CHREST+ : Investigating How Humans Learn to Solve Problems Using Diagrams

Peter C.R. Lane, Peter C.-H. Cheng and Fernand Gobet

ESRC Centre for Research in Development, Instruction and Training,

School of Psychology, University of Nottingham,

University Park, NOTTINGHAM NG7 2RD, UK

`{pcl,pcc,frg}@psychology.nottingham.ac.uk`

Abstract

This paper describes the underlying principles of a computer model, CHREST+, which learns to solve problems using diagrammatic representations. Although earlier work has determined that experts store domain-specific information within schemata, no substantive model has been proposed for learning such representations. We describe the different strategies used by subjects in constructing a diagrammatic representation of an electric circuit known as an AVOW diagram, and explain how these strategies fit a theory for the learnt representations. Then we describe CHREST+, an extended version of an established model of human perceptual memory. The extension enables the model to relate information learnt about circuits with that about their associated AVOW diagrams, and use this information as a schema to improve its efficiency at problem solving.

Introduction

Earlier work has established the general form of a model of problem solving with diagrams and the kinds of internal representations required for effective performance. This project attempts to combine these ideas into a learning-based framework, so developing a computational model which can learn effective representations for problem solving with diagrams. Our framework is based on Chase and Simon's (1973) *chunking theory*, which hypothesises that an expert's memory contains a large number of perceptual chunks, which can be used for problem recognition and decomposition. CHREST (Gobet, 1998; Gobet & Simon, in press) is a computer model which uses a simulated eye and Short-Term Memory (STM) for learning a discrimination network of perceptual chunks. This model has been shown to provide an excellent fit to human data in the task of recalling chess positions.

CHREST shares a number of components with models of reasoning and inferencing with external representations (e.g. Tabachneck-Schijf, Leonardo & Simon, 1997). These include an external representation, a simulated eye and a STM. In general, the STM may contain a number of components, such as: a perceptual memory, for visuo-spatial information; a verbal memory, for propositional or sentential information; and also a memory for information relating to the current goals of the system. The Long-Term Memory (LTM) of CHREST is a discrimination network of perceptual chunks. In this paper we extend CHREST with an output device, a pen, for adding information to the external representation. More importantly, we include suitable internal representations

for associating problem states with domain-specific information about possible solutions. Detailed work in this area has identified *schemata* as the primary representation of skilled knowledge (e.g. Koedinger & Anderson, 1990). However, in spite of this theory being the basis of intelligent tutoring systems (e.g. Koedinger & Ander-
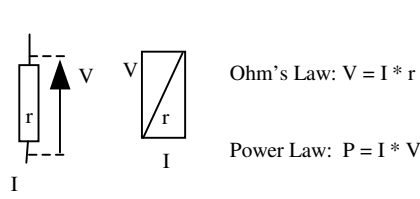

Figure 1 : An AVOW box for a resistor.

Ohm's Law: V = I * r

Power Law: P = I * V

(a) Series Resistors
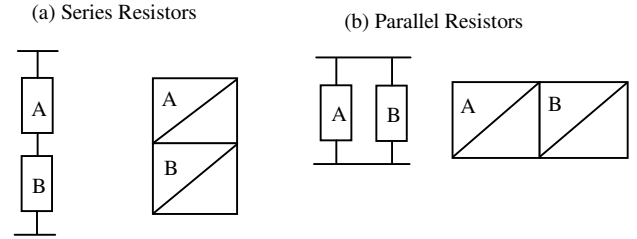
(b) Parallel Resistors

Figure 2 : Composing AVOW boxes for
(a) series and (b) parallel resistors.

son, 1993), no substantive computational model has been proposed for learning such representations. This question is addressed in the remainder of this paper.

This paper focuses on a specific diagrammatic representation, AVOW diagrams, for problem solving in the domain of electric circuits. Diagrammatic representations support efficient problem solving by humans because they encode domain-specific information in constraints such as the topology or geometry of the diagram (e.g. Larkin & Simon, 1987). AVOW diagrams are themselves an example of a wider class of diagrammatic representations, known as Law Encoding Diagrams (Cheng, 1996), which encode physical laws of scientific domains in the properties of the diagram. The construction of AVOW diagrams makes a suitable domain for our project because, firstly, subjects require a few hours to demonstrate developing expertise, and, secondly, the performance of subjects can be shown to conform with the chunking and schema theories of expert representations.

We begin by describing the diagrammatic representation used as a target domain for our model, and how data from subjects' protocols provides support for the chunking and schema theories of expert representations. Then we describe a concrete model for learning perceptual chunks, and extend it to support the acquisition of multiple representations. The extended model, known as CHREST+, learns information which supports a growing expertise in problem solving with diagrams.

Representing Circuits as AVOW Diagrams

We focus in this paper on a specific task requiring subjects to construct a diagrammatic representation for electric circuits. This task has the advantage of requiring relatively little training time before signs of expertise are observed in subjects. This rapid development is in part due to the fact that diagrammatic representations index information in a manner which supports useful and efficient computational processes (Larkin & Simon, 1987; Tabachneck-Schijf, Leonardo & Simon, 1997; Zhang & Norman, 1994). Cheng (1996) has introduced a range of representations for problem solving and learning in science which rely on geometric or topological properties of the diagrams to encode domain-specific laws. These kinds of representation enable perceptual information to de-

termine the conceptual similarity of separate instances. These points are manifested in our example domain, using a diagrammatic representation for electric circuits known as AVOW diagrams (Cheng, 1998, 1999).

An AVOW diagram is composed of a number of AVOW boxes, with each AVOW box in the diagram representing a specific resistor (or load) within an electric circuit. An example is shown in Figure 1. An individual resistor has the properties of voltage (V), current (I) and resistance (r), and these properties are represented diagrammatically in the AVOW box by scaling the indicated dimension, voltage being the height, current the width, and resistance the gradient of the box's diagonal. The relation of the gradient to the box's height and width encapsulates Ohm's law, and the area of the box also represents the power expended in the resistor.

Simple rules of composition are used to combine separate AVOW boxes into a complete diagram for an electric circuit. The composition process relies on breaking the circuit into collections of parallel and series resistors. Two series resistors are represented by aligning two AVOW boxes vertically, as shown in Figure 2(a), and two parallel resistors by aligning the boxes horizontally, as shown in Figure 2(b). Although simple, these rules encapsulate Kirchhoff's Laws which govern the flow of current and distribution of potential differences in electric circuits. For the completed AVOW diagram to be a well-formed representation of the circuit, it must be a rectangle completely filled with AVOW boxes with no overlap or gaps. These rules additionally capture an important abstraction often used in circuit analysis: a collection of resistors in a circuit can be regarded as equivalent to a single resistor, and formulae exist to compute this single resistor's resistance from that of its components. In the same way, a collection of AVOW boxes can be regarded as equivalent to a single AVOW box. The difference is that, just as with the single AVOW box, the resistance of the total AVOW diagram can be found by measuring the gradient of the total rectangle's diagonal, irrespective of the layout of the separate boxes which comprise it; in the equivalent algebraic case, separate formulae must be used for each arrangement of resistors. This, coupled with the geometrical nature of the composition rules, in large part explains computational benefits of working with this representation.

The Problem to Solve: Constructing an AVOW Diagram

The construction of an AVOW diagram for a given circuit requires the subject to obey two sets of constraints simultaneously: the first is to form an accurate representation of the circuit, and the second is to construct a well-formed AVOW diagram. For an ideal problem solver, defined as one for which no resource constraints apply, the problem becomes a technical one: for each resistor, a separate AVOW box must be drawn, and its dimensions (height, width, gradient) can be computed from the circuit using appropriate algebraic equations. This is possible because the AVOW diagram is an equivalent representation for the information in the circuit. However, the AVOW diagram, as a diagrammatic representation, provides some assistance with the necessary computations, and the way in which human learners take advantage of this is what makes their behaviour interesting and worth simulating.

Essentially, the human learner relies on the constraints imposed by the rules for composing AVOW diagrams. The geometric and intuitive nature of these constraints lead to the computational benefits when working with this diagrammatic representation. For instance, the compositional rules for AVOW boxes mean that the size of a box will be constrained by any neighbouring boxes, and so not require computing from the circuit diagram. This also means that each problem solver can adopt a different construction strategy, depending on which of these constraints is used at any time: either information is explicitly taken from the circuit diagram, or else the evolving AVOW diagram itself is used to constrain the construction process. In consequence, a rich variety of strategies is observed in human subjects, even with relatively simple problems. This can be seen in Figure 4, where we illustrate the separate steps taken by three subjects in constructing an AVOW diagram for the circuit illustrated in Figure 3(a); the complete AVOW diagram is illustrated in Figure 3(b).

The subject (S15) in Figure 4(a) begins by drawing the AVOW box for one of the resistors in the diagram; most often subjects start with the top-left one. Because the only knowledge about the resistor immediately available is that its resistance is 1 ohm, S15 draws a square AVOW box. Next, S15 applies the same reasoning to the adjacent resistor, but this time, because the two resistors are in parallel, the AVOW boxes are aligned horizontally. Finally, S15 can draw the third resistor, an AVOW box which is

(a) Circuit Diagram   (b) AVOW Diagram

$r_A = r_B = r_C = 1$ ohm

Figure 3 : Equivalent circuit and AVOW diagrams.

constrained to be aligned with the lower edge of the previous two boxes, and also a square, because its resistance is again 1 ohm. Once the AVOW diagram is complete, S15 can use a ruler (or a background grid/mental calculations) to find the quantities in the diagram; the total height of the AVOW box represents 12V, the given value of the source. Therefore, by measuring and rescaling, the rest of the quantities in the circuit can be simply obtained.

The second subject (S6), shown in Figure 4(b), exhibits a similar pattern but begins by only drawing diagonal lines for the resistance of each of the three resistors. These lines constrain the shape of the entire AVOW diagram, and the final step is to fill in the implicit bounding squares, completing the diagram. Radically different is the progress of the third subject (S11), shown in Figure 4(c). S11 begins by drawing a single vertical line to represent the voltage across the entire circuit, making this line a multiple of 12 grid units. The next piece of information to be filled in is the current flow through the left-hand of the top two resistors. This is followed by a line for the resistance of the lower resistor. At this point S11 now has a fully constrained diagram, and so proceeds methodically to complete it.
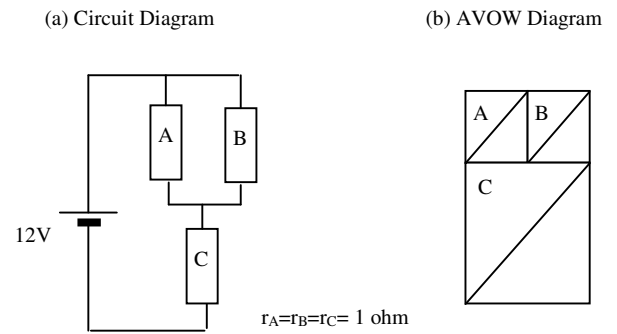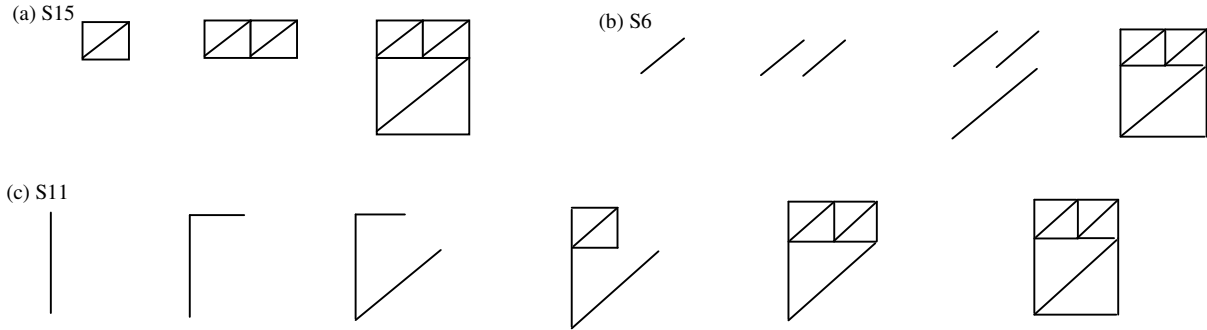
Figure 4 : Solutions by three subjects for the circuit in Figure 3(a). (This data has been taken from a larger study, Cheng, 1999). Each set of diagrams shows the separate steps in the construction of the solution.

The first thing to note is that the subjects use the sheet of paper as a store for known information, i.e. the representation is used as an external memory aid which reduces demands on STM. The difference between the subjects may be explained on two distinct dimensions: first, subjects differ in their level of experience with the domain; second, subjects differ individually in the sequences of actions used to construct the AVOW diagram. The different strategies used by the subjects S15 and S11, illustrated in Figures 4(a) and 4(c) respectively, may be explained with the theory of perceptual chunking (Chase & Simon, 1973; Egan & Schwartz, 1979; Koedinger & Anderson, 1990). For instance, S15 draws components of the circuit at the single resistor level, whereas S11 begins by drawing a line representing the voltage for the entire circuit and proceeds by filling out key lines to constrain the diagram. Individual differences can be seen in how S6 and S11 fill out critical information to constrain the full AVOW diagram before completing the details, whereas S15 carefully completes each AVOW box before moving on to the next. Taken together, this suggests that subjects use perceptual cues from the circuit diagram to form an internal representation, or mental impression, of how the completed AVOW diagram should look.

The basic elements for modelling such behaviour are an eye, a STM and an appropriate long-term perceptual memory. We restrict our attention in this paper to the acquisition of appropriate perceptual information. We begin by discussing the chunking theory for perceptual memory, from which CHREST was developed. Later we show how CHREST's learning operations and the use of an appropriate STM and directable eye model the acquisition of chunks of perceptual information.

The Chunking Theory of Memory

The chunking theory of memory is based on EPAM (Elementary Perceiver and Memoriser), a well-known computer model of a wide and growing range of memory tasks. The basic ideas behind EPAM include mechanisms for encoding chunks of information into long-term memory (LTM) by constructing a discrimination network. The EPAM model has been used to simulate the learning of verbal material (Feigenbaum & Simon, 1962, 1984) and expert digit-span memory (Richman, Staszewski & Simon, 1995). EPAM has been expanded to use visuo-spatial information, as in MAPP (Simon & Gilmartin, 1973). CHREST (Gobet, 1998) is a further extension of EPAM which includes the ability to learn templates and semantic links between nodes. Next we describe the learning mechanisms used to construct a discrimination network and explain how CHREST can be extended to be a model of problem solving. In a later section we describe an extension of CHREST where nodes can be linked to represent equivalences between multiple representations.

CHREST organises memory into a collection of chunks, where each chunk is a meaningful group of basic elements. For example, in chess, the basic elements denote the pieces and their locations; the chunks are collections of pieces, such as a king-side pawn formation. These chunks are developed as the discrimination network grows through the processes of *discrimination* and *familiarisation*. Essentially, each node of the network holds a chunk of information about an object in the world. The nodes are interconnected by links into a network, with each link representing the result of applying a test to the object. When trying to recognise an object, the tests are applied beginning from the root node, and the links are followed until no further test can be applied. At the node reached, if the stored chunk matches that of the object then familiarisation occurs, in which the chunk's resolution is increased by adding more details of the features in that object. If the current object and the chunk at the node reached differ in some feature, then discrimination occurs, which adds a new node and a new link based on the mis-matched feature. Therefore, with discrimination, new nodes are added to the discrimination network; with familiarisation, the resolution of chunks at those nodes is increased.

The experiments in the recall of chess positions reported in Gobet (1998; Gobet & Simon, in press) show that CHREST captures the main features of perceptual memory gathered in experiments with human subjects; the difference between expert and novice behaviour is explained by the size of the discrimination network, i.e. the number of stored chunks of information. However, CHREST as it stands is not a model of problem solving behaviour. For instance, CHREST does not play chess as it lacks a mechanism for handling the construction of game trees and the interaction of various chunks. Although CHUMP (Gobet & Jansen, 1994), a program based on CHREST, does play chess, it does so purely by pattern recognition and without search. The important ability required in complex problem solving, which CHREST lacks, is the ability to form a plan. Accordingly, we adapt CHREST to handle multiple external representations, and apply it to acquiring perceptual chunks of electric circuits and their associated AVOW diagrams. The advantage of this domain is that a visual image of the target diagram can be used as a plan for problem solving, which is of the same type as the perceptual memory being acquired, whereas in chess, plans require a separate type of knowledge. In the remainder of this paper we describe some of the details of our current implementation of this model, and also show how our approach forms the basis of a larger model of problem-solving behaviour.

Learning Internal Representations for Problem Solving

For the model to learn to solve problems using diagrams, it must first extract information from the external representation. We avoid an inappropriate amount of low-level simulation by assuming a set of basic visual primitives. These include identification of separate rectangles and shapes in the diagrams, as well as their relative positioning, alignment and interconnections. These primitives are a subset of those described in, for example, Lindsay (1988), and enable the model to parse the diagram into separate resistors or AVOW boxes; no experimental subject has shown difficulty in such parsing. We describe next how the model combines this information into larger

chunks for whole diagrams, the learning mechanism for associating problem states and their solutions, and how the learnt information generalises to assist in improving problem-solving ability.

Acquiring Large Perceptual Chunks

Small chunks are those of the order of the size of the visual field. These can be learnt and recognised by passing the information retrieved by the eye directly to the perceptual memory (LTM), where the standard familiarisation/discrimination process will apply. In order to acquire chunks for visual images which extend beyond the visual field, an interaction is needed between the eye, STM and LTM. The procedure here is the same as that used in CHREST (Gobet, 1998). The visual STM contains a queue of pointers to the last chunks observed. One of these chunks has a special status, and is known as the 'hypothesis', the largest chunk currently considered. Information retrieved from the visual field is passed to the LTM and familiarisation/discrimination will occur if appropriate. A pointer to the chunk indexed by the current visual object is placed in the STM queue. The hypothesis chunk is then combined with the most recent chunk stored in STM, and this new chunk will be used for further learning in LTM.

   This combination process depends upon two assumptions: firstly, that the two chunks combined will be overlaid without duplication, and secondly that the image of the combined chunks will be the same if the chunks were combined in a different order. This latter assumption deals with the situation of visual chunks being retrieved by a different sequence of eye movements. Each of these assumptions is handled in our current implementation by retaining the identity of the separate elements in the visual field: a sorting process on these elements then removes the duplications and ensures all chunks with the same elements are stored in an identical fashion.
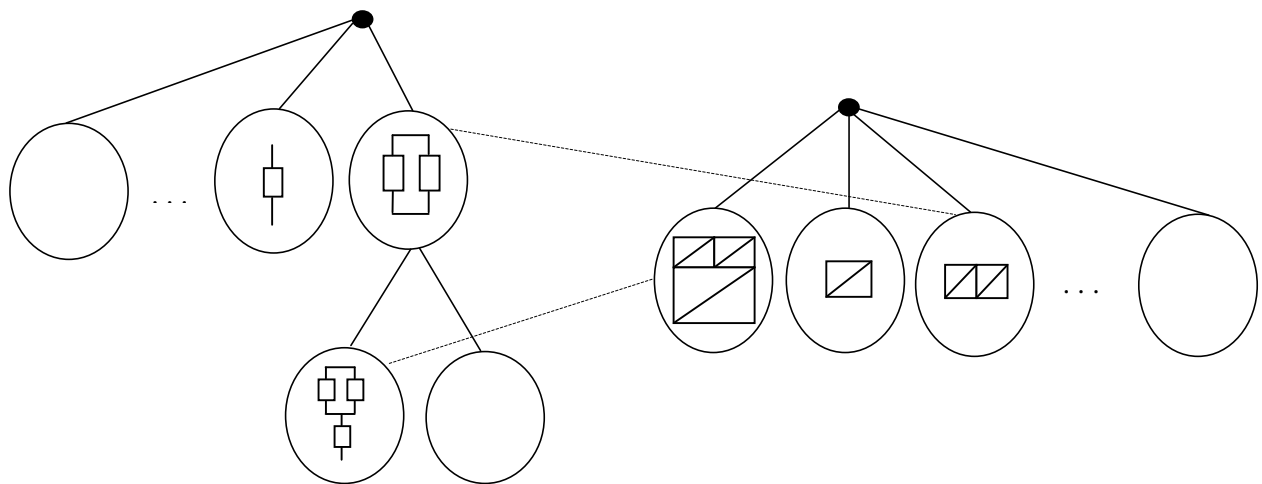
Figure 5 : A discrimination network for multiple representations, showing two equivalence links and the inheritance structure. The chunks at some of the nodes are illustrated, dark lines represent test links and dashed links equivalence links.

Combining Multiple Representations

In order to work with more than one representation, a method is needed for indexing information in different parts of a discrimination network other than with the links representing tests. Gobet (1996) describes such a method, showing how production rules and semantic networks can be learnt by combining multiple networks. All that is required is the use of additional links, which can be used when traversing the network.

This method can be used to combine chunks for multiple representations; we assume that every node in the network will refer to a separate representation, which has a similar effect to a high-level test in the discrimination network for separating out the representations. Figure 5 depicts part of the network learned from the circuits in Figures 1-3 (the full network contains around 20 nodes): the left-hand part of the network represents information about electric circuits, and the right-hand information about AVOW diagrams. The relationship between these two representations can be shown by an *equivalence link* which connects a node from the electric circuit network with its equivalent AVOW representation in the AVOW diagram network. The dashed links in Figure 5 show examples of this.

Before constructing an equivalence link, the two separate representations must be fully learned. This condition is easily identified during the recognition of a visual chunk; if no discrimination or familiarisation occurs at the node reached during recognition, then that node has fully learned the current chunk. Consider now that the system is presented with a circuit, recognises it, and so places in STM a pointer to a fully learned chunk representing that circuit. On looking at the AVOW diagram, no training during its recognition will occur if a fully learned representation for the AVOW diagram exists. Because of this, the pointer to the node representing the circuit diagram will not be displaced from STM. Once the AVOW diagram has also been fully recognised, an equivalence

link can be formed between the two nodes. This is the critical mechanism for learning multiple representations and associations between them in CHREST+, which is an important extension of CHREST.

Solving Problems like Humans

The LTM acquired for circuit diagrams and their associated AVOW diagrams has two features which generate human-like problem solving. First, the associations between the circuit and AVOW information learnt by the model form *schemata*, i.e. domain-specific information which supports problem solving (Sweller, 1988). The circuit information provides a template, or set of pre-conditions, for determining the set of circuits to which the schema applies, and the AVOW information provides a template for constructing specific AVOW diagrams for a range of possible quantities specified in the circuit.

The second feature of our model is the efficiency with which these schemata support problem solving. This support comes in two parts: flexibility in creating specific solutions, and generalisation to complex circuits. Firstly, let us assume a solver matches a given circuit to a chunk in its memory and retrieves information about the corresponding AVOW diagram. The AVOW diagram must then be drawn as a series of separate lines, each line of the length determined by the quantities in the circuit diagram. As there is no constraint in the AVOW diagram's representation on the precise sequence in which to do this, the solver is free to fit its drawing sequence with the information provided in the circuit. For example, information about current will lead the solver to begin by drawing the width of the appropriate box to scale; information about voltage will lead the solver to scale the height of the appropriate box. Our model therefore exhibits the step-skipping behaviour described by Koedinger and Anderson (1990, 1993) as a characteristic of human solvers.

A second support for problem solving is that each schema may be matched to sub-circuits as well as entire circuits, and it is a feature of the representation that perceptually similar circuits will have perceptually similar AVOW diagrams. Therefore, our model can use the perceptual hierarchy in the discrimination network to retrieve subgoals; if the circuit were that in Figure 3(a) but the lower equivalence link in Figure 5 had not been learnt, the model could instead use the link from its parent node to retrieve a diagram for a sub-component of the circuit. This feature also models individual differences observed in the way subjects tackle a problem. Essentially, because the discrimination network is built up in an incremental manner, individual differences between subjects arise naturally due to differences in the training sequence, caused, for instance, by subjects following different attentional strategies. These differences lead to different discrimination networks, which means different sub-circuits will be retrieved from LTM, and therefore a different solution strategy will be followed.

Conclusion

This paper has described the underlying principles of a computational model, CHREST+, for simulating the learning process by which expertise at problem solving with diagrams develops in human subjects. CHREST+ is

based on an established model of perceptual memory which learns perceptual chunks. The key feature of CHREST+ is an ability to associate information about two different representations with equivalence links, and we use this ability to learn representations for problems and their solution. We have shown how the model applies to subjects learning a diagrammatic representation for electric circuits known as AVOW diagrams. CHREST+ has a number of desirable features for this level of cognitive modelling: firstly, the use of a perceptual memory and the learning of perceptual chunks is supported by observations of such behaviour in human subjects; secondly, the use of the AVOW diagram as a template for the solution process means that behaviour such as step-skipping is observed in the model. Future work will focus on a more complete comparison of CHREST+'s behaviour with human data and also its application to further domains.

References

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.

Cheng, P. C.-H. (1996). Scientific discovery with law-encoding diagrams. *Creativity Research Journal*, 9, 145-162.

Cheng, P. C.-H. (1998). A framework for scientific reasoning with law encoding diagrams: Analysing protocols to assess its utility. In M. A. Gernsbacher & S. J. Derry (Eds.) *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society* (pp. 232-235). Mahwah, NJ: Erlbaum.

Cheng, P. C.-H. (1999). Electrifying representations for learning: An evaluation of AVOW diagrams for electricity. *Technical Report No.64, ESRC Credit, University of Nottingham.*

Egan, D. E., & Schwartz, B. J. (1979). Chunking in recall of symbolic drawings. *Memory and Cognition*, 7, 149-158.

Feigenbaum, E. A., & Simon, H. A. (1962). A theory of the serial position effect. *British Journal of Psychology*, 53, 307-320.

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 Second International Conference of the Learning Sciences* (pp. 398-403). Evanston, III: Northwestern University.

Gobet, F. (1998). Memory for the meaningless: How chunks help. In M. A. Gernsbacher & S. J. Derry (Eds.) *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society* (pp. 398-403). Mahwah, NJ: Erlbaum.

Gobet, F. & Jansen, P. (1994). Towards a chess program based on a model of human memory. In H. J. van den Herik, I. S. Herschberg, & J. W. Uiterwijk (Eds.) *Advances in computer chess 7*. Maastricht: University of Limburg Press.

Gobet, F. & Simon, H. A. (in press). Five seconds or sixty? Presentation time in expert memory. *Cognitive Science*.

Koedinger, K. R., & Anderson, J. R. (1990). Abstract planning and perceptual chunks: Elements of expertise in geometry. *Cognitive Science*, 14, 511-550.

Koedinger, K. R., & Anderson, J. R. (1993). Reifying implicit planning in geometry: Guidelines for model-based intelligent tutoring system design. In S. P. Lajoie & S. J. Derry (Eds.), *Computers as Cognitive Tools*, Lawrence Erlbaum, New Jersey.

Larkin, J. H. & Simon, H. A. (1987). Why a diagram is (sometimes) worth ten thousand words. *Cognitive Science*, 11, 65-99.

Lindsay, R. K. (1988). Images and inference. *Cognition*, 29, 229-250.

Richman, H. B., Staszewski, J., & Simon, H. A. (1995). Simulation of expert memory with EPAM IV. *Psychological Review*, 102, 305-330.

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

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science*, 12, 257-285.

Tabachneck-Schijf, H. J. M., Leonardo, A. M., & Simon, H. A. (1997). CaMeRa: A computational model of multiple representations. *Cognitive Science*, 21, 305-350.

Zhang, J., & Norman, D. A., (1994), Representations in distributed cognitive tasks. *Cognitive Science*, 18, 87-122.