Assigning information to modalities: comparing graphical treatments of the syllogism

Keith Stenning and Richard Tobin
Human Communication Research Centre
Edinburgh University
*

February 7, 1994

1 A Cognitive Theory of Modality Assignment

The goal of a cognitive theory of modality assignment is to give a general account of the cognitive effects of assigning information to different modalities. This is an ambitious goal, some might say too ambitious, but without some general framework, the multi-disciplinary effort which will undoubtedly be necessary for progress is likely to remain scattered and unrecognised. A more modest aim of this paper is to illustrate how at least the disparate techniques of semantic analysis and experimental psychology are implicated into the development of a general theory.

Expressing the general goal also makes it clear how important a practical problem this is. New technology is making available new media and modalities of information presentation and new modes of interaction with information, but the technology makes the modality assignment issue critical. First, the new opportunities include assignments of information to modalities which are actually pathological, and the craft design skills which avoid these pitfalls in established technologies have not yet developed. Second, the cost of authoring presentations in the new technologies is very

*The support of the Economic and Social Research Council UK (ESRC) is gratefully acknowledged. The work was partly supported by the ESRC funded Human Communication Research Centre. Partial support from the following grants is also acknowledged: Grant # 9018050 from the Joint Councils Initiative in Cognitive Science and HCI; and Grant # GR/J58612 .

high and so any body of knowledge which helps avoid poor designs or speed up the process will be useful. We may expect the development of theory in this area to interact strongly with empirical experience of developing design in the new technologies and the interaction to benefit both.

The kernel of the theoretical approach is the idea that when graphical presentation is useful it is because graphics are logically *inexpressive* but therefore tractable. The observation that graphics have this property of enforcing information representation which leads to concreteness or inexpressiveness is classical (see e.g. Berkeley 1710). The novelty is to combine it with modern computational complexity theory to explain why concreteness leads to usefulness. There are several parallel developments of related ideas in the literature. A notable one in logic is Levesque's (1988) concept of 'vividness' and his use of it to study the computational complexity of logical fragments. In AI, Larkin & Simon's (1987) explanations of the usefulness of graphics in terms of the different computational characteristics of representations which are informationally equivalent is also related. Our approach seeks to explain the computational differences in terms of differences in the expressiveness of the *systems of representation* in which the token representations are embedded. The general semantic approach is developed in more detail in Stenning & Oberlander (1991) and its relation to the psychological literature on comprehension in Stenning & Oberlander (1994).

A theory of modality assignment will need a robust grasp of what modalities are, and this conceptual requirement is plagued by a rich set of cross-disciplinary confusions of current terminology. The first section establishes the terminology which will be used here, and then goes on to summarise an account of the differences between graphics and language as modalities of information presentation. This conceptualisation and semantic account is developed in more detail in (Stenning, Neilson and Inder 1993). The section ends by sketching a methodology for approaching a general theory of modality assignment through studies of example graphical and linguistic systems. This methodology combines semantic analysis of graphics with experimental psychological evaluation of the cognitive impact of contrasting presentations of the same information.

The remainder of this paper consists of semantic presentations and comparisons of four treatments of the categorial syllogism embedded within our framework of studying information enforcement. Some available evidence of their usability will be briefly mentioned but the focus of the paper will be on the different semantic/computational properties of the systems. In fact carefully controlled behavioural evaluation of alternative presentations is currently scant. We will argue that this is precisely because the required semantic analysis and conceptual frameworks have not been available to psychologists for use in the design of systematic evaluations.

The categorial syllogism is neither at the forefront of logical development nor at the cutting edge of pedagogical practice, and so it may appear a strange choice of theoretical testbed for a theory of modality allocation. The major part of the choice is of course driven by the need for a domain which is both susceptible to semantic analysis and tractable for behavioural evaluation. Syllogisms have a long history of graphical pedagogy, and have received much attention from psychologists who have formulated theories of the 'internal' mental representations involved in students reasoning (when no external medium is available) in terms of theoretical appeals to graphical methods. Other psychologists have claimed to be able to demonstrate that such appeals must be wrong on the basis of arguments about graphical semantics. It turns out, however, that these arguments are based on mistaken semantic analyses of the graphical methods. So here is a domain in which we can make a good case that it really is important to get the semantic details right in order to be able to develop or apply cognitive theory.

The syllogism is also a sufficiently simple domain that we can invent systems which have certain semantic properties which our theories make predictions about and hope to evaluate them with reasonable resources. Of the four treatments considered here, three are inventions made for the purpose of developing this theory. When this approach is possible it becomes painfully obvious that simplistic ideas that graphical presentation is 'good' because it is natural, or analogical, or concrete, or vivid or ... are exposed for what they are—intuitions probably containing some germ of truth but wholly unsatisfying as cognitive explanations.

However, we do not adopt the syllogism purely for theoretical convenience. The domain is extremely simple but it is nevertheless sufficient to bring out many of the most fundamental problems for students learning about deductive discourse. Because of their simplicity, syllogisms allow us to proceed very rapidly from the teaching of object language inference to meta-language principle—very much more rapidly than is possible in more serious logical languages. Much of the cognitive interest of logic, and pedagogical interest for that matter, lies in these general ideas of object and meta levels and the selection of appropriate representations and reasoning methods. It is no accident that the categorial syllogism was historically the first area of logic in which the very idea of a meta-theory was developed (see Lear 1980 for development of this theme). We are concurrently engaged in studies of graphical methods of teaching first order logic (Cox, Stenning & Oberlander (1994)), but syllogisms are a more appropriate domain for the focus of the current study—the comparison of alternative graphical treatments.

2 Modalities of information presentation

In attempting to integrate ideas about information presentation from logic and computer science with those from psychology we encounter severe terminological conflict. Psychology uses the term *sensory modality* for what computer science uses the term *information medium.* Logic uses the term *modality* for a variety of uses which relate to the *interpretation* of signs rather than their *perception.* HCI researchers often refer to systems which can process both speech and written language as multi-*modal* systems, whereas from a logical point of view speech and text are generally (though not unproblematically) to be seen as information in the different media but in the *same* modality.

There are good intra-disciplinary reasons for all these usages but the situation demands some definition of terms if communication is to succeed. We adopt a usage in which *media* are the physical/perceptual aspect of representation systems, and *modalities* are kinds of interpretation function. Braille and printed text are the same modality of information in different media (tactile and visual), whereas embossed diagrams and braille contrast in their modality, one being graphical and the other linguistic. The same modality contrast in a different medium holds between printed diagrams and text.

This sense of medium is perhaps closer to the computer scientists', but retains the cognitive requirement for abstraction over physical media along perceptual lines. Writing in smoke trails in the sky, ink on paper and electrons on phosphor are all read visually and though they may differ in their ergonomic detail, an abstract interpretation of 'medium' allows us to capture their commonalities.

It is much harder to give an adequate working definition of the concept of modality. We can cite the difference between language and diagram as a paradigm example of a modality contrast, but a succinct account of the difference is not so easy. A deeper analysis in fact reveals that our intuitive grasp of what is graphical and what is linguistic draws a somewhat different line than the one we need for theoretical purposes. The most expressive 'graphics' (systems such as unrestricted semantic networks) are best treated as languages, whereas the least expressive languages are best treated as essentially 'graphical'. Our concept is based on the directness and uniformity of the semantic interpretation of representing relations in graphics contrasted with the indirectness of interpretation in languages that have abstract syntax. The mediating abstract syntax leads to heterogeneous interpretation of the same representing relation in such languages.

This account is developed in more detail in Stenning, Neilson & Inder 1993) but here we can give examples which should suffice as illustrations for our current purposes. To take the linguistic case first, in languages a single spatial/temporal relation is uniformly interpreted as concatenation. On the basis of this relation, and the classification of concatenated symbols into types, strings of symbols are parsed into an abstract syntax which then defines the many different relations which are semantically interpreted. Hence a single representing relation (concatenation) is heterogeneously semantically interpreted. In the degenerate case of finite state generated or regular languages which have no abstract syntax, concatenation is given direct semantic interpretation (generally as temporal succession). It is notable that these 'languages' exhibit information enforcement (they cannot represent anything other than total orderings of event tokens) which is the property of graphics on which our theory focuses.

In graphical representations, where no abstract syntax intervenes, information enforcement is the rule. Introducing a symbol representing an element will typically determine the relations of that element to all other elements already represented. Most graphical systems will employ more than one directly and uniformly semantically interpreted relation unlike the single but variously interpreted concatenation relation in languages. Our claim is that usability of representations can be related to the tractability of inference that this information enforcement brings with it. To support this claim we need to examine alternative representations of the same information and assess their usability for performing the same task.

2.1 A methodology for investigating modality effects

Predictions of peoples' performance in tasks using different modalities is a complex affair affected by many factors other than the computational expressiveness of the representations employed.

Making predictions of cognitive usefulness of modalities requires us to relate representations to the information processing tasks they are to be used for. If the benefit of graphics is that enforcement of information representation means inexpressiveness and inexpressiveness means tractability, this will have different impacts on usefulness for different tasks. A task that requires no abstraction will not suffer because a representation cannot express this unnecessary abstraction, and will gain from its tractability. A task that requires arbitrary and unpredictable abstractions will suffer if some of them are unavailable within a representation system.

Although graphical systems enforce information representation to some extent, they generally can express some abstractions. Stenning & Oberlander (1991) classify systems into minimal abstraction representation systems (MARS), limited abstraction representation systems (LARS), and unlimited abstraction representation systems (UARS). The most fruitful ground for developing the theory will evidently be the middle ground in which tasks demand some abstraction, and graphics can express some but not all abstractions. Where the match is good we would expect graphics to prove useful. The study of the 'tricks' whereby actual graphical systems do express albeit limited abstractions reveals how important it is to have an accurate characterisation of the semantics of these systems. As we shall see in the case of the syllogism, quite subtle changes in the semantics can radically affect the tractability of a task by allowing or disallowing the abstractions required. Quite apart from task variables there are differences between people in their degree of experience and their cognitive styles which may make very gross differences to the usefulness of modalities, and it is well known that these factors interact with the task variables mentioned above.

One way of avoiding some part of all this complexity it to study domains and tasks in which the information expressed by representations in different modalities can be accurately equated; where there are reasonably measurable indices of performance (and understanding); and where the level of expertise can be controlled. Having found such a domain we can then compare the processing of the same information during the same task presented in different modalities by users of similar degrees of expertise, and we can at least analyse results to see if we find uniformity from individual to individual.

It is an added advantage if we can find domains in which carefully informationally equated contrasting modalities of information presentation are both plausibly useful. If both are intuitively plausible (if different peoples' intuitions differ somewhat about usefulness, still better) then there is a good chance that more careful study will reveal task differences or cognitive styles which discriminate between the alternatives. These in turn will provide further evidence about why representations are useful and therefore help to support or undermine theoretical explanations in terms of expressiveness and computational complexity.

It is quite hard to find domains in which truly equivalent expressions of the same information in different modalities can be studied, at least if one wants to maintain any plausibility that both of the modalities are useful. The area in which most work has probably been done is what I will call here *matrix* graphics—histograms, pie charts, line-graphs—graphics which represent a population of data points on a set of dimensions of variation(Stenning, Neilson, Inder 1993). As we shall see, the high degree of information enforcement in these representations means that they are quite limited in the types of data-set they can represent, but for those data, for most inferential tasks, they are very clearly superior to linguistic modality representations. Where there is cross-over it is usually for rather obvious task determined reasons.

We gain more methodological elbow-room if we can find a domain and task which require slightly more expressive resources where opinion is somewhat more evenly balanced about what representation is best.

3 The Categorial Syllogism

There is a long history of graphical treatments of the syllogism used for teaching elementary logic. Euler invented his system in about 1770. The very longevity of this system suggests its usefulness, yet there are also accurately equivalent sentential methods of teaching syllogisms (see e.g. Boolos & Jeffrey 1989). Furthermore, there are graphical variants with different semantics and different expressive power (e.g. Venn (1894), Dodgson (1896)). The domain is sufficiently simple that it is possible to adapt standard systems of representation such as matrix graphics to represent the required information, or even to invent novel graphical systems with different combinations of properties to investigate theoretical predictions.

In keeping with our focus on expressiveness, we will progress from the least expressive treatment to more expressive ones. We start by considering a matrix graphics system which is MARS. We then consider how it might be made more usable by extending it to be a LARS. We progress through Euler's LARS system and circle-based variants, to a network based system which is potentially UARS. The three systems will each be described in turn before comparisons are made in Section 4.

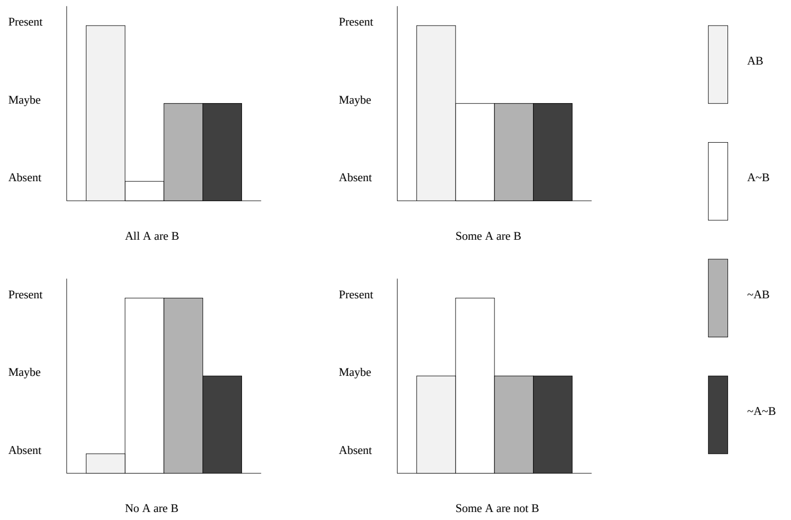

Figure 1: Representing syllogistic premisses with histograms

3.1 Matrix graphics: a treatment in depth

Any treatment of the syllogism requires representations for the possible premisses and procedures for relating the information from the two different premisses in a way that allows either the drawing of a valid conclusion or the decision that there is none. Representing the premisses requires that they be related to the fundamental semantic entities which are collections of types of individual specified in terms of combinations of properties. The premisses contain quantifiers which express relations between two sets—e.g. *All A are B*. Since there are four types of individual defined by two sets (ABs, $A\neg Bs$, $\neg ABs$ and $\neg A\neg Bs$), each quantifier is a relation between these four types of individual and three possible states (existence entailed by the premiss; non-existence entailed by the premiss; existence/non-existence not determined by the premiss).

These relations could be represented in a variety of ways using matrix graphics. For example, they could be represented as histograms as illustrated in Figure 1. Procedures for combining information from histograms representing a pair of premisses could be defined, but it is difficult to see how procedures for drawing inferences could be any more systematic than a look-up table.

Another way to represent these relations which is rather more promising is to use two linear dimensions for the sets A and B. The four types

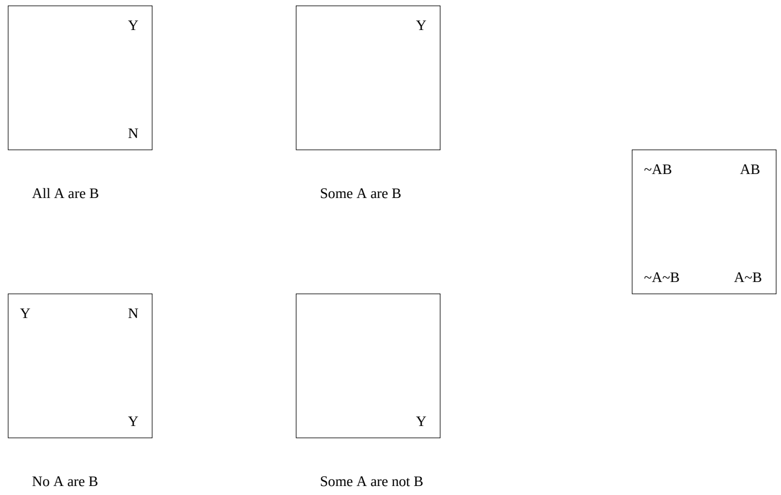

Figure 2: Representing syllogistic premisses with squares

of individual defined by the two predicates correspond to the corners of a square. We use labels Y and N to represent 'must exist' and 'must not exist' respectively, and leave unlabelled those combinations not determined by the premiss. This representation is similar to a table, but it also allows us to represent the sets A and B and their complements, as well as the four intersections, by using the edges of the square[1]. Figure 2 shows how each of the four quantifiers is represented in this system.

One way of exploiting spatial representation systematically is to use the third dimension, and we can use this to combine the two premises of a syllogism. Whereas each premiss is represented as a square, their combination can be represented as a cube. The eight types of individual defined by the three predicates correspond to the vertices of the cube. The edges represent pairwise intersections, and the faces the individual sets and their complements. Figure 3 illustrates this. Each premiss square is a projection of the cube onto a plane – the edges of the cube are projected onto the corners of the square.

A fairly straightforward algorithm allows us to use the cube to draw syllogistic inferences from the two premisses. Labels are propagated from

[1]This system bears some resemblance to that of Lewis Carroll (Dodgson 1896) though his system does not appear to have exploited the possibility of expressing abstraction by use of a third dimension. His trick of placing notations on borders owes more to Venn's system than to the one described here.

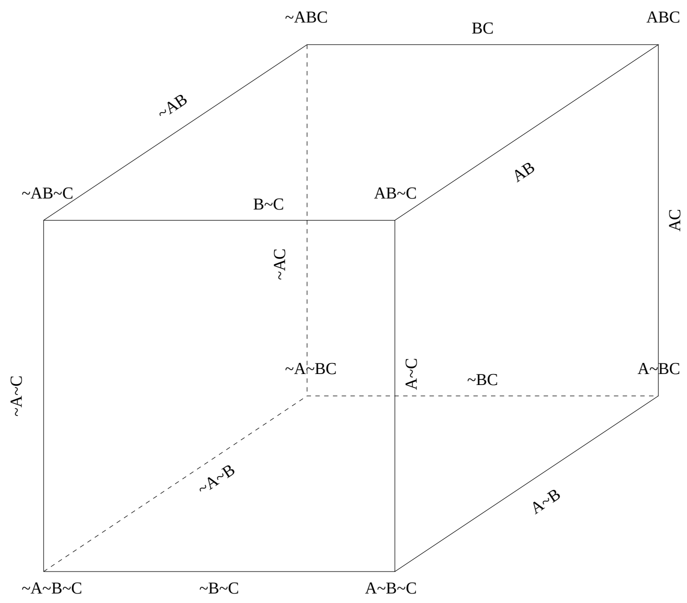

Figure 3: Representing syllogisms with a cube

the edges to the adjacent vertices, and from the vertices to the remaining edges. The cube can then be projected onto a third square which expresses the relation between A and C, and the conclusion (if any) can be read off.

The rules for propagating labels are simple and fairly well motivated. Any edge labelled N requires both its end-vertices to be labelled N – if there are no individuals of the type specified by two predicates, then there are no individuals of the type specified by those two predicates *and the third predicate*. Any edge labelled Y requires that at least one of its end-vertices be labelled Y – if there are individuals of a type specified by two predicates, then there must be individuals of one or other of the two types formed by adding the third property or its negation to the type-description. Conversely, if either end of an edge is labelled Y, the edge must be labelled Y, and if both ends of an edge are labelled N, the edge must be labelled N. The algorithm is illustrated in Figure 4 for the syllogism *All A are B, Some C are not B.*

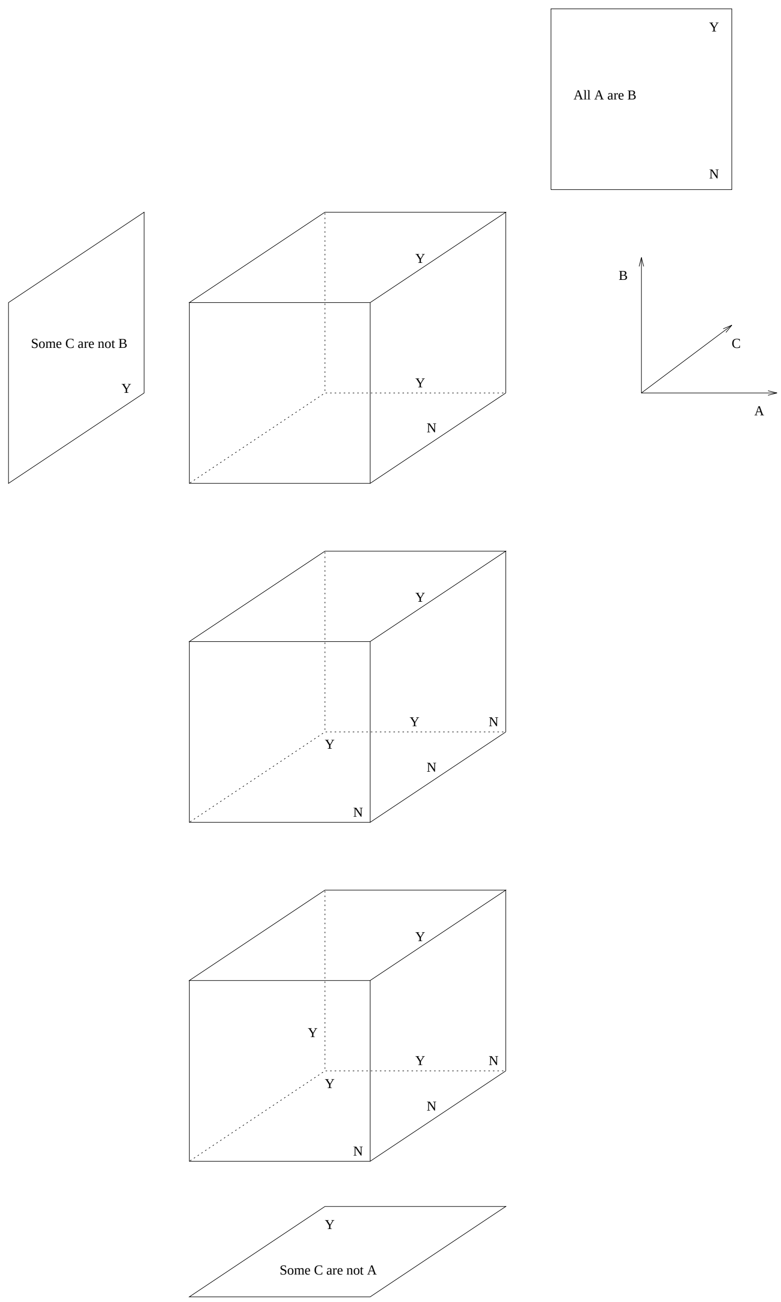

Figure 4: Solving a syllogism with the cube

3.2 Set Diagrams: Euler's Circles

There are several diagrammatic systems for logical reasoning which rely on the analogy between spatial inclusion and set inclusion. Two of the best known are the eighteenth century system known as Euler's circles (Euler 1772) which was later modified by Venn 1894. These systems were originally specifically targeted at syllogism solution. Although they are based on the same analogy, their respective semantics are quite distinct (see Stenning (in press)), and though they have never been systematically investigated for their comparative cognitive properties, there is quite a large difference in their expressiveness which theory would predict would lead to didactic differences.

We will concentrate here on Euler's system because it uses to a greater degree the geometrical constraints of circles to model the logical constraints of the syllogism. There are misinterpretations of Euler's method in the literature but we have given a detailed account of the semantics of the system in Stenning & Oberlander (1991). Here we repeat a brief summary for the reader's convenience.

A specification of the algorithm for solving syllogisms with Euler's circles requires strategies for:

- choosing premiss representations
- unifying premiss representations
- deciding whether there are valid conclusions
- formulating valid conclusions

An example will illustrate the method.

Figure 5 illustrates the method applied to the syllogism *All A are B. Some C are not B.* Each premiss is first represented by the diagram of its *maximal* model ie. the one containing the most types of individual (henceforward *types*) consistent with the premiss. Within these diagrams, sub-regions representing the *minimal* models of the premiss are marked with an 'x'[2]. The two premiss diagrams then have to be 'registered'—their B circles superimposed and their A and C circles arranged, again to form the maximum number of regions consistent with both premisses. A simple topological rule

[2]The syllogism is generally interpreted with a 'No-empty-sets' axiom which asserts that there are As, Bs and Cs.

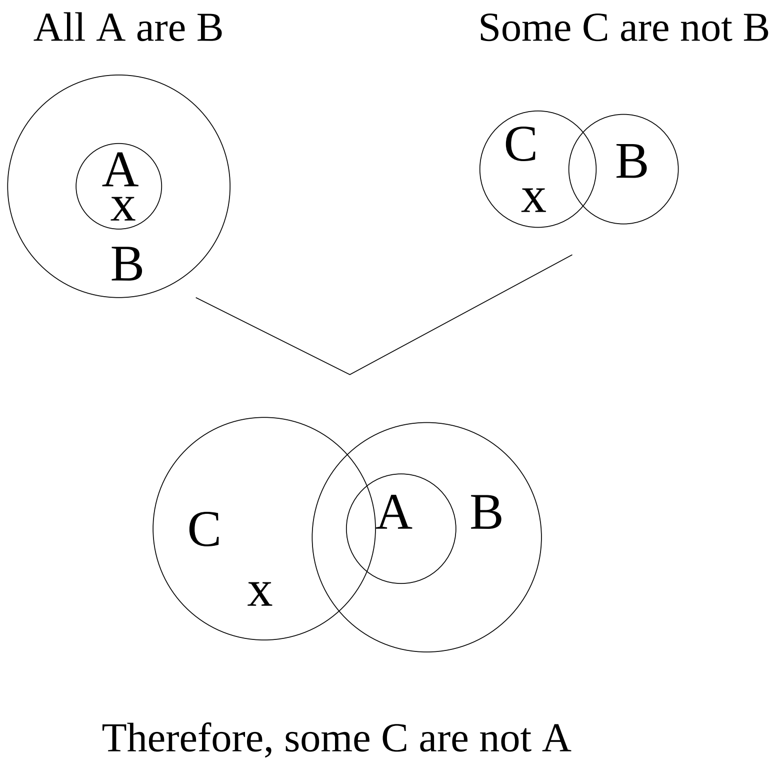

Figure 5: An example of the graphical algorithm applied to the syllogism *All A are B. Some C are not B.*

indicates the fate of the regions marked with xs in the premiss diagrams. If their region is bisected by the third circle when the premiss diagrams are registered, then the x is removed. This ensures that the remaining xs still mark the minimal models of the conjunction of the two premisses.

Having constructed a representation of the premiss information, it is now possible to determine whether there are any valid conclusions and to formulate them. If and only if there is a region marked with an x is there a valid conclusion. We will call such regions *critical* regions. Existential conclusions are formulated by taking the 'type descriptor' for a critical region, removing the middle B term, and existentially quantifying. So, if the critical region represents As that are B but not C, then conclude *Some A are not C*. Universal conclusions are only valid if the critical region is circular, in which case the label on the critical region's circle becomes the subject of a universal conclusion. If the critical region is wholly included in the other end term circle, the valid universal conclusion is positive. Otherwise it is negative. This graphical algorithm is summarised in Figure 6.

This completes our description of this reconstruction of Euler's system. A few remarks about the relation between this system and that of Venn may be in order here as they provide a link between matrix graphics and Euler's system.

Venn's system uses only one diagram, the seven-region one which represents all possible types. The system augments this with annotations which distinguish the status of the type of individual corresponding to the annotated region. Slightly different systems of annotation coexist. For an insightful exposition with a full formal semantics see Shin (1991). Because the system allows negative annotations which are interpreted to mean that no individuals of a type corresponding to the annotated region exist, the diagrams do not exploit any geometrical constraint in defining combinations of types. In fact, as Figure 7 shows, these diagrams are merely unconventionally organised tables. In contrast, Euler's system uses geometrical constraints to express combinations of (possible) types of individual. Whereas there is a Venn diagram (including suitable annotation) for any syllogistic model, only about half such models have a unique Euler diagram. See Stenning (in press) for a more extended discussion of the relation between Euler and Venn systems. We return to these points when comparing the system below.

1. Form characteristic diagram for each premiss;
2. Register B circles of the characteristic diagrams of the premisses and arrange A and C circles with most types consistent with the premisses.
3. If no shaded region from a component premiss remains non-intersected, then exit with No Valid Conclusion response. If there is one, then it is the *critical* region.
4. If such a region does exist but both premisses are negative, then exit with a No Conventional Valid Conclusion response. (If task permits, conclude that *Some non-As are not Cs*).
5. Formulate conclusion:
   (a) Take the description of the individual type represented by the critical region of the diagram (e.g. $A \neg BC$)
   (b) Eliminate the B term from this description
   (c) Existentially quantify the remaining description for an existential response
   (d) Is the critical region circular?
      i. If so, it is the subject term of a universal conclusion
      ii. If not, there is no universal conclusion

Figure 6: Summary of a graphical algorithm for solving syllogisms using Euler's Circles

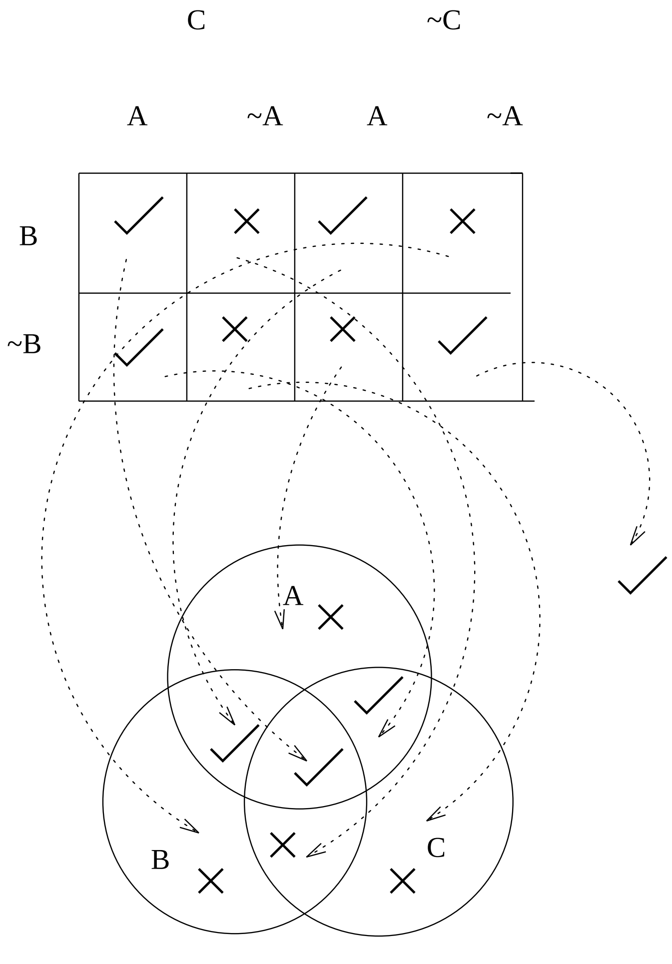

Figure 7: Venn Diagram as a variant of a table: dotted arrows indicate mapping from table cells to diagram regions

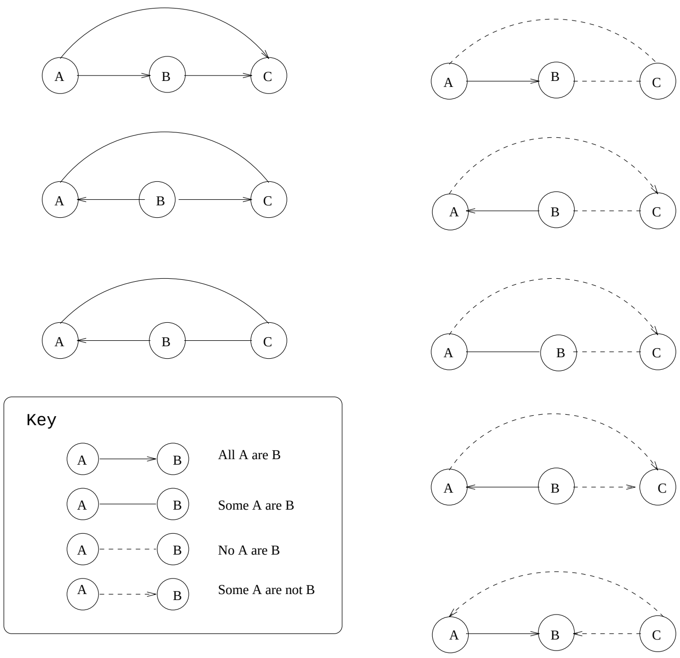

Figure 8: A network method of syllogistic reasoning

3.3 A network treatment

Figure 8 illustrates a graphical method for solving syllogisms using a network notation. Each premiss type is represented by two predicate nodes linked by one of four link-types corresponding to the four possible quantifiers. A pair of syllogistic premisses is therefore represented by three nodes joined by two links, and reasoning to a conclusion is a matter of deciding whether the two as yet unlinked nodes can be joined by a new link, and of what type such a link can validly be.

Figure 8 represents all the syllogisms which have conventionally valid conclusions[3] and the strongest valid conclusion that can be drawn. Although we have labelled the outer nodes of each case 'A' and 'C' for convenience of reading, the other half of the valid syllogisms can be read off by switching the labels.

Figure 9 presents a method of reasoning from representations of pairs of premisses to conclusions. First some terminology: we call both *arrows* and

[3] *All not A are not C* conclusions are not included though they could be if a fifth link-type were permitted for expressing conclusions

symmetrical links links; *rerooting* adds a link to the diagram connecting the outer nodes by taking the root of a link from its old node and rooting it in the third node. Extending arrows adds a link to the diagram by taking the head of an arrow from the node it points to, to the third node.

4 Comparing the three graphical treatments

What does our general theory have to say about these three systems of syllogistic reasoning? Unfortunately little systematic experimental investigation of the usability of even the oldest of them is available in the psychological literature and so we have to rely, until such evidence becomes available, on intuition. Intuition is an unreliable guide, but at least intuition about concrete systems of representation and inference worked out in detail is considerably more reliable than intuition about vague generalities about what a good thing graphics is. At least these examples may show where theory has to be further developed before predictions can be made.

To begin with the last described, the network method make its closeness to linguistic formulations transparent. The semantics of the representations does not exploit the logical properties of the spatial relations in anyway other than to define a graphical 'concatenation' operator. Inference rules for any other set of quantifiers could be defined equally well on the same representations. Of course, it is very hard to be sure that there is not a simpler, more perspicuous set of rules definable on these basic representations, but we have certainly not gone out of our way to bias matters against networks. These are the simplest rules we have been able to define.

One observation about our network system is that it makes the symmetry/asymmetry of the quantifiers more explicit than does the natural language formulation. Symmetrical links represent symmetrical quantifiers: asymmetrical arrows represent asymmetrical quantifiers. Linguistic forms are inherently asymmetrical since temporal or spatial ordering is inherent in the modality. So this asymmetrical nature has to be 'overcome' in learning that some quantifiers are logically symmetrical. The network system can represent quantifiers by primitively symmetrical or asymmetrical signs thus highlighting a logical distinction. In fact, the network system *must* represent logically asymmetrical relations asymmetrically because otherwise the viewer has no way of telling which way to assign the relation (where in language that could never be left undetermined).

Symmetry or asymmetry are important features which one learns ab-

Rules for rerooting solid arrows:

1. Reroot a solid arrow back along any link except against the arrow of a negative link
   (a) if the extending link is dashed the established connection is dashed
   (b) if you're extending back going against an arrow, then the new link remains an arrow, otherwise it loses its arrowhead
   (c) if you're extending back along a dotted link, then the new link remains an arrow, pointing towards the dotted link

A rule for extending solid arrows forward:

1. Extend a solid arrow forwards along a link only if the link is dotted and it isn't in the direction of the extending arrow. The new link is of the same type as the dotted link.

Finally, the only rule for extension not involving a solid arrow:

1. If both links are symmetrical, one dotted and one not, then draw a dotted arrow pointing to the dotted side

If none of these rules apply, there is no valid conclusion.
Conclusion read-off is a matter of encoding the relation between the two outer nodes according to the key.

Figure 9: Summary of rules for using network method of syllogism solution

out quantifiers in learning syllogistic logic. We have independent evidence that this may be a real problem for at least one sub-group of undergraduates when faced with the task of learning this domain (see Stenning & Cox forthcoming). But apart from this possible gain of emphasis and some saving in numbers of rules, it is hard to see what the network system gains over any linguistically based set of rules for solving syllogisms. With this one minor caveat, we would predict no benefit of the system because it does not exploit graphical constraints in expressing logical constraints. Both the cube and circle treatments exploit geometrical constraints to represent logical constraints, and both seem to be the more usable for it.

The more interesting comparisons are between the set diagram systems and our 'cubist' treatment. Our general intuition is that the cube system is less usable than the Euler system. Part of this intuition may be attributable to the use of two dimensional diagrams to represent three dimensional representations—a wire-frame model might be easier to operate, but this is only part of the story. Euler's Circles seems to us the most thoroughgoing exploitation of geometrical constraints of all the three systems and to be the most usable system both for teaching and for actually performing syllogistic reasoning. The circles are appreciably simpler than the cubes—one has only to try using the systems without external aid to appreciate this fact. There are several reasons. Euler's system actually exploits the directionality of inference by jettisoning information in constructing a representation of a premiss pair. It is not always possible to get back to the original premisses from a composite diagram. The cube system preserves all premiss information. Euler's system is forced to take this parsimonious approach because it has no resources for representing individuals not specified by all the represented properties. Once three circles are present, every represented individual is maximally specified in terms of all three predicates. In contrast, the cube represents partially specified individuals as faces and edges and so can maintain premiss information.

To assess the expressiveness of the cube system, and the extent to which it exploits spatial constraints, it is necessary to place it within a family of possible representational systems constrained by the graphical nature of the central device. We must consider two aspects of each system: the number of possible representations of quantifiers, and the nature of the transformation operations.

With respect to quantifiers expressible, there are 3^4 possible representations of quantifiers in the cube system (four corners each containing one of three annotations, 'y', 'n' or blank). With ECs, there are five diagrams depicting topological relations between two sets. Each minimal region of each diagram (including its background) can be marked with an X or not. This allows 44 quantifiers to be expressed, about half the number for the cube.

Both these numbers are substantially reduced by semantic constraints which the user may know (or be told). For example, none of the distinctions between models containing or not containing the 'background type' are ever used. In ECs this corresponds to never placing Xs in the background. In cubes, it corresponds to never annotating the 'origin'. Another example, we assume the no-empty-sets axiom, and this reduces the number of quantifiers expressible in both systems, but still leaves the cubes more expressive than the circles.

But it is with respect to the transformations that the circles are really seen to be less expressive and more geometrically constrained. The propagation rules in the cube system are unconstrained symbol rewriting rules. Even though the particular rules used can be related to model-theoretic properties of the system, these rules are of a *form* which is quite arbitrary and unrelated to the spatial structure of the diagrams. The spatial structure of the diagram is only exploited to organise the application of the rewriting rules by defining which symbols are antecedents to which applications. This situation is reminiscent of our brief discussion of Venn diagrams as unconventionally laid out tables. In ECs, the transformation operations are those involved in 'registering' pairs of circles. These operations are topological operations which are constrained entirely by spatial considerations save for the rule about removing Xs from bisected regions. Even this elimination rule is of a simpler form that unconstrained rewriting rules. On both counts the circles are more constrained than the cubes, though our intuition is that it is the constraint on rules of transformation which is important.

Analysis of the semantics of the cube system raises the interesting question to what degree it can be regarded as matrix-graphical. The 2-D representations of quantifiers are effectively tabular[4] and therefore matrix graphical. However, when we consider the 3-D innovation of using edges to represent combinations of two properties and faces to represent combinations of three properties, we go beyond matrix graphics. Representation of partially specified items is not possible in matrix graphics. Note that these 'tricks' do not allow general representation of abstraction, at least not unless arbitrary faces or edges are assumed to be available for labelling inside

[4]Note the correspondence to Lewis Carroll's tabular system mentioned above.

the cube. These small extensions of expressive power of a fundamentally matrix graphical system are typical of the tricks which make real systems usable for many domains/tasks. They are akin to the 'empty cell' notations discussed in Stenning & Oberlander 1991, and the 'x' chains of Shin's (1991) extended Venn system. They are what make the cube system usable at all, but according to our intuition do not bring it to the level of tractability of the ECs.

Of course, systematic observation may prove us wrong in this intuition. Perhaps some particular task can be found for which the cube system is better fitted than the circle one. In constructing it we discovered several curious inferences which we had not previously noticed because the Euler system does not deduce them[5]. If the additional abstractions and inferences that the cube allows were vital to some task, then the cube system could find a use. Similarly, it will be interesting to see whether the network representation has some use for teaching the importance of symmetrical and asymmetrical quantifiers.

Our main contention is that something like this level of specification of alternative systems for reasoning in the same domain is required if a serious theory of usability is to be based on a theory of representations.

References

Berkeley, G. (1937) *The Principles of Human Knowledge: the text of the first edition (1710)*. London: A. Brown.

Boolos, G. S. and R. C. Jeffrey (1989) *Computability and Logic*, 3rd ed. Cambridge: Cambridge University Press.

Cox, R., K. Stenning and J. Oberlander (1994) Graphical effects in learning logic: reasoning, representation and individual differences. In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*. To appear.

Dodgson, C. (1896) Symbolic Logic. Chap. 3 in *Lewis Carroll's Symbolic Logic*. Hassocks, Sussex: Harvester Press.

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

Lear, J. (1980) *Aristotle and Logical Theory*. Cambridge: Cambridge University Press.

[5]For example, from *All A are B* and *No B are C* it follows that *Some not A are not* **B**

Levesque, H. (1988) Logic and the Complexity of Reasoning. *Journal of Philosophical Logic* **17**, 355–389.

Shin, S-J. (1991) A Situation-Theoretic Account of Valid Reasoning with Venn Diagrams. In J. Barwise, J. M. Gawron, G. Plotkin and S. Tutiya, eds., *Situation Theory and its Applications*, Vol. 2, pp. 581–605. Stanford, Ca.: Center for the Study of Language and Information.

Stenning, K. (in press) Logic as a foundation for a cognitive theory of modality allocation. In M. Masuch, ed., *International Colloquium 'Logic at work'*, Amsterdam University.

Stenning, K., I. Neilson and R. Inder (1993) Applying semantic concepts to the media-assignment problem in multi-media communication. Research Paper. Human Communication Research Centre, University of Edinburgh.

Stenning, K. and J. Oberlander (1991a) A cognitive theory of graphical and linguistic reasoning: logic and implementation. Research Paper HCRC/RP-20, Human Communication Research Centre, University of Edinburgh.

Stenning, K. and J. Oberlander (1991b) Reasoning with words, pictures and calculi: computation versus justification. In J. Barwise, J. M. Gawron, G. Plotkin and S. Tutiya, eds., *Situation Theory and its Applications*, Vol. 2, pp. 607–621. Stanford, Ca.: Center for the Study of Language and Information.

Stenning, K. and J. Oberlander (1994) Spatial containment and set membership: a case study of analogy at work. In J. Barnden and K. Holyoak, eds., *Analogical Connections*, pp. 446–486. Hillsdale, N.J.: Lawrence Erlbaum Associates.