In N. R. Shadbolt, H. O'Hara, & G. Schreiber (Eds.), *Advances in Knowledge Acquisition: 9th European Knowledge Acquisition Workshop, EKAW'96.* (pp. 179-194). Berlin: Springer-Verlag.

Diagrammatic Knowledge Acquisition: Elicitation, Analysis and Issues

Peter C-H. Cheng

ESRC Centre for Research in Development, Instruction and Training,
Department of Psychology, University of Nottingham, Nottingham, NG7 2RD.
email: peter.cheng@nottingham.ac.uk

Abstract. This paper considers the acquisition of knowledge that experts would naturally express in the form of diagrams — diagrammatic knowledge acquisition, DKA. The implications for DKA of previous research on diagrammatic representations and reasoning are considered. Examples of knowledge elicitation and knowledge analysis, with two different diagrammatic representations, are given. They demonstrate the feasibility of DKA. Issues raised by the analysis of the examples are discussed and consideration is given to the development of DKA tools and methodologies. DKA is distinguished from knowledge visualization, which attempts to design effective visual presentations of given information that is already expressed as propositions.

1 Introduction

What are the potential roles for diagrammatic knowledge representations in knowledge acquisition (KA)? What are the problems to be addressed when attempting to use diagrams in eliciting and analysing expert knowledge for use in knowledge based systems, KBSs? The desire of experts to draw diagrams whilst solving problems or explaining some aspect of their specialist domain will be a familiar phenomenon to knowledge engineers. Similarly, figures, charts, flow diagrams and a host of other diagrammatic forms are found in the documents of most domains, such as text books, instruction manuals, working sketches. However, use of diagrammatic representations for KA has been rather neglected in this field. By overlooking diagrams we may not only be missing important knowledge and powerful problem solving methods, in some domains, but also missing efficient methods for eliciting, analysing and implementing knowledge.

Diagrammatic representations have an established place in knowledge engineering, but this has mainly been in the form of data visualization aids, rather than in processes to acquire knowledge (see Jones, 1988 for a review). In cognitive science much is known about how humans reason with particular diagrams and what makes effective diagrammatic representations. Several AI systems have demonstrated the feasibility and benefits of using diagrammatic representations, which could be potentially incorporated into KBSs.

Diagrammatic representations are being considered rather than *visual representations*. Diagrammatic representations are a subset of all visually perceived representations, and are graphical notational systems that involve artificial, symbolic or simplified visual depiction, which are not organised in the form of a linear verbalisable or propositional structures (Kulpa, 1994). In this view, most KA techniques are not diagrammatic, although some exploit visual features, including some implementations of repertory grids and concept laddering. The graphical aspects of these techniques aids the visualization of the information rather than deeply encoding the knowledge. *Propositional representations* will to refer to conventional knowledge representational formalisms, such a logic, semantic networks and frames/schema.

Shaw and Woodward (1990) present a framework describing knowledge acquisition as modelling expert knowledge. The adequacy of their model will not be discussed here, but it is merely adopted as a framework that provides a convenient context and consistent terminology, for the purposes of the present analysis. Central to the framework are the processes involved in the production of (i) intermediate knowledge bases and (ii) computer knowledge bases from (a) mental models, (b) conceptual models and (c) semi-formal conceptual models[1]. A computer knowledge base is essentially a KBS and intermediate knowledge base is an organized and refined, but unimplemented, body of knowledge. Mental models are internal to experts and conceptual models are informal externalizations of the mental models, which are communicated to others. Semi-formal conceptual models are final working models or operational models. Elicitation and analysis procedures are involved in the production of intermediate knowledge bases. Elicitation follows introspection of mental models and conceptual models. Analysis procedures draw upon communication processes with conceptual models and the semi-formal conceptual models. Analysis procedures are also required for the production of computer knowledge bases, in conjunction with implementation procedures that use formalisations of semi-formal conceptual models. The development of intermediate knowledge bases and computer knowledge bases inform each other by means of the analysis procedures. Shaw and Woodward view the processes and procedures as bi-directional; the development of the knowledge bases and the various levels of models mutually aid each others' development. The image of acquisition presented has multiple layers of knowledge, which are closely interconnected by cognitive processes and acquisition procedures.

Now, for the considerations of diagrammatic KA (DKA) it is a matter of identifying diagrammatic versions for each of the knowledge and procedural components of Shaw and Woodward's framework. What are the diagrammatic forms of mental models, conceptual models and semi-formal conceptual models? How can knowledge elicitation, analysis and implementation procedures be conducted with diagrams? Are diagrammatic representations suitable formalisms for building, or incorporating into, intermediate knowledge bases and computer knowledge bases? The focus of this paper will be on elicitation and analysis procedures from conceptual models and semi-formal conceptual models.

---

[1]Shaw and Woodward's term is *models of the conceptual models,* but *semi-formal conceptual models* is used as it is less cumbersome and somewhat more descriptive.

In the paper research on diagrammatic representations and reasoning that is relevant to DKA is first considered. The main body of the paper will consider the potential of two forms of diagrammatic representations for KA, as a means to being to identify some of the issues for DKA. Conventional Cartesian graphs and a special class of diagrams, called Law Encoding Diagrams (Cheng, 1994, 1995a, in press), will be considered. They are both well suited to expressing qualitative and quantitative relations among multiple variables. A discussion of the issues raised by the examples then follows, with brief consideration of how to develop methodologies and tools for DKA.

2 Diagrammatic Knowledge Representations and Reasoning

There has been substantial interest in reasoning and problem solving with diagrammatic representations in Artificial Intelligence, cognitive science and cognitive psychology. There have been various attempts to produce generative taxonomies of visual representations (e.g., Lohse *et al.*, 1994; Bertin, 1983), but these are too general to be useful for examining DKA. Work that can address components of Shaw and Woodward's (1990) framework will be briefly reviewed to provide some insights about the benefits of diagrams for KA and to flag some potential difficulties.

To begin, consider the nature of expert mental models. Although there has been much work on expertise (e.g., Chi *et al.*, 1988) and diagrammatic reasoning (e.g., Narayanan 1992), only a few studies of expert reasoning with diagrams have been conducted. Most notable is Koedinger and Anderson's (1990) work on *diagrammatic configuration schemas* (DCSs). DCSs are diagrammatic perceptual chunks. By comparing novice and expert geometry problem solvers, they found that structural constraints in the form of diagrammatic whole-part relations were used by experts as an efficient knowledge representations. Experts performed better, because they have schemas that combine diagrams with information relevant to common problem states (configurations). They solve problems by searching for DCSs that are applicable to the problem situation, matching parts of the given situation with the diagrams in DCSs. The information in the slots of active DCSs becomes available for inferences. Some implications for DKA flow from the existence of DCSs. By assuming that knowledge in a domain is encoded in DCS-like representations, this can guide and productively constrain knowledge elicitation, analysis and implementation. When an expert draws a diagram during elicitation, the general form of DCSs may be taken as a template within which to try to organise the expert's knowledge. During analysis DCS-like templates may suggests how to maximize the completeness of the knowledge, by explicitly identifying empty slots and seeking information to fill them. Given representations in the form of DCSs, the implementation of a KBS should be largely based around the search of the space of DCSs for those that are applicable to the target problem. These are fairly strong recommendations, but it should be noted that the applicability of Koedinger & Anderson's theory beyond geometry problem solving is yet to be fully demonstrated.

Moving from mental models to conceptual models and semi-formal conceptual models, consider the seminal paper on diagrammatic representations by Larkin &

Simon (1987), entitled 'Why a diagram is (sometimes) worth 10,000 words'. By building and comparing computational models of problem solving with diagrammatic and sentential representations, they demonstrated that there are benefits in the processes of search and recognition in problem solving with diagrams. In diagrams information needed for particular inferences is often found at the same location. This locational indexing makes the search for pieces of relevant data easier and reduces the effort to recognise what rules (productions) are applicable, by limiting the amount of symbolic matching that is required. This work is important to DKA, because it implies that consideration must be give to the way information is indexed within diagrams. Locational indexing is important, but this may not be apparent from experts' own descriptions of their diagrams or from observations of diagrams in use. Thus, during elicitation in which diagrams are drawn, care must be taken to understand how and where the experts' focus of attention resides in the diagram, at different stages in problem solutions.

On the machine side of the KA process, there are various lines of related work. There has been much work on graphical interfaces for knowledge engineering (see, Jones, 1988, for a review). Eisenstadt and colleagues (Eisenstadt *et al.,* 1990) use the term *visual knowledge engineering* (VKE) to refer to the synthesis of visual programming tools, program visualization tools, a sound knowledge engineering methodology, and AI programming support tools. The difference between DKA and these approaches to visual knowledge engineering is that DKA is attempting to obtain knowledge from experts that would normally be expressed in diagrams, and that may necessarily require diagrams. VKE, on the other hand, aims to re-present information, which is already expressed in a conventional propositional notation, in a graphical or visual format, to aid KBS development. Thus tables, abstract labelled networks and flow charts are common representations in VKE, but (Cartesian) graphs and structural or configurational diagrams will be used in DKA.

Of interest to DKA are techniques for parsing diagrams. Individual techniques exist for particular classes of diagrams, mostly network or node-link diagrams (e.g., Lutz, 1986). The development of techniques for parsing graphs and configurational diagrams would enhance diagrammatic elicitation, analysis and implementation procedures. However, the prospect of effective engines that would translate diagrams into standard propositional representations is remote, because of the high information content of diagrams and the wide variety of forms of problem solving they support, especially those that exploit perceptual inferences. An alternative is to directly incorporate diagrams in intermediate or computer knowledge bases. For intermediate knowledge bases a partial solution is to store diagrams alongside related knowledge from the target domain (perhaps as bit maps).

For computer KBSs there is work that suggests that systems may in the future use diagrammatic representations. For example, Koedinger & Anderson (1990) have built a simulation model that uses diagrammatic configuration schemas for geometry problem solving, to demonstrate the sufficiency of DCS. In addition to processes to search its space of schemas, the model also possesses the capability to parse a limited set of geometrical objects. Furnas (1992) has built a rule based system, with an architecture similar to a standard production system, that reasons with bitmaps. Rule conditions identify small bitmap patterns and rule actions specify changes to those patterns. Other problem solving systems that possess diagrammatic representations are described by Funt (1977), Novak (1977) and Shrager (1990). Unlike Furnas's system, which reasons with diagrams only, these other systems integrate propositional representations with the diagrams. This raises the important issue of multiple representations in expert reasoning, which was directly addressed by Tabachneck *et al.* (1994). They found that an expert uses the unique features of each representation for different purposes. Diagrams are used as place holders and to summarise information, whereas verbal expression are used to give semantic meaning and for making causal explanations.

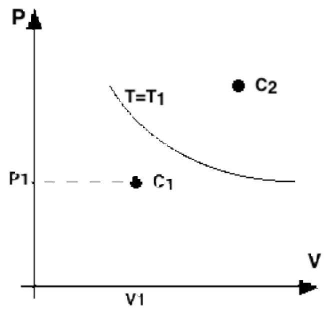

Fig. 1. A simple Cartesian Graph

Diagrammatic representations and reasoning has been studied in some detail, but the problem here is to understand the possible uses of diagrams in the processes of KA. Some of the research has direct implications for DKA, but this section has been somewhat speculative. In the remainder of the paper, the focus will be on concrete examples of diagrammatic representations and how they may be used for KA.

3 Elicitation with Cartesian Graphs

Cartesian graphs are a common form of diagrammatic representation, which can be found in a wide variety of domains. In Shaw and Woodward's terms, graphs are often used as conceptual models but are also commonly found as semi-formal conceptual models. For example, an expert's pencil sketch graph on the back of an enveloped to explain a process to colleague would be a conceptual model, but a printed graph in an plant operating manual is a semi-formal conceptual model. We will consider how the knowledge in such graphs can be elicited for the development of intermediate knowledge bases.

Typical Cartesian graphs define at set of axes so that points in a co-ordinate space can be identified. Figure 1 shows a *P*-*V* co-ordinate space (say pressure and volume). The magnitudes of a point in the graph is obtained by reading values off the appropriate axes; for example, for case C_1 the values are (V_1, P_1), as shown by dashed lines in the Figure 1. Families of points can be identified by the curves in the space; for example, the curve in Figure 1 is a contour for a particular value of another variable *T*. Similarly, families of curves can be shown by multiple lines in the graph; in Figure 2, three curves for different values of *U* are shown.

Simple graphs may be used in various ways for KA. Because the axes define a co-ordinate space, it is possible to distinguish different states. For example, suppose an expert sketches Figure 1 during a knowledge elicitation session. It is apparent that the space can be divided into two regions either side of the central curve, and that cases C_1 and C_2 fall in different regions. By further questioning the expert, it may be possible to find distinguishing characteristics for the regions, for instance, that points above the curve represent "critical" cases but those below are "normal" cases. There is information in the graph which may help to refine the distinction between critical and normal cases, by considering the shape and location of the curve in the co-ordinate space. Thus, critical cases are likely when both *V* and *P* are large, but unlikely when both *P* and *V* are small. Further, when either *V* or *P* is large but the other is small, there is a moderate possibility that the cases will be critical.

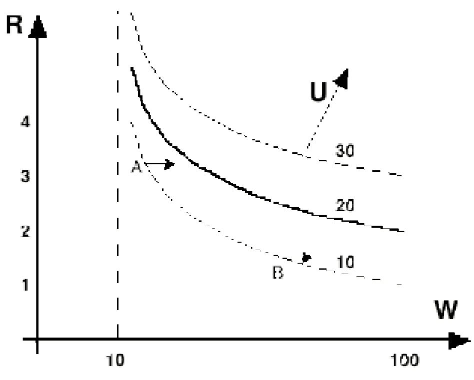

Fig. 2. A more complex graph

The purpose of a graph like Figure 1 may not be to distinguish different state spaces or regions, but could be to specify the relation between the variables *P* and *V*. If the graph is only a rough sketch, it is possible to infer that the *P* appears to be a monotonically decreasing function of *V*. If the graph is drawn more accurately, with scales to the axes, it may be possible use a curve fitting technique to find a equation describing the relation. Guidance about the appropriateness of different techniques can come by further questioning the expert. For example, a simple linear approximation may be sufficient when the curvature of the line is low and the expert is uncertain about its precise shape. Fitting polynomials to curves is a somewhat more sophisticated option for graphs that experts can draw more accurately. With further elicitation, it may be possible to select a suitable order for the polynomial, using the expert's knowledge of the number of characteristic operating modes of the target system, which might correspond to the number expected roots. If the expert says that the relation is periodic, then Fourier analysis may be used. The usefulness of the outputs of such analytical techniques for the eventual KBS will depend on the nature and tasks of the domain.

A greater variety of knowledge can be obtained from more complex graphs, like Figure 2, which an expert may have carefully drawn or that might be found in the documentation for a domain. In Figure 2 three curves for particular values of *U* are shown in a *W-R* co-ordinate space. The values on the axes permit quantitative facts to be inferred from the graph; for example, when *W*=100, it happens that *R* and *U* are linearly related. Although the graph's axes indicates that *R* and *U* continue to infinity, we may take the values on the axes as delimiting likely ranges of values that are of interest in the domain, unless otherwise specified. The graph not only gives the relations among variables but also the ranges over which the relations are known to hold. For example, the graph shows that there is a limiting minimum value of

W(=10), when *R* is large. It may be the case that there is a relation for *R* and *W* to the left of the asymptote, but this cannot be inferred from the graph.

Suppose that the line for *U*=20, in Figure 2, is a boundary between critical and normal cases, as in the previous graph. In this graph the contours for values of *U* provides useful information about the proximity of particular states to the boundary. For example, an expert may indicate that the area between curves for *U*=10 and *U*=20 should be treated as a warning zone. A rule may be formulated which states that if the *R* and *W* are such that the warning zone is entered, then corrective action should be taken to prevent the system moving nearer criticality. Further, given the shape of the warning zone it is possible to infer, when *R* is constant and *U*=10, that a small change in *W* is more dangerous at point *A* than it is at *B*, as shown by the arrows in Figure 2. For case *B* action must be taken promptly, even for small increases of W.

The examples illustrates some of the forms of knowledge that are encoded in graphs and that may be elicited using them. These include: qualitative relations among variables; rates of change of variables with respect to each other; algebraic approximations to relations; different systems states or cases of phenomena, plus boundary conditions for them; and, sensible ranges over which to consider the variables, for which the relations are valid. Although all these forms of information can be obtained by conventional propositional KA techniques, there are benefits in using a graph. In graphs, as with most diagrams, information that is often needed for particular inferences is often found at the same location in the diagram (Larkin and Simon, 1987). Typically information is neatly integrated within a diagram, so it is likely to be easier to obtain related information from a graph or diagram than from propositional knowledge structures. Experts can effectively use the conventions for constructing graphs as constraints to produce such compact and coherent packages of knowledge. Some properties of diagrams are perceptually obvious and require little effort to identify, such as the gradients at different points on a curve. Further, important information in diagrams often appears as emergent features (Koedinger, 1992), such as the set of curves tending towards an asymptote in Figure 2. This means acquisition of derived or inferred knowledge can be easier with diagrams than propositional representations, which may require more cognitive effort.

We now move from considerations of elicitation with single graphs to deal with knowledge analysis with multiple graphs.

4 Analysis with Cartesian Graphs

Cartesian graphs are a flexible yet formal form of knowledge representation, so they can also be used in knowledge analysis procedures. Such procedures are used for the refinement of intermediate knowledge bases and for the development of computer knowledge bases. An example is examined in which graphs are used for consistency and completeness checking.

Consider a simple hypothetical example of KA for the dynamic control of a fermentation process, in which the yield rate, *Y*, of the process is proportional to the product of three variables, say, temperature deviation from some optimum, *dT*, nutrient concentration, *C*, and oxygen concentration, *G*. Suppose the expert has

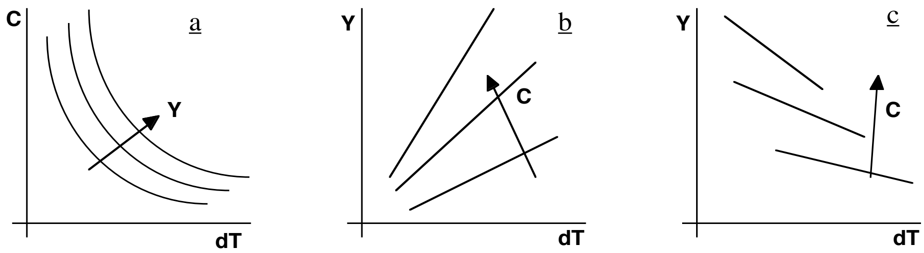

Fig. 3. Different perspectives in multiple graphs

during an initial elicitation phase identified Y, dT and C as pertinent variables, and has drawn different graphs to show the relations among the variables under various conditions. During an analysis phase the consistency of the knowledge may be tested by examining the form of the graphs. If the expert had drawn Figure 3a and 3b for the process then there is no problem, because the graphs are consistent; both show that $Y \quad C \cdot dT$. However, if Figures 3a and 3c were drawn, then comparing these graphs reveals a difficulty, because Y and dT do not both increase together (with C constant). There are at least three possible reasons for the inconsistency. First, one of the graphs could be wrong. On telling the expert about the problem, the expert may simply recognise that one of the graphs is in error. Second, the model of the process may be incomplete, perhaps a hidden variable is influencing the process. For example, under Figure 1c the oxygen concentration, G, may be to high, so causing temperature dependent aerobic degeneration of the product to occur. Finally, the graphs may not apply over a consistent range of one of the variables, but by extending that range the incompatibility might be resolved. For example, all three graphs in Figure 3 would be consistent, if Figures 3b and 3c apply disjointly to low and high value ranges of dT, respectively. This would imply that Figure 3a is only showing the left hand side, low dT values, of a larger U-shaped curve. More direct inferences could be made about the ranges given information regarding the values on axes of the graphs.

Analysis in DKA is feasible with Cartesian graphs. The potential benefits are similar to the benefits noted for knowledge elicitation. Graphs allow powerful perceptual inferences to be made, such as visually joining Figure 3c to the right of Figure 3b and imagining that the sets of lines form inverted 'U' shaped curves. Equivalent inferences can be made with algebraic representations of the processes, but it is not simple to generate a formula that will have the correct form (shape) for the full range of all three variables. Similar, it is likely that inconsistencies will be more readily spotted in a set of graphs than in a set of abstract equations or other propositional representations. Graphs are a common form of representation, that most experts may find more natural to use than existing propositional representations in KA.

Graphs are not the only form of diagrams. The next two sections consider a class of diagrams that may provide other benefits to KA.

5 Law Encoding Diagrams

This is a class of diagrammatic representations with some interesting properties and is used for various forms of expert problem solving. Problem solving and learning with Law Encoding Diagrams (LEDs) has been studied empirically (Cheng 1994, 1995a). The role of LEDs in some important discoveries in the history of science has been investigated and computationally modelled (Cheng, in press; Cheng and Simon, 1995). Interactive computer based discovery learning environments have been built using LEDs for various domains (Cheng 1995a, 1995b). This section describes LEDs and considers some of the forms of problem solving that may be done with a particular LED. The next section considers how LEDs may be used for KA.

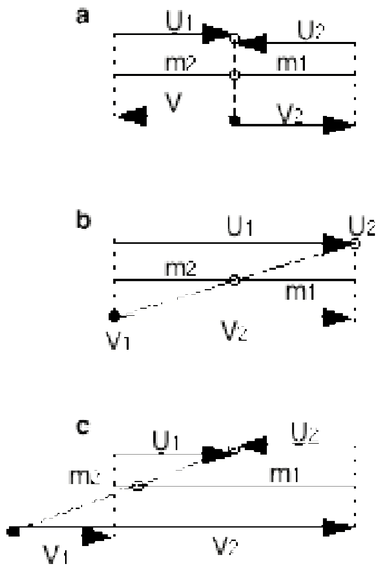

Fig. 4 . 1DP Diagrams

A Law Encoding Diagram is a representation that encodes the underlying relations of a law, or a system of simultaneous laws, in the structure of a diagram by the means of geometric, topological and spatial constraints, such that each instantiation of a single diagram represents an instance of the phenomenon or one case of the law(s).

LEDs are effective for various forms of reasoning and problem solving, because they are representations at an intermediate level abstractions, which bridge the conceptual gulf between abstract general laws and descriptions of the behaviour of phenomena. LEDs are specialized computational devices for particular domains, that exploit the benefits of reasoning with diagrams, whilst ensuring that problem solving is grounded in the correct laws of a domain.

Figure 4 shows three examples of a LED, call the One-dimensional property diagram, 1DP diagram. Each 1DP diagram represents a single one-dimension head-on collision between two elastic bodies. In this domain there are six variables of interest, which are represented by diagrammatic elements in 1DP diagrams. The lines (vectors) u_1 and u_2 are the velocities of the two bodies before impact, and the lines v_1 and v_2 are the velocities afterwards. The subscripts indicate the two bodies, body-1 and body-2. The masses of the bodies are shown by the lines m_1 and m_2. Their relative sizes show the magnitudes of the variables. Each instantiation of a LED, a single diagram, represents one instance of the phenomenon; the 1DP diagrams in Figure 4 depict three collisions. For example, in Figure 4a the two bodies have the same mass and bounce off each other with the same speeds but in opposite directions.

The collisions domain is not one that we would attempt to learn about by doing KA, as it is already well understood in physics. However, it is a good domain to introduce LEDs, to show some of their properties and demonstrate the kinds of reasoning that can be done, and thus to see some of their potential for DKA. The underlying laws of the domain are momentum conservation and energy conservation. An impression of the complexity of relations that may encoded in LEDs can be seen by inspecting the algebraic forms of the laws;

$$m_1u_1 + m_2u_2 = m_1v_1 + m_2v_2 \ , \qquad \ldots 1$$

and

$$\tfrac{1}{2}m_1u_1^2 + \tfrac{1}{2}m_2u_2^2 = \tfrac{1}{2}m_1v_1^2 + \tfrac{1}{2}m_2v_2^2 \ , \qquad \ldots 2$$

respectively.

Laws are captured in LEDs by rules that constrain the geometric, spatial or topological structure of the diagrams. There are three such law-encoding constraints for the 1DP diagram. (i) The tails of the initial velocity arrows and the tips of the corresponding final velocity arrows must be in line vertically, so the total length of the u_1—u_2 line equals the total length of the v_1—v_2 line. (ii) The total length of the mass line, m_1—m_2, equals the length of the velocity lines. (iii) The small circles indicate the ends of the lines that are not fixed by the previous two constraints; the circles must lie in a straight vertical or diagonal line.

Cheng (1995a) describes the forms of reasoning and problem solving that are feasible with LEDs. These include: quantitative calculations; qualitative reasoning; configurational reasoning; debugging and trouble shooting; reasoning about extreme cases; analysis of complex interactions; extending LEDs to incorporate wider ranges of phenomena; and, high level conceptual explanations, which use notions of physical and temporal symmetry. Two examples of problem solving with 1DP diagrams are considered to give an impression of the capabilities of LEDs.

A common problem, especially in text books, is to find values of dependent variables (unknowns) given the values of independent variables (knowns). For example, what are the final velocities of two bodies in an impact, given that they approach from opposite directions with the same speed, but one is ten times heavier than the other (m_1/m_2=10)? Quantitative problems are relatively easy to solve with 1DP diagram, by drawing a diagram to scale, using the constraints given above, and measuring the required values. Figure 4c shows the diagram for the given problem, with the mass line divided in the ratio 10:1, showing the final velocities are v_1=0.6 and v_2=2.6. Given the simple geometric constraints that define 1DP diagrams, it is possible to compute precise answers using simple geometry. It is as easy to solve quantitative problems for other combinations of unknown variables, such as m_1 and m_2 given the 4 velocities.

For a representation to be generally useful in the present domain it should cope with collisions involving three or more bodies. For example, how can the behaviour of the balls in Newton's Cradle be explained (the executive toy with five suspended balls). What will happen when two balls hit three coming in the opposite direction? 1DP diagrams apply only to pairs of colliding bodies, so it initially seems that they cannot be used to explain complex interactions in Newton's cradle. However, the idea of pairs of impacts suggests that 1DP diagrams might be used in a (de)compositional manner, by treating the whole as a series of independent pair-wise collisions. This is

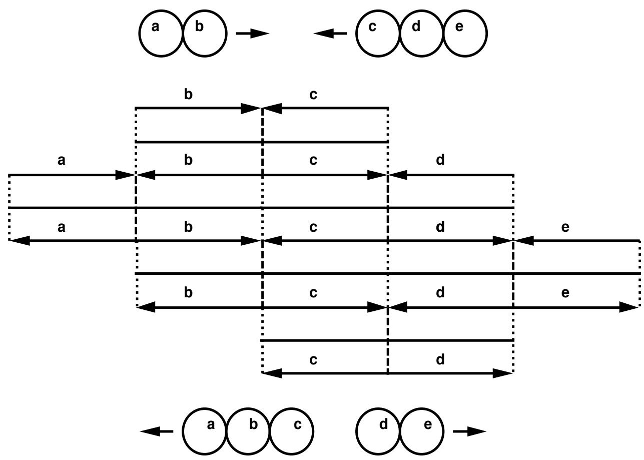

Fig. 5. Explaining Newton's Cradle

possible because a single diagrammatic element may be shared in more than one 1DP diagram. Figure 5 is a diagram for the given problem composed of six 1DP diagrams from Figure 4a. The first row shows ball *b* hitting ball *c*, and rebounding equally. The second row shows balls *a* and *b* colliding and similarly balls *c* and *d*. The collisions propagate until finally balls *a*, *b* and *c* go in the opposite direction to balls *d* and *e*. This is a good example of the compositionality of LEDs, a property that will be useful in DKA.

The two examples of problem solving with 1DP diagrams illustrate how LEDs can be the basis for some types of expert problem solving. A partial explanation of the effectiveness of LEDs, which is somewhat different to those mentioned above, is to view LEDs as perceptual chucks, similar but not identical to Koedinger & Anderson's diagrammatic configuration schemas. Although the DCS have only, so far, been found in geometry problem solving, it seems that much of Koedinger & Anderson's findings are applicable to LEDs, as they employ significant geometric constraints. Given the similarity between LEDs and DCSs, the observation in section 2, that it might be appropriate to use DCSs for DKA, may also be valid for LEDs. The next section considers DKA with LEDs.

6 LEDs for KA

The 1DP diagram is a highly specialised LED for a particular set of relations (Equations 1 and 2), so it is not suitable for general use in KA. There are other classes of LEDs that may be used, although one will be considered here — Algebra Triangle, AT, diagrams. (A specialised subset of AT diagrams have been used in a computer based tutoring system for electrical circuits, which allows the user to interactively manipulate the diagrams on screen, Cheng, 1995b.)

Figures 6a and 6b show the basic form of AT diagrams, using the fermentation example introduced in section 4 above. The height of the triangle is the value of *dT* and the length of the base is *Y*. The bold line beginning at the apex is a unit line and at its other end is a perpendicular line, *unit perpendicular*, whose length is *C*. These AT diagrams encode the relation $Y=dT \cdot C$. Thus as *C* or *dT* independently increase *Y* will increase. Each diagram represents one set of values of the variables; in Figure 6a *dT* is greater than unity but in Figure 6b it is less than unity. Though mathematically equivalent, it is possible to conceptualize these AT triangles as encoding the relation $C=Y/dT$. Thus, we would consider that *C* will increase as *Y* increases with *dT* held constant, but *C* will decrease as *dT* increases with *Y* constant.

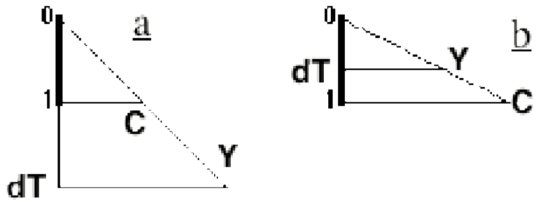

Fig. 6. AT diagrams

DKA with this class of LEDs could progress by the identification or construction of AT diagrams, to find structures in which the changes to the variables are consistent with the expert's expectations. It is assumed that the AT diagrams are directly manipulable objects on a computer screen. For the fermentation process example, an expert would initially identify the relevant variables and attempt to map them into AT triangles. Knowing how AT triangles encode qualitative relations among the variables, such as those mentioned in the previous paragraph, this would suggest which variables should be mapped onto the height, base and the unit perpendicular. The acceptability of the mapping can be checked by generating different cases, particularly extreme ones, to see whether they match the expert's predictions. Given an acceptable AT diagram, it is possible to infer a variety of rules to be added to an intermediate or computer knowledge base. A rule could be quantitative/algebraic, if the expert is able to make a close numerical match between the AT diagram and the process (e.g., If the process is fermentation and *dT*=<u> and *C*=<v>, then *Y*=<u>*<v>). Alternatively, if the AT diagram is only an approximation to the behaviour of the process, qualitative rules could be inferred (e.g., If the process is fermentation and *Y* is increasing and *dT* is decreasing, then *C* is rapidly increasing).

Like the 1DP diagrams, AT diagrams can be used in a compositional fashion to deal with complex interactions or relations. As the lines in the AT diagrams represent variables, it is possible for a line to be shared by more than one AT triangle. Suppose an expert introduces a new variable, *G*, to be incorporated into the model, Figure 7

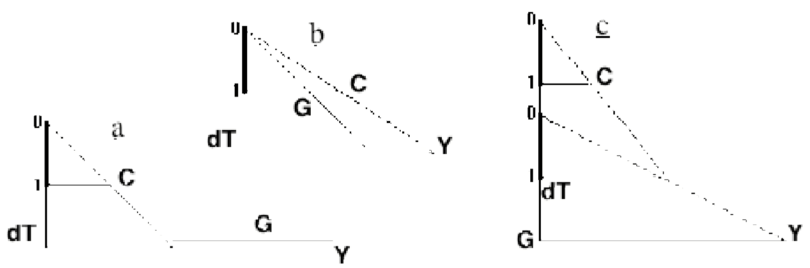

Fig. 7. Composite AT diagrams

shows three examples of composite AT diagrams that include G. Algebraically the relations that they encode are (a) $Y=dT \cdot C+G$; (b) $Y=dT(C+G)$; and, (c) $Y=dT \cdot C \cdot G$. This is a small sample of the possible AT diagram structures for four variables, but the examples give an impression of the expressive power of this class of LEDs for encoding relations. Again, by interactive manipulation of the diagrams, experts would be able to judge whether the changes among the variables are consistent with their knowledge of the target process. Quantitative or qualitative rules may be inferred directly from the structural form of the composite AT diagrams.

The examples demonstrate the feasibility of KA with LEDs, but the full potential for elicitation and analysis with AT diagrams, and LEDs in general, is yet to be investigated. However, it is possible to make some predications of how and why LEDs may be effective for KA based on the existing research with LEDs for problem solving and learning. The fundamental purpose of LEDs is to make the important relations of a domain readily accessible using the structure of diagrams. The potential uses and benefits of LEDs flow from this. LEDs may be effective for: (i) finding qualitative relations among variables; (ii) building algebraic models of systems; (iii) examining extreme and special cases; (iv) identifying general constraints among variables; (v) exploring rates of change among variables; and, (vi) if scales are provided for their elements, LED may be used for quantitative problem solving. As the relations among the variables are apparent from the structure of the diagram, it is likely that LEDs will be good for eliciting and analysing the forms of knowledge just listed. The cognitive effort required to generate and make inferences with LEDs is less than that of other representations, because they are diagrammatic and are at an intermediate level of abstraction. They are models that bridge the gap between abstract relations and concrete instances. The rules governing the structure of composite AT diagrams are simply defined and can be translated into algebraic formulas in a routine and straightforward manner. Particular instances can be examined by manipulating the shape of a LED within its diagrammatic constraints and unusual cases spotted as they appear as distinctive patterns.

7 Discussion: Research Issues for DKA

The possibility of DKA has been demonstrated using graphs and AT diagrams and some of the potential benefits have been discussed. Research on DKA is in its infancy, but the examples and the review of research on diagrammatic reasoning highlights various research issues for the field. This section deals with five of them.

The first issue concerns the need for tools to support DKA. KA can be greatly facilitated by computer based tools, which may improve the effectiveness of construction and revision of knowledge bases. In the case of DKA, tools will be particularly important, because diagrams are external representations that require a physical medium in which to be drawn. Tools will greatly reduce the effort needed for the drawing and re-drawing many different versions of graphs or composite AT diagrams. Direct on screen manipulation of LEDs will be essential to enable experts to simply examine and refine the diagrams. DKA tools are quite feasible with modern computers and can draw upon the same technology used in existing discovery environments that deploy interactive LEDs for learning (Cheng, 1995a, 1995b). For example, in a graph based DKA tool, should the users be (i) allowed to sketch curves on screen, or (ii) should they select curves from an on-line library? The first option gives the user greater flexibility, but requires computational machinery to interpret and encode the shape of the curves. In the second option, pre-coded descriptions of the curves could circumvent some of the problems of on line parsing of the graphs. However, this option relies on the library containing all the conceivable forms that a user might require.

Following DKA tools is the issue of devising effective ways to use the tools. Sketches of possible procedures were given in the above examples, but theoretically and empirically founded methodologies are needed. A systematic approach to the analysis of graphs is required and could be loosely based on the guidelines for structured interviews (mentioned by Shadbolt & Burton, 1989). These guidelines or protocols might consist of series of queries to the expert, to ensure the correctness and completeness of the information needed to interpret graphs. The protocols will need to consider the nature of the axes, the form of the curves, the role of other features in the co-ordinate space (e.g., asymptotes), and typical cases of reasoning with the graph. For example, take the axes, it is necessary to know whether the expert has assumed that they are linear or logarithmic, whether the origin coincides with zero values of the variables, and whether positive changes map onto upward and rightward movements along the axes.

A general protocol for DKA would be useful, but the diversity of diagrammatic representations makes this a remote possibility. An alternative is to decompose diagrams into separate problem solving functions and develop protocols associated with each function. A set of such functional roles of diagrams has been devised by Cheng (1996b). Previous work on reasoning with diagrammatic representations, see section 2, will also help to guide the development of methodologies. For example, perceptual schemas may be used as templates to encode knowledge associated with emergent features at particularly in locations in a diagram.

The third issue is the need for procedures to generate propositional rules from diagrammatic representations for use in existing KBSs. The translations will need to maintain the accuracy, completeness and consistency of the expressions generated, with respect to the original diagram. Some of the kinds of information that are available from graphs was considered in sections 3 and 4 above, but there are others, including: areas "under" the curve to the abscissa or ordinate (integration); interpolation to points between curves; possible extrapolation beyond the given ranges; and, the existence of symmetries which are often useful in reasoning. Achieving an appropriate balance the breadth of information and economy of knowledge in the translation is hard problem. Methods to match particular kinds of information in diagrams to specific forms of problem solving will have to be considered.

The fourth issue concerns the use of special representations for DKA, such as AT diagrams. The cost of learning to conceptualizing relations in terms of unfamiliar representations, implemented as contrived (rather than natural) elicitation techniques, may outweigh the potential benefits of having DKA tools in the first place. However, the diversity of kinds of diagrams suggests the possibility of marrying particular diagrams to different forms of knowledge or problem domains, or even to cope with the preferences of different experts. For example, a knowledge engineer might use AT diagrams with experts who are familiar with geometric reasoning, but adopt a graph based DKA tool for those who are not.

The final issue is mentioned to acknowledge its importance in the long term, but it is not considered in detail. This is the role of role of multiple representations in expertise. In particular, the integration of diagrams and propositional knowledge will be necessary for complete diagrammatic KBSs. Investigating DKA in relative isolation is the first step towards the more complex issue of integrating representations.

8 Conclusion

Diagrammatic knowledge acquisition has been somewhat neglected in KA. Here the feasibility of DKA has been shown, with some support from previous work on diagrammatic representations and reasoning. The identification and discussion of issues for study provides a context for future research. The development of graph and LED based KA tools, to investigate the issues discussed in the paper, has just been awarded funding. The construction of graph and LED based tools will involve the modifications to existing mechanisms used in interactive graphical learning environments. The work will be mainly focused on procedures for knowledge elicitation and knowledge analysis (to use Shaw and Woodward's conceptualization for the last time).

References

Bertin, J. (1983). *Seminology of Graphics: Diagrams, Networks, Maps.* Wisconsin: University of Wisconsin Press.

Cheng, P. C.-H. (1994). An empirical investigation of law encoding diagrams for instruction. In *Proceedings of the 16th Annual Conference of the Cognitive Science Society*. (pp. 171-176). Hillsdale, NJ: Lawrence Erlbaum Associates.

Cheng, P. C.-H. (1995a). Law encoding diagrams for instructional systems. *Journal of Artificial Intelligence in Education*, **6**(4).

Cheng, P. C.-H. (1995b). *AVOW tutor version 1.0: Instruction manual* (Technical Report No. 31). ESRC Centre for Research in Development, Instruction and Training, University of Nottingham.

Cheng, P. C.-H. (in press). Scientific discovery with law encoding diagrams. *Creativity Research Journal.*

Cheng, P. C.-H. (1996a). Problem solving and learning with diagrammatic representations. In D. Peterson (Eds.), *Forms of Representation* Intellect Books.

Cheng, P. C.-H. (1996b). Cognitive analysis of the functional roles of diagrams. Submitted to the *18th Annual Conference of the Cognitive Science Society.*

Cheng, P. C.-H., & Simon, H. A. (1995). Scientific Discovery and Creative Reasoning with Diagrams. In S. Smith, T. Ward, & R. Finke (Eds.), *The Creative Cognition Approach* (pp. 205-228). Cambridge, MA: MIT Press.

Chi, M. T. H., Glaser, R., & Farr, M. J. (Eds.). (1988). *The Nature of Expertise*. Hillsdale, NJ: L. Erlbaum Associates.

Eisenstadt, M., Domingue, J., Rajan, T., & Motta, E. (1990). Visual knowledge engineering. *IEEE Transactions on Software Engineering*, **116**(10), 1164-1177.

Funt, B. V. (1977). Whisper: A problem-solving system utilizing diagrams and a parallel processing retina. In *Fifth International Joint Conference on artificial Intelligence*, (pp. 286-291). San Mateo, CA: Morgan Kaufmann.

Furnas, G. W. (1992). Reasoning with diagrams only. In H. Narayanan (Eds.), *AAAI Technical Report on Reasoning with Diagrammatic Representations (Report No. SS-92-02)* (pp. 118-123). Menlo Park, CA: AAAI.

Jones, S. (1988). Graphical interfaces for knowledge engineering: an overview of relevant literature. *Knowledge Engineering Review*, **3**(3), 221-248.

Koedinger, K. R. (1992). Emergent properties and structural constraints: Advantages of diagrammatic representations for reasoning and learning. In N. H. Narayanan (Eds.), *AAAI Technical Report on Reasoning with Diagrammatic Representations (SS-92-02)* Menlo Park, CA: AAAI.

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

Kulpa, Z. (1994). Diagrammatic representations and reasoning. *Machine Graphics and Vision*, **3**(1/2), 77-103.

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

Lohse, G., Biolsi, K., Walker, N., & Rueter, H. (1994). A classification of visual representations. *Communications of the ACM*, **37**(12), 36-49.

Lutz, R. (1986). *Diagram parsing - A new technique for artificial intelligence* (Research Paper No. CSRP.054). School of Cognitive and Computing Sciences, University of Sussex.

Narayanan, H. (Ed.). (1992). *AAAI Technical Report Reasoning with Diagrammatic Representations (SS-92-02)*. Menlo Park,CA:AAAI.

Novak, G. S. (1977). Representations of knowledge in a program for solving physics problems. In *Proceedings of the 5th International Joint Conference on Artificial Intelligence*, (pp. 286-291).

Shadbolt, N., & Burton, M. (1989). Knowledge elicitation: a systematic approach. In J. R. Wilson & N. E. Corlett (Eds.), *Evaluation of Human Work* London: Taylor & Francis.

Shaw, M. L. G., & Woodward, J. B. (1990). Modeling expert knowledge. *Knowledge Acquisition*, **2**(3), 179-206.

Shrager, J. (1990). Common-sense perception and the psychology of theory formation. In J. Shrager & P. Langley (Eds.), *Computational Models of Scientific Discovery and Theory Formation* (pp. 437-470). San Mateo, CA: Morgan Kaufmann.

Tabachneck, H. J. M., Leonardo, A. M., & Simon, H. A. (1994). How does an expert use a graph? A model of visual and verbal inferencing in economics. In A. Ram & K. Eiselt (Eds.), *Proceedings of the 16th Annual Conference of the Cognitive Science Society*. (pp. 842-847). Hillsdale, NJ: Lawrence Erlbaum Associates.