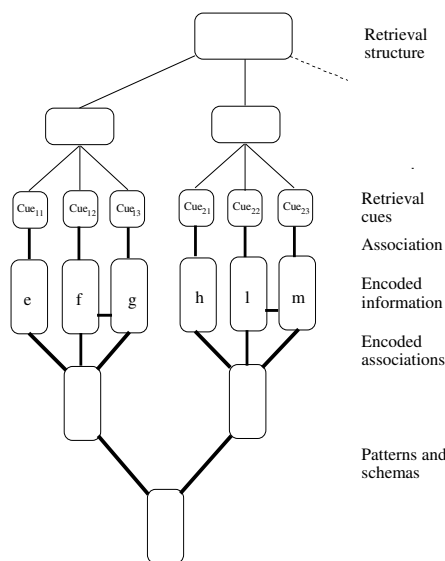


Figure 1. Illustration of the two types of LT-WM encodings. The upper part depicts a hierarchical organisation of retrieval cues associated with units of encoded information. The lower part shows knowledge-based associations relating units of encoded information to each other along with patterns and schemas. (Adapted from Ericsson and Kintsch, 1995.)

LT-WM conflicts with most traditional theories of working memory, which emphasise its transitory storage capacity. Given that LT-WM unifies empirical results from domains that are normally treated separately, it is quite impressive in its scope. However, it has also been criticised on several grounds (e.g., Gobet, 1998): LT-WM in general, and the concept of retrieval structure in particular, are not sufficiently well specified to allow precise empirical predictions; the concept of retrieval structure is used inconsistently in the theory; and the empirical support for LT-WM is weaker than claimed by Ericsson and Kintsch.

Application of long-term working memory to chess expertise

Chess has been a very popular topic for developing theories of expertise, and has played an important role in the development of theories such as the chunking theory, the template theory, and the constraint attunement theory. Two advantages of chess, as compared to other domains of expertise, are that it already offers a large amount of empirical data and that its rating scale allows a precise and quantitative measure of players' skill. Charness (1992) goes so far as to propose that chess is a model task in the study of expertise. In their paper on LT-WM, Ericsson and Kintsch (1995) also highly value results from chess expertise and state that "research on planning and memory of chess positions offers some of the most compelling evidence for LT-WM" (p. 238). Chess is therefore a suitable domain to investigate both the strengths and the weaknesses of LT-WM.

Ericsson and Kintsch (1995) suggest that strong chess players use a hierarchical retrieval structure corresponding to the 64 squares of the chess board. This structure, which both associates individual pieces to their respective locations and relates pieces to each other, allows a position to be represented as an integrated hierarchical structure. This structure allows a rapid encoding into LTM, where patterns and schemas are stored. Patterns and schemas can also encode new information rapidly, independently of the use of the retrieval structure, through knowledge-based associations between items and/or patterns and schemas. Ericsson and Kintsch suggest that the concept of retrieval structure accounts for chess masters' excellent memory for chess material, as well as for their ability to plan and evaluate alternative sequences of moves. They support their claim by considering several types of data, including: recall experiments with game and random positions; experiments where a chess master has to rapidly access the location of pieces of a position he has just memorised; experiments manipulating the way pieces are grouped during presentation; experiments where multiple boards are presented; and anecdotal evidence showing that masters can play

blindfold chess, which requires them to mentally manipulate and update the information contained in a position (or indeed in several).

Limits of LT-WM as an explanation for chess expertise

Ericsson and Kintsch thus suggest that chess provides some of the strongest support for LT-WM. However, as I have noted elsewhere (Gobet, 1998), their account of chess expertise is stated in rather general terms and does not specify many parameters and mechanisms that would seem crucial in accounting for the data. It is the goal of this paper to offer a first step in embodying LT-WM explanation of expertise as a computer program (see Kintsch, 1998, for a similar attempt with text comprehension).

A simulation of Saariluoma's (1989) dictation experiment using LT-WM

The lack of specificity of LT-WM for chess expertise has the consequence that there is an indefinite number of models that satisfy Ericsson and Kintsch's (1995) description. Only some of them correspond to what Ericsson and Kintsch really have in mind. Even so, it is worthwhile to write a computer program that offers a plausible implementation of Ericsson and Kintsch's description, in order to understand how the two critical LTM storage mechanisms (retrieval structures and elaboration of schemas and patterns) interact. One needs more specifications than offered by Ericsson and Kintsch if one wants to write a running computer program; hence, several decisions, sometimes arbitrary, had to be taken.

Saariluoma's (1989) dictation task

Saariluoma's (1989) dictation task has been chosen for illustration, as it is discussed at length in Ericsson and Kintsch. In this experiment, which uses a technique similar to that of the digit-span task, a chess position is dictated auditorily piece-by-piece at a rapid rate (typically, a rate of one piece every 2 seconds). Once all the pieces have been presented, the subject has to recall as many of the pieces as she can. Saariluoma found that masters were able to recall game positions almost perfectly, but that, with random positions, they did not recall more than around 60% (Saariluoma 1989). Using the same technique, Saariluoma (1989) has found that masters had big difficulties at memorising four random positions presented in sequence (around 10% correct per position), while they obtained reasonably good results with game positions (about 60% per position).

In the version of the task modelled here, each piece of a game position or a random position is dictated at a rate of 2 seconds per piece, using the algebraic notation widely used in the chess community (e.g., "white king on g1"). Since the positions had on average 25 pieces, the dictation of one position took 50 seconds.

LT-WM account of the dictation experiment

In their discussion of Saariluoma's (1989) experiment with auditory presentation, Ericsson and Kintsch (1995, p. 237) apply LT-WM as follows:

"If, on the other hand, chess experts had a retrieval structure corresponding to a mental chess board, they could store each piece at a time at the appropriate location within the retrieval structure. After the end of the presentation the experts would be able to perfectly recall the entire position if the presentation rate had been slow enough."

Random material¹ is critical for evaluating the retrieval structure hypothesis, because, with meaningful material, it is difficult to tease apart the role of retrieval structures, on the one hand, and that of patterns and schemas, on the other. Ericsson and Kintsch (1995, p. 237) are very clear about chess experts' ability to store information from random positions: "skilled chess players are able to encode and store the locations of individual chess pieces of a chess position in the absence of meaningful configurations of chess pieces."

Ericsson and Kintsch also state that meaningful patterns of relations between pieces can be encoded with game positions, which makes it possible for the position to be stored as an integrated structure. This is not possible with random positions. The ability to find higher-order relations in game positions, but not in random positions, explains why the former are easier to recall than the latter.

Although it explains why game positions are easier to remember than random positions, this explanation is not quite satisfactory from a theoretical point of view, because Ericsson and Kintsch are explicit that Saariluoma's (1989) experiments were discussed to show that "the ability to store random chess positions

¹ In chess research, random positions are typically created by taking a game position and randomly reassigning its pieces to new squares.

provides particularly strong evidence for the ability to encode individual chess pieces into the retrieval structure” (Ericsson and Kintsch, 1995, p. 237). However, the empirical data hardly support this point: the recall for random positions dictated auditorily is far from perfect even with masters, who do not recall more than about 60% (about 15 pieces) after a 50-second presentation. As noted by Gobet (1998), this level of recall is roughly predicted by the chunking theory (Chase & Simon, 1973), a theory that Ericsson and Kintsch criticise for encoding information into LTM too slowly. It would seem that the importance of encoding into the retrieval structure is overestimated in LT-WM. Computer modelling can help us to disentangle the role of elaboration encoding and retrieval structure.

Components of the model

The model consists of three components: an articulatory loop, LTM patterns and schemas, and a retrieval structure. The articulatory loop is similar to that proposed by Baddeley (1986) and can store up to 3 pieces. Patterns and schemas, for which Ericsson and Kintsch do not provide any definition, are assumed to refer to the chunks used in traditional computer models of expert perception and memory in chess, such as MAPP (Simon & Gilmarin, 1973) and CHREST (Gobet & Simon, in press). In the simulations discussed below, we will focus on master-level performance, requiring around 100,000 chunks, the upper bound of the range proposed by Simon and Gilmarin (1973) for the number of chunks acquired by chess masters. These chunks were obtained from discrimination nets grown by CHREST, using a database of masters’ games as input.

Finally, the 64-square hierarchical retrieval structure² consists of the “square” level, where individual pieces may be associated with their location, and of higher-levels, where chunks of pieces (e.g. “white King on g1; white Pawn on g2; white Pawn on h2”) can be stored.³ Levels in the hierarchy are defined by the size of chunks, counted in pieces.

Mechanisms of the model

During the two seconds allotted for the dictation of each piece, it is assumed that the following steps take place:

1. The dictated piece is stored in the articulatory loop.
2. The dictated piece is stored in the corresponding square in the square level of the retrieval structure, with a probability of p_{cue} .
3. Using the dictated piece and the pieces stored in the retrieval structure that are ± 2 squares away from it, the model attempts to match an LTM chunk. If successful, the chunk is added to the higher level of the retrieval structure corresponding to its size, with probability p_{cue} .
4. The model attempts to create a new chunk by combining the last chunk matched with the largest chunk in the retrieval structure. This mechanism corresponds to Ericsson and Kintsch’s idea of storage in LTM of new structures. The new chunk is created with probability $p_{elaboration}$ for last chunks > 3 , and with probability ($p_{elaboration} * bias$) with last chunks ≤ 3 , where *bias* is equal to .25, .50, and .75 for chunks of size 1, 2, and 3, respectively. This bias then favours elaboration of large chunks, which seems consistent with Ericsson and Kintsch’s description.

During the recall phase, the model places pieces using information from three sources: the articulatory loop, the square level and the higher levels of the retrieval structure, and the chunks that were created during the presentation of the position.

In order to understand the interaction between retrieval-structure encoding and LTM-elaboration encoding, I now explore the effects of the two parameters of the model:

1. p_{cue} , the probability that a piece or chunk is successfully stored in the retrieval structure
2. $p_{elaboration}$, the probability that a chunk is elaborated successfully

Consistent with Ericsson and Kintsch’s (1995) analyses, decreases in the values of these two parameters can be seen as reflecting the effects of proactive and retroactive interference.

Simulations

² For other ways to implement the concept of retrieval structure in chess, see Gobet (1993) and Saariluoma and Laine (in press).

³ A chunk may contain a single piece; for example: “white King on g1”.

The model was tested on the recall of 1,000 master game positions and 1,000 random positions. The number of pieces in each position varied from 15 to 32 with an average of 24.6 (in Saariluoma's experiments, this number varied from 18 to 28). The parameters p_{cue} and $p_{elaboration}$ were systematically varied from 0 to 1.

Results

Estimated from the graphs of Saariluoma's (1989) experiments 1 and 2, skilled subjects got 91% and 82% correct with game positions, and 62% and 57% correct with random positions. Figure 2 illustrates the results obtained in the simulations. Several features may be mentioned: (a) recall is easier for game positions than for random positions; (b) with $p_{cue} = 1$, indicating that all attempts to encode information in the retrieval structure are successful, the model obtains 100% correct both with game and random positions; (c) recall of random positions is higher than observed by Saariluoma (1989) with runs where $p_{cue} \geq .40$; (d) in general, the higher the probability of successful encoding into the retrieval structure, the smaller the role of elaboration becomes; (e) with random positions, p_{cue} plays an important role even with high values of $p_{elaboration}$; (f) many values which get a good fit with game positions do not match the data well with random positions; (g) with random positions, the model covers a large portion of the possible outcomes; even when one omits the results where any of the two probabilities is equal to zero, the model covers about half the space of possible outcomes.




Figure 2. Percentage correct as a function of the probability of making an elaboration $p_{elaboration}$ and the probability of storing information in the retrieval structure (p_{cue}). Left panel: game positions. Right panel: random positions. Asterisks indicate the results of the skilled subjects in Saariluoma's (1989) experiments 1 and 2.

Discussion of Simulations

In this paper, I have attempted to provide a computer implementation of the model presented by Ericsson and Kintsch for chess expertise. The simulations capture the important difference in recall performance between game and random positions. They also show that, with most values of p_{cue} , the computer model overestimates human recall with random positions, and that the best fit is obtained with a low probability of encoding information into the retrieval structure ($p_{cue} = 0.0, 0.2, \text{ or } 0.4$). They also illustrate that Ericsson and Kintsch's model, even with the additional assumptions made here, predicts a large number of possible outcomes, in particular with random positions. Without actually running the model and setting the values of the two parameters in question, it is impossible to estimate which value fits the data reasonably well. In particular, some values which fit the recall of game positions well do poorly with random positions. Using the verbal description of the model given by Ericsson and Kintsch, which does not explicitly consider the probabilities of encoding into the retrieval structure and of making LTM elaborations, one runs into the danger of inadvertently running "mental simulations" where these probabilities are not kept constant. This danger, obviously, is not specific to LT-WM, but to any theory formulated informally.

While several pairs of values account for the recall of game positions, in particular if one takes into account the variability of human data, the random position data are more selective. Saariluoma does not report the variance or the mean square error of his data, but one can, as a first rough approximation, assume that the band 10% above or below the mean value of the two experiments (58%) includes the "true" value. Using this criterion, the following pairs of $\{p_{elaboration}, p_{cue}\}$ lead to a result contained in this band: $\{0.2, 0.2\}$, $\{0.2,$

0.4}, {0.4, 0.2}, {0.4, 0.4}, {0.6, 0.2}, {0.8, 0.0}, {0.8, 0.2}, {1.0, 0.0}, and {1.0, 0.0}. All these values also lead to acceptable results with game positions. In none of these cases, the average of the two values is above .5, and that, in many cases, it is well below. This suggests that, as proposed elsewhere (Gobet, 1998), LT-WM has to use relatively low probabilities of encoding, either through the retrieval structures or through LTM elaborations, in order to account for human data.

It is of interest to compare the results of these simulations with the simulations of the digit-span task carried out by Richman, Staszewski, and Simon (1995). Their model, EPAM-IV, incorporates both retrieval structures and LTM schemas, as well as a limited-capacity articulatory loop. They found that a forgetting rate of 25% in the retrieval structure accounted well for the learning of their human subject. EPAM-IV did not incorporate forgetting in the semantic net, but assumed that its construction is time-consuming: 8 seconds to create a new chunk, and 2 seconds to add information to that chunk. These time parameters, which are also used in the CHREST models of chess memory, have similar effects as $p_{elaboration}$, the probability of elaborating items together or with LTM schemas and patterns. Thus, like the model discussed in this paper, EPAM-IV incorporates parameters which seriously either slow down or make less reliable the use of LT-WM. As noted by Richman, Staszewski, and Simon (1995), the simulations are quite sensitive to these parameters, a finding that was observed as well in the simulations discussed in this paper.

Conclusion

This paper has offered a computer implementation for LT-WM account of chess memory. In order to do so, several assumptions were made. Among the most important, one can mention: (a) an articulatory loop stores up to three dictated pieces; (b) patterns and schemas consist of chunks; (c) the hierarchical structure consists of a “square” level and of higher levels; the latter are defined as a function of the size of chunks; and (d) encoding into the retrieval structure is modulated by a parameter (p_{cue}), and so is creation of new LTM structures (elaboration).

Three important conclusions may be drawn. First, even though the computer model is more specific than Ericsson and Kintsch’s description, it predicts a wide range of results, which are quite sensitive to the value of the model’s parameters. Second, the parameters that best fit one set of data may not do so with another set; without a quantitative model, there is no safeguard against inadvertently changing these parameters from one task to the other. Third, the sets of parameters best fitting the data suggest that the probability of encoding information in the retrieval structure and the probability of making new LTM associations are relatively small (not more than 0.5, on average). Taken with the simulations of Richman, Staszewski and Simon (1995), these results suggest that the two mechanisms proposed by LT-WM have to be constrained by assumptions about processing limitations, using either time parameters or probability parameters, in order to fit the empirical data.

References

- Anderson, J.R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard Univ. Press.
- Baddeley, A. (1986). *Working memory*. Oxford: Clarendon Press.
- Chase, W.G., & Ericsson, K.A. (1982). Skill and working memory. In G.H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 16). New York: Academic Press.
- Chase, W.G., & Simon, H.A. (1973). The mind’s eye in chess. In W.G. Chase (Ed.), *Visual information processing*. New York: Academic Press.
- Ericsson, K.A., & Kintsch, W. (1995). Long-term working memory. *Psychological Review*, *102*, 211-245.
- Gobet, F. (1998). Expert memory: Comparison of four theories. *Cognition*, *66*, 115-152.
- Gobet, F., & Simon, H.A. (in press). Five seconds or sixty? Presentation time in expert memory. *Cognitive Science*.
- Kintsch, W. (1998). *Comprehension. A paradigm for cognition*. Cambridge, UK: CUP.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard Univ. Press.
- Richman, H. B., Staszewski, J., & Simon, H. A. (1995). Simulation of expert memory with EPAM IV. *Psychological Review*, *102*, 305-330.
- Saariluoma, P. (1989). Chess players’ recall of auditorily presented chess positions. *European Journal of Cognitive Psychology*, *1*, 309-320.
- Saariluoma, P. & Laine, T. (in press). Novice construction of chess memory. *Scandinavian Journal of Psychology*.
- Simon, H.A., & Gilmarin, K.J. (1973). A simulation of memory for chess positions. *Cognitive Psychology*, *5*, 29-46.