Problem Solving with Diagrams :

Modelling the Learning of Perceptual Information

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
United Kingdom
email: {pcl,pcc,frg}@psychology.nottingham.ac.uk

ESRC Centre for Research in Development, Instruction and Training
June 1999
Technical Report No. 59

Abstract

This paper presents work on a computer model of how humans learn to solve problems with diagrams. Although earlier work has discussed the kinds of representations which experts use when solving problems, not much work has been done on how these representations may be learned. We focus on a particular problem solving domain and describe a specific diagrammatic representation for it. We define a novel computer model which is able to use the perceptual information available in typical problems in this domain to learn appropriate representations to assist the problem solving process. Specifically, we consider the different strategies used by subjects in constructing a diagrammatic representation of an electric circuit known as an AVOW diagram. Experiments offer evidence that subjects learn perceptual chunks about the domain, and also show how subjects exploit the external representation. Of primary importance for a more comprehensive model is how the problem solver creates an appropriate internal representation of the domain. We use a model for Long-Term Memory based upon an extended version of EPAM, which can handle multiple representations. The model as currently implemented retrieves visuo-spatial information using a directable eye and obeys the constraints of Short-Term Memory to learn perceptual chunks about circuits and their associated AVOW diagrams.

Email: {pcl,pcc,frg}@psychology.nottingham.ac.uk
Phone 0115 9515151, Fax 0115 9515324

1 INTRODUCTION

A number of studies have shown that humans learn to solve problems more efficiently when the problem is in an appropriate diagrammatic representation. In order to better understand the kinds of information learnt from diagrams and also how to create better representations for different tasks, we aim to construct a computer model which simulates this ability. For this model to be effective, certain features should be present in the problem solving domain. For instance, the domain should require the use of visuo-spatial information and show a rich variety of behaviour from different subjects. In order to assess the development of a subject's ability, an effective grading system should exist and timing and response data should be measurable. For a rapid turnover in experiments, an observable development in the subject's ability should require a training time of the order of hours and not years. In this paper we describe part of an on-going project which aims to develop a simulation of the use of diagrammatic representations in problem solving using AVOW diagrams, a novel representation for electric circuits. AVOW diagrams are highly suited for the study of problem solving with diagrams as they are designed to be a good diagrammatic representation for solving problems relating to circuits. We begin with a brief overview of the main components in the computational model to be developed for the simulations.

The general form of a model of problem solving with diagrams has been well established by earlier work on reasoning and inferencing with external representations (for example, Tabachneck-Schijf, Leonardo & Simon, 1997). Firstly, there is the external representation itself, which in our case is a (computer representation of a) sheet of paper containing line drawings of the circuit and AVOW diagrams. Secondly, if the system is to interact with its external representation, it also requires some input and output devices. In our case these are a directable eye, for retrieving information from various parts of the paper, and a pen, for adding information to the paper. Thirdly, to match the cognitive processes of humans, there should be a Short-Term Memory (STM). 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. Finally, the system must possess a Long-Term Memory (LTM), which also may have different components, the counterparts to those in STM.

However, this earlier work has not addressed the question of learning information about these external representations in order to support a growing expertise in problem solving. Accordingly, we focus in this paper on how the long-term perceptual memory is acquired whilst learning to construct AVOW representations of electric circuits. We show that the observed behaviour of human subjects in constructing AVOW diagrams is consistent with the perceptual chunking theory (Chase & Simon, 1973; Egan & Schwartz, 1979; Koedinger

& Anderson, 1990). One of the better models for this memory is that of the EPAM-chunking

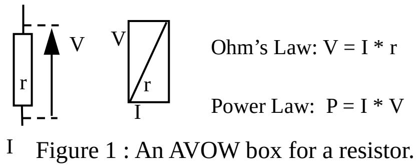

Figure 1 : An AVOW box for a resistor.

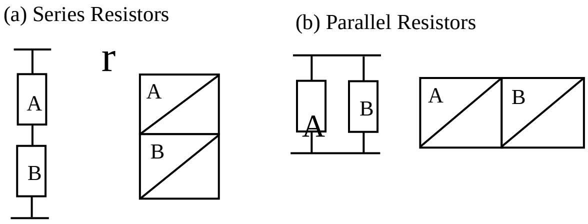

Figure 2 : Composing AVOW boxes for

(a) series and (b) parallel resistors.

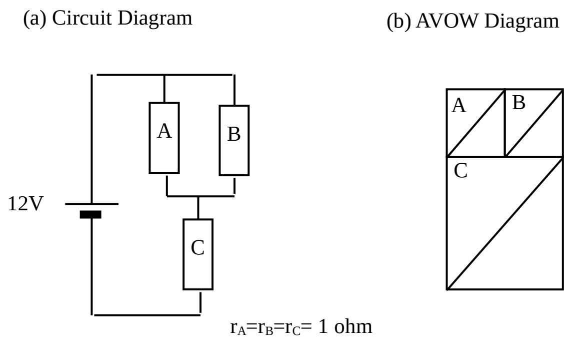

Figure 3 : Equivalent circuit and AVOW diagrams.

theory. Our current implementation of these ideas uses a graphical computer environment with a directable eye for retrieving diagrammatic information from circuit and AVOW diagrams. A visual STM is used in conjunction with an extended form of EPAM to acquire perceptual information about multiple external representations, i.e. the circuit and AVOW diagrams. Finally, we show how this memory can be used as a component in a more comprehensive computer model of how humans learn to solve problems with diagrams.

2 REPRESENTING CIRCUITS AS AVOW DIAGRAMS

In this paper we are interested in how subjects learn to use a novel diagrammatic representation for electric circuits, AVOW (Amps, Volts, Ohms, Watts) diagrams (Cheng, 1998, submitted). The advantages of diagrams as representations stem largely from their indexing of information in a manner which supports useful and efficient computational processes (Larkin & Simon, 1987; Tabachneck-Schijf, Leonardo & Simon, 1997). AVOW diagrams are designed so that these computational properties of diagrams aid problem solving with electric circuits, and they are one example of a range of such representations for problem solving and learning in science known as Law Encoding Diagrams (Cheng, 1996).

An AVOW diagram is composed of AVOW boxes, each AVOW box being a diagrammatic representation of a resistor (or load) within an electric circuit, as shown in Figure 1. A resistor has the properties of voltage (V), current (I) and resistance (r). These properties are represented 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. It can be seen that the gradient encapsulates Ohm's law, and also that the area of the box represents the power (P=I*V) expended in the resistor.

The AVOW boxes are combined into an AVOW diagram for an entire circuit using simple rules of composition. In order to represent two series resistors, two AVOW boxes are aligned vertically, as shown in Figure 2(a). Similarly, two parallel resistors are represented by aligning the boxes horizontally, as shown in Figure 2(b). The alignment rules encapsulate Kirchoff'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. This requirement captures an important abstraction 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, the composite AVOW diagram is also an AVOW box, containing all the information for this equivalent single resistor. 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. The geometrical nature of this constraint on the final AVOW diagram and the rules for composing separate AVOW boxes mean that it is very natural for humans to work within this representation, a fact which has important pedagogical implications. For the purposes of constructing a computer model, the geometrical nature of this knowledge means that the long-term memory can be assumed to be perceptual.

3 CONSTRUCTING AVOW DIAGRAMS

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. The geometric and intuitive nature of these constraints lead to certain 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.

We illustrate this by describing the construction of an AVOW diagram for a simple circuit by an ideal problem solver and some human subjects. We then discuss how a computer model, based on earlier work with models of perceptual chunking, can be developed to simulate these different problem-solving strategies.

3.1 The Ideal Problem Solver

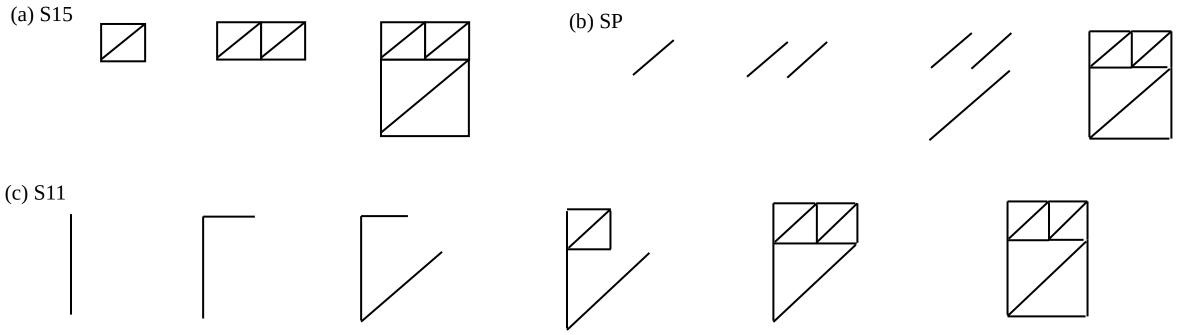

Figure 4 : Solutions by three subjects for the circuit in Figure 3(a).

Each set of diagrams shows the separate steps in the construction of the solution.

An ideal problem solver is one for which no resource constraints apply, so that the solver can extract and manipulate whatever information is required for solving the problem of constructing an AVOW diagram from a given circuit diagram. A typical circuit is that shown in Figure 3(a) which contains three resistors. In order to construct the AVOW diagram, three AVOW boxes must be drawn, each requiring its dimensions (height, width, gradient) and position in the diagram to be determined. The ideal problem solver first calculates, for each resistor, the amount of voltage drop across it and the amount of current flow through it. This can be done in a number of ways, but all use equivalent applications of Ohm's law and formulae for parallel and series resistors. The position of each box must then be determined to satisfy the compositional rules: the AVOW boxes for the upper two resistors must be placed side-by-side on top of the lower resistor's AVOW box. In our implementation of an ideal problem solver, an AVOW diagram such as that in Figure 3(b) is constructed.

3.2 Human Subjects

A number of subjects have been trained in our laboratory (Cheng, submitted) to construct AVOW diagrams from circuit diagrams. From such studies we can observe how different subjects use different strategies whilst solving the same problem. In Figure 4 we illustrate this variety by showing the separate steps taken by three subjects in constructing an AVOW diagram for the circuit illustrated in Figure 3(a).

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

The second subject (SP), 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 lines, 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.

3.3 Modelling the Human Behaviour

A number of differences can be observed in the problem-solving strategies of the various solvers, ideal and human. Firstly, there is a large difference between the ideal and the human solvers in that, for the latter, the sheet of paper is used as a store for known information, i.e. the paper is used as an external representation in order to reduce demands on STM. Secondly, it is evident that human subjects display a great variety in their approach to a given problem. There are at least two distinct dimensions explaining this difference: 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; 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 SP 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 can form an internal representation for the whole circuit diagram, and this is in the form of a 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 EPAM-chunking theory for perceptual memory, and later show how its learning operations and the use of an appropriate STM and directable eye model the acquisition of chunks of perceptual information.

4 THE EPAM-CHUNKING THEORY OF MEMORY

EPAM (Elementary Perceiver and Memoriser) is 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. We describe next the learning mechanisms used to construct a discrimination network, which are common to all EPAM variants, and we show how CHREST needs extending to be a model of problem solving. In a later section we describe such an extension of CHREST, where nodes can be linked to represent equivalences between multiple representations.

EPAM The EPAM model organises memory into a collection of chunks, where each chunk is a meaningful group of basic elements. For example, in chess, the basic elements are the pieces and their locations; the chunks are collections of pieces, such as a king-side pawn formation. These chunks are developed as the EPAM 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.

CHREST The experiments in the recall of chess positions reported in Gobet (1998) show that CHREST captures all 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.

5 CREATING INTERNAL REPRESENTATIONS

Several components are required in a simulation of learning to solve problems using diagrams. Of primary importance is the internal representation of the domain created by the problem solver; the difficulty is, because this perceptual memory is acquired implicitly whilst the subject is observing the diagrams, the exact form of the memory can only be known indirectly. We argue here that this representation is created through an interaction of the information acquired by the eye, the structure of the internal STM and LTM mechanisms, and the demands of the problem-solving process.

In the next four sections we discuss our current model for the creation of this internal representation. We begin with a discussion of the level of visual information required from the eye, and how this is built up into perceptual chunks for entire diagrams using EPAM and a finite STM. Our domain for problem solving requires associations to be formed between two distinct external representations, circuit and AVOW diagrams. We describe an extended EPAM model for handling multiple representations, one of the advantages of which is that an inheritance structure is formed for the perceptual chunks. This structure can be used to generalise information about simpler diagrams to more complex examples. Also, because the memory is acquired dynamically, the precise form of the network will vary depending on the order of presentation of the training examples. We briefly discuss the prospects for the next stage of this project, an integrated model of problem solving based on this memory for multiple representations.

5.1 Retrieval of Diagrammatic Information

In order to work with diagrams, the information acquired from the eye must contain an appropriate level of diagrammatic information; too abstract a representation will lose the benefits of working with diagrams, and too fine a level of detail will involve the model in an inappropriate amount of low-level simulation. Therefore, although our model contains a directable eye, it is not our intention to simulate the visual process beginning from a low-level image at the retina. Instead, an appropriate level of visual input for this project is the diagrammatic information contained within the field of view.

The basic primitives we use allow for identification of the separate rectangles and shapes in the diagrams, and the relative positioning, alignment and connections between them. These primitives are a subset of those described in, for example, Lindsay (1988). Therefore, when looking at part of a circuit diagram, the information retrieved will describe the resistors, their interconnections and their relative positions within the visual field; similarly for the AVOW diagram, the separate AVOW boxes and adjacency relationships will be described for those boxes within the current field of view. This approach is further justified by the observation that human subjects have no difficulty in parsing a diagram, so that a circuit is readily seen as a collection of connected resistors.

5.2 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 received 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 arises if the visual chunks are 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.

5.3 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

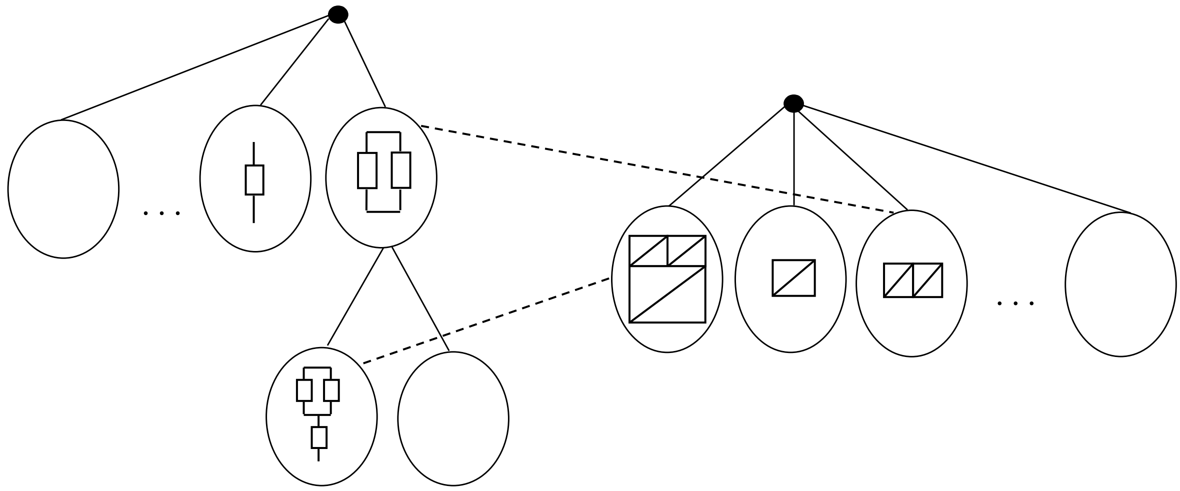

Figure 5 : Multiple discrimination networks, showing two equivalence links and the inheritance structure. The chunks at some of the nodes are illustrated.

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, for simplicity, that a high-level test in the discrimination network separates out the two representations; although this is not strictly necessary, as all that is required is for the chunk at each node to be either that of an electric circuit or that of an AVOW diagram, and this condition should arise automatically as the visual procedure extracts information from the diagrams. Figure 5 depicts part of the network learned from the circuits in Figures 1-3 (the full network contains around 70 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 dotted links in Figure 5 show examples of this.

In order to construct 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. We imagine 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. Turning now to 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.

5.4 Constructing a Solution

The perceptual memory described so far has some useful properties for the larger system we are developing for solving problems using diagrams. Firstly, the equivalence links within the memory can be used to retrieve plans to guide the construction of a solution. For example, once a system recognises a circuit, such as that in Figure 3(a), its equivalent AVOW representation can be retrieved from memory. This representation can then be instantiated by drawing lines of the appropriate length and layout on the sheet of paper.

The second useful property lies in the type of generalisations possible in the LTM discrimination network. Because this memory uses tests to discriminate between different objects, the nodes in the network are organised into an inheritance structure for diagrams based on their perceptual similarity. This structure can be used to generalise information about simpler diagrams to more complex examples. For example, when recognising a diagram in one representation (a circuit), an equivalent diagram in another representation (its AVOW diagram) should be indexed. If, however, a node is reached where no equivalence links are present, it is possible to consider equivalence links from the node's ancestors. For example, the parent node for the circuit in Figure 3(a) might contain a chunk for a sub-component of the whole circuit; in Figure 5, a chunk for two parallel resistors is its parent. Therefore, where no equivalent AVOW representation has been learnt for a given circuit, representations of its sub-components can be used to create a partial AVOW diagram. The diagrammatic constraints imposed by the AVOW representation will then help guide the solver towards a correct representation for the total circuit.

Finally, the dynamic nature of learning in EPAM-type models means that individual differences in perceptual memory occur quite naturally due to differences in the training sequence. The separate system required to instantiate the visual plans will also provide scope for individual differences, depending on the (acquired) heuristics and the data observed in the circuit. These considerations show that the model we propose contains the required richness and flexibility for modelling the different problem solving strategies observed in Figure 4.

6 CONCLUSION

This paper has described a project to simulate human subjects learning to solve problems with diagrams; specifically, subjects must construct an AVOW diagram representation for a given electric circuit. This project uses an extended version of the CHREST model of perceptual memory in order to learn equivalences between multiple diagrammatic representations. This model extends on earlier work by addressing the question of learning information about external representations in order to support a growing expertise at problem solving within our domain.

Our current implementation uses a graphical computer environment with a directable eye for retrieving diagrammatic information from circuit and AVOW diagrams. The model builds up chunks of perceptual information in LTM, which appear of the right kind to explain the behaviour of different problem solvers. Also, because the incremental learning used by EPAM can construct different networks depending on the order of the training data, this model explains fairly directly how individual differences in problem-solving behaviour can arise. We have shown how this model for long-term perceptual memory can be used to provide a visual plan to guide the construction of an AVOW diagram. This project will continue by developing a drawing module by which the system can interact with its external representation, at which point we should have a computer model which can be trained to solve problems using diagrams.

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. (submitted). Electrifying representations for learning: An evaluation of AVOW diagrams for electricity.

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. & 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. (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.

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

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.

Qin, Y. & Simon, H. A. (1995). Imagery and mental models in problem solving. In J. Glasgow, N. Hari Narayanan and B. Chandrasekaran (Eds.) *Diagrammatic reasoning*, AAAI Press, Menlo Park, California.

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.

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.