LOGIC AS A FOUNDATION FOR A COGNITIVE THEORY OF MODALITY ASSIGNMENT

KEITH STENNING

May 30, 1994

Abstract

A cognitive theory of modality allocation has to explain the different cognitive consequences of assigning the same information to different modalities of expression—crudely, which picture is worth which thousand words? Most theories of the benefits of graphical representations have focussed on *visual* properties of graphics. This talk will sketch a theory which focusses on general logical/computational properties of graphics. The theory is based on the proposition that graphical representations are weakly expressive, in a logical sense, and that this explains graphics' processing advantages. The theory distinctively predicts that graphics will be worse than language when graphics' logical inexpressiveness impedes abstraction that is necessary for a task.

The general theory will be illustrated by an analysis of the use of Euler's Circles to solve syllogisms. Earlier misinterpretations of Euler's system by psychologists illuminate the distinctions the present theory draws between minimally abstractive representation systems (MARS), limited abstraction representation systems (LARS), and unlimited abstraction representation systems (UARS). The analysis reveals model-theoretic properties of this logical fragment which explain why it can be captured by Euler's system. Euler's and Venn's diagrammatic systems are both based on the analogy of spatial containment for set membership. The basis of the difference in their expressive power is revealed by this approach. Euler's graphical algorithm is then generalised and the cognitively fruitful geometrical limitation to circles is related to the logical limitations of the generalised system.

The example analysis of Euler's Circles is then related back to the cognitive theory of graphics and the role that logic plays in this approach to modality allocation.

1 Perspective

Psychologists have tended to employ an anachronistically mechanistic view of logic in their theorising about human reasoning (see e.g. Johnson-Laird 1986, Rips 1986). Instead of regarding a logic as an abstract consequence relation which can be implemented in many mechanisms, they have tended to assume that a direct implementation of their favourite pencil-and-paper method of doing logic (often a natural deduction system) is the only mechanism of reasoning related to logic. Alternatively, some psychologists have claimed that psychological theories of reasoning correspond to model theory in logic, thus neglecting the fact that model theoretic reasoning generally takes place in some metalanguage.

Elsewhere in this volume (Gabbay, Chapter ??), proposes moving more implementational detail inside what we count as a specification of a logical system. Since this makes explicit the many alternative methods of implementation and focusses attention on their differences, it can be viewed as complimentary to the current approach. That is, it makes explicit what are implementation details in the classical view by importing them within the specification of logical systems.

Alongside the assumption that logic is a mechanism, a further assumption widespread in the psychological literature is that logic is a *sentential* mechanism. It is easy to see where this assumption comes from—since the nineteenth century, sentential methods of specifying logics have completely dominated the field. But this obscures the fact that there was a strong vein of graphical thinking in the development of logic and particularly in the teaching of logic. Recently this graphical perspective on reasoning has gained interest both from a didactic point of view (e.g. Barwise and Etchemendy 1990)

and from a more purely logical point of view among the theorem proving community (e.g. Barker-Plummer & Bailin 1992).

A modern, more abstract view of logic which makes implementation explicit has a lot to offer to cognitive theorising about human implementations of logical fragments. The view leads to a focus on substantive empirical questions about how logics are implemented in human memory rather than on trying to model issues of performance directly through competance theories of reasoning. This paper argues that there is a substantive role for logical techniques in analysing non-sentential systems of representation and predicting their cognitive properties.

The remainder of this section will sketch an approach to a general cognitive theory of modality allocation which sees graphical systems as weakly expressive systems. Graphics are seen as cognitively useful when their weak expressive power is sufficient to the task at hand and offers the inferential tractability which comes with weak expressiveness. Graphics are seen as pathological when the task at hand requires abstractions which cannot be graphically expressed, except perhaps by unweildy disjunctions of graphical representations.

This approach starts by describing three grades of expressiveness and claims that actual graphical systems are generally intermediate and that their cognitive properties can be explained on that basis. The next section, takes Euler's Circles, an eighteenth century system of teaching elementary logic, as an example and by providing an analysis of the system's use of graphics shows how the approach works in one simple case. Even this elementary example reveals some intriguing logical properties of the syllogism, particular a property called *case identifiability* which is the property which permits Euler's particular graphical treatment of the logical fragment. This model theoretic property also reveals a relation between syllogisms and the more natural task of text comprehension which perhaps explains the historical prominence of this logical fragment.

The third section explores generalisations of the graphical fragment and their relation to sentential fragments. Geometrical properties are related to logical limitations. Some remarks about possible human memory implementations of these graphical fragments are intended to suggest the cognitive relevance of these logical observations. Finally, the analysis of Euler's Circles is related to the general framework for studying non-sentential expressions of information, and the role of logic in developing this approach.

1.1 A general theory of modality allocation

Information can be allocated to expression in different media and modalities. Choice of expression can have substantial impact on the cognitive processing of the same information. There is a sort of folk-lore, often expressed in the adage that 'a picture is worth a thousand words', that graphics are a 'good thing'. There are substantial literatures on graphical and 'direct manipulation' interfaces to computers (Salvendy 1987), on visual programming languages (Myers 1990),

on the design of graphics (Tufte 1983, Cleveland 1985, Twyman 1979), on media allocation (Larkin & Simon 1987, Hovy & Arens 1990) and text layout (Hartley 1987, 1992).

All of these are concerned with improving the consumption of information. However, there is very little theory which underlies choice of design and spans a wide domain of information expressions. It seems highly questionable whether graphics are always a 'good thing', and certainly *why* they are good when they are. A theory of which picture is worth which thousand words would be of immense practical value, quite apart from its theoretical interest.

Both *medium* and *modality* are concepts which deserve close scrutiny (see e.g. Hovy & Arens 1990, Coutaz 1992, Stenning, Neilson & Inder (1993)), but for present purposes, we will

concern ourselves here with contrasts between sentential and diagrammatic representations, regarding these as contrasting modalities.

The key idea in this approach to understanding graphical representations is that they *enforce* the determination of certain classes of information. Or, to put it another way, graphics are concrete in the sense that they cannot express certain abstractions, and it is this inexpressiveness that makes them inferentially tractable.

We distinguish *primitively* interpreted graphics which use a single diagram to represent one model. But there are, of course, conventions for introducing some abstractions. A substantive claim is that *not all* abstractions can be reintroduced by conventions of interpretation. We therefore distinguish three grades of expressiveness of representation systems: MARS (minimally abstractive representation system) represent exactly one model for each representation in the system, under the intended interpretation. LARS(limited abstraction representation system) representations can abstract over several models but cannot express arbitrary partitionings of the space of models. UARS(unlimited abstraction representation system) representations can express any arbitrary partitioning of the space of models. Stenning & Oberlander 1992 presents a more detailed exposition of these distinctions.

An example will clarify these intuitions. Consider a two-by-two tabular system of representations for a domain of cell contents [X, 0] where every cell contains one and only one symbol. Then a simple tabular representation is a MARS.

X	0
X	0

It represents just one of the 16 possible models from the domain. A simple convention for representing some abstractions is introduction of a meta-symbol '?' interpreted to indicate a disjunction of values in a cell.

X	?
X	0

This convention turns the system into a LARS. One 'diagram' abstracts over more than one model, but the system still cannot express any one of the 2^{16} partitionings of the space of 16 models. For example, to capture the partitioning which includes just

X	0
X	0

and

X	X
0	0

requires something more than this meta-symbol convention because it requires contingencies between cells' contents to be expressed. It is not easy to see how to introduce a natural convention which will permit expression of this abstraction[1], though allowing arbitrary indexing of 'key assertions' to cells (see Brachman & Levesque 1985) will allow any

abstraction to be captured, but leaves the graphic doing no work. (See Stenning & Oberlander 1992 for a more detailed development of this framework of grades of expressiveness).

The observation of the concreteness of graphics on which this general approach to understanding their cognitive efficacy is not original. It has a long history going back at least to Berkeley's (1710)

discussion of Locke's theory of word meaning. Berkeley also noted that it was nevertheless possible to base some general reasoning on graphical illustration, and he discussed this in connecion with geometry proofs and the restrictions on appeal to particular properties of figures in the derivations of general proofs. One reason these observations have never been built into a general theory of graphical representations is that Berkeley's two observations appear to cancel each other out. Making an empirical theory out of the observations hinges on showing that not all abstractions can be expressed by conventions of interpretation for graphics. The theory can then acknowledge that real graphical systems are not interpreted completely concretely, but that they are are also limited in their expressiveness, and hence amount to systems of intermediate expressive power. It then remains to demonstrate that this intermediate band of expressiveness gives cognitive advantage.

We conjecture that the general computational theory of the complexity of inference can be applied to explaining cognitive complexity. A complete theory would of course require specifying details of the particularities of human computational resources. Nevertheless, an architecture free approach through general computational complexity can perhaps go some way toward explaining some of the broader categories of task complexity as it is affected by the nature of external representations. In order to explore this possibility it is necessary to analyse *real* examples of graphical systems that have proved their usefulness in performing tasks and to see whether their interpretations accord with the general approach proposed. We take a classical graphical system for teaching elementary logic as a simple example in which to develop a detailed application of general theory.

2 Euler's Circles—really useful graphics

We choose the domain of elementary logic diagrams based on the analogy of spatial inclusion for set membership firstly because the several systems have stood the test of time, and so we can have some confidence in their usefulness; and secondly because the task of syllogistic reasoning is well defined and well understood in a sentential framework. We choose Euler's

[1] An indifference-to-rotation convention captures these two but also allows in two more models

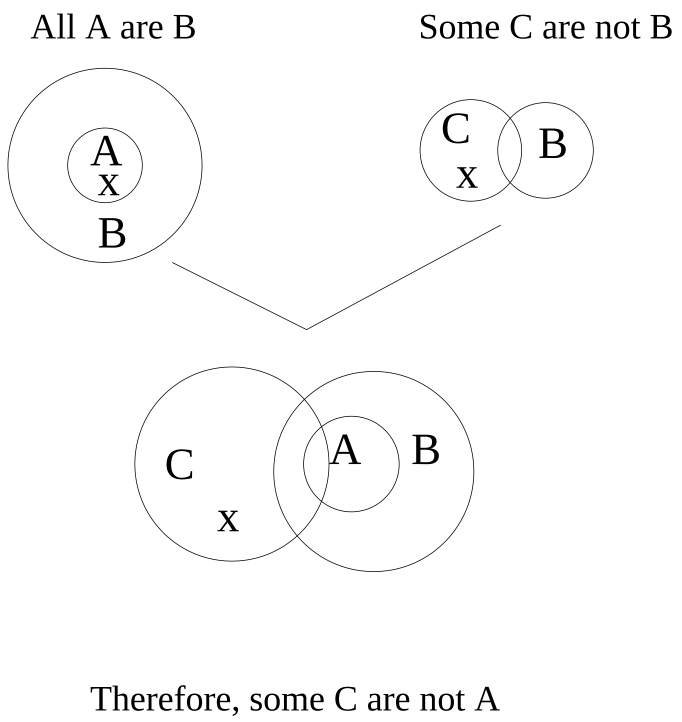

Figure 1: An example of the graphical algorithm applied to the syllogism *All A are B. Some C are not B.*

particular system because it is even weaker in its expressive power than, for example, Venn's better known system, and because it is, in a sense which will become plain, more *graphical*. It exploits the concreteness of graphical representations more thoroughly than Venn diagrams.

Our main concern here is the relation between graphical and sentential formulations of fragments of logic. For a more detailed treatment of the relation between this graphical system and psychological data relevant to the internal representations people use in solving syllogisms without recourse to external diagrams see Stenning & Oberlander 1993.

We first present a rational reconstruction of Euler's method. This is not intended as a historical exegesis, but we believe does not depart from the spirit in which Euler intended his system to be used. Some subsequent interpretations of his systems by psychologists (e.g. Erickson 1975, Johnson-Laird 1983) interpret the system differently.

The differences turn out to bear directly on our distinction between MARS and LARS.

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 1 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 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 2.

2.1 Case Identifiability and Graphicality

This graphical algorithm for solving syllogisms brings out some interesting model-theoretic properties of the syllogism. The fragment is *case-identifiable*—if there is a valid conclusion, then there is some maximal type of individual (case) whose existence is entailed by the premisses. The converse almost holds but for a group of syllogisms for which it is valid to conclude *Some not As are not Cs*. We will call these *U-conclusions* after the A, E, I and O conclusions of the medieval notation. An example appears in Figure 3. The existence of these inexpressible conclusions is of some historical interest. Aristotle presumably would not have accepted that these conclusions related the sets A and C because he did not distinguish between the domain of interpretation and the universal domain—an achievement which awaited twentieth century metalogic. These conclusions, were they admitted, would violate Aristotle's principle that two negative premisses never licence a valid conclusion. Interestingly, these conclusions are easier to find than some of the conventionally valid ones (Yule & Stenning 1992).

The property of case-identifiability is critical to Euler's invention of his graphical method. Syllogistic conclusions require the entailment of maximal types. These graphical representations can only represent maximal types because any point in a plane is either inside or outside any circle in the plane. Syllogistic conclusions are not contingent on relations between types in models, and these contingencies, in the form of disjunctions, are just what diagrams cannot freely represent. It is possible to define Euler's algorithm involving a single progressive construction for every syllogism because it is possible to simultaneously portray the maximal and minimal models. This, in turn, is possible because both are always unique. In fact,

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

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$\neg B$C)
    (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 2: Summary of a graphical algorithm for solving syllogisms using Euler's Circles

No A are B. No C are B.

Therefore Some not-A are not-C.

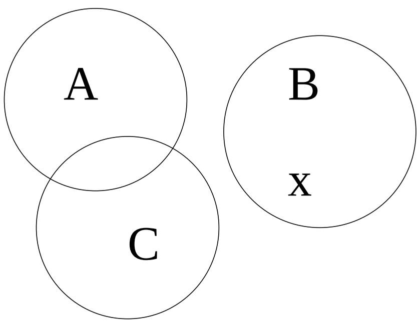

Figure 3: Example of a syllogism with a U-conclusion

reasoning could proceed by simply representing minimal models. Maximal models are only of importance in deriving the minimal models for conjunctions of premisses.

2.2 Euler's Circles as LARS

Returning to our general categorisation of representation systems we can see that Euler's system is a LARS—each diagram abstracts over several possible models—and that this depends on:

- the interpretation of regions as representing *possible* types;
- the strategy of selecting the maximal model to represent each premiss;
- the device of simultaneously distinguishing minimal models;
- the strategy of registration to form *maximal models*; and
- the topology of circles.

This LARS interpretation of Euler's diagrams contrasts with a MARS interpretation which has been current in the psychological literature on syllogism solution. This interpretation treats sub-regions of the diagrams as representing types existing in the model, and therefore represents one model per diagram. Figure 4 shows the resulting many-to-many mapping from diagrams onto premisses. Using this interpretation to find that a syllogism such as *Some A are B. Some B are C.* has no valid conclusions requires examining the registration of 16 pairs of premiss diagrams, a feat well beyond the working memory capacity of human subjects. This combinatorial explosion lead Erickson (1974) into a stochastic account of reasoning, and Johnson-Laird (eg. 1983) into a rejection of graphical methods. At least for this example, it seems that our distinction between MARS an LARS interpretations is critical to the usability of graphical systems of reasoning.

Conversely, we can see in this example that the invention of a graphical system is dependent on the weakness of the fragment of logic to be implemented. There are many syllogistic models which cannot be uniquely represented by an Euler diagram, but the inaccessible distinctions are ones which the syllogistic fragment does not require. This line of thought raises questions about just what graphical methods can achieve in this domain.

3 Euler's Circles Generalised

Graphical methods based on the analogy of spatial inclusion for set membership are clearly capable of capturing more than the syllogism. We have already seen that the method reveals informative inferences from syllogistic premisses which Aristotle gives us no way of expressing sententially. What fragment of logic can be implemented by our graphical system and its natural extensions?

To impose some structure on the space of graphical systems based on the analogy between spatial containment and set membership, it is useful to contrast Euler's method with Venn's. Euler represents all and only *possible* types by regions in diagrams and distinguishes *necessary* types by marking some sub-regions. So *impossible* types are eliminated *topologically*. Any type inconsistent with the premisses is simply not represented in the diagram. Venn, in contrast, employs a single diagram, namely the one with a minimal region for every type.

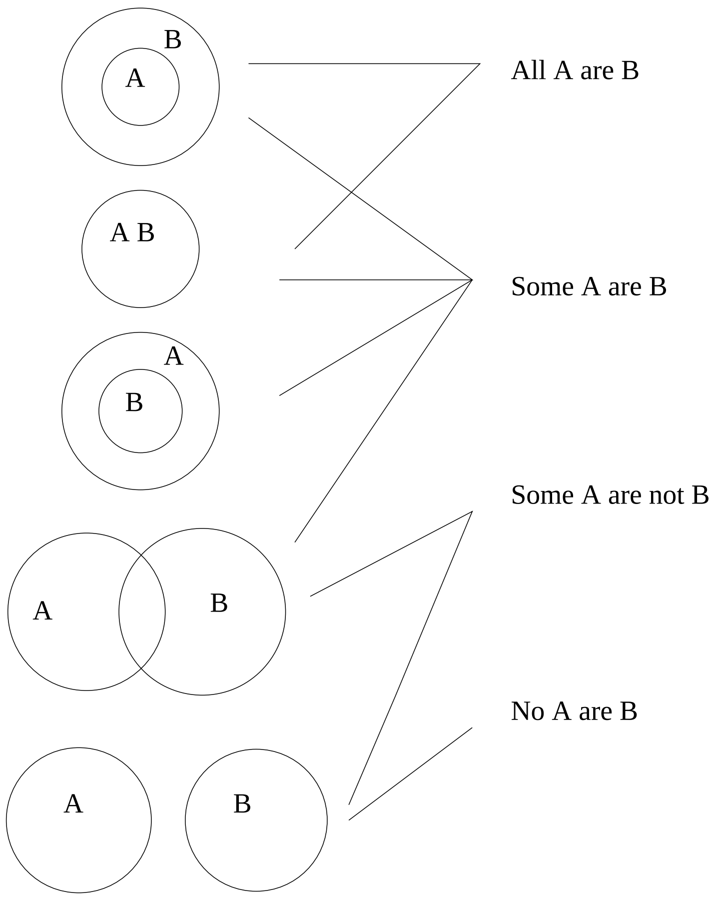

Figure 4: The many-to-many mappings of the topological relations between two circles onto syllogistic premisses.

He distinguishes *necessary*, *possible* and *impossible* types by annotations in sub-regions, or sometimes on the borders between sub-regions. Venn's method therefore does not exploit topological constraints to express logical exclusions. At most, geometry aids *search* of a somewhat unconventionally layed out table.

Because Venn's method does not exploit the graphical nature of its representations to exclude possibilities, there are less geometrical limitations on its extension. A theorem due to V. S. Poythress and H. S. Sun (1972) states that for N sets, there is a convex connected

plane figure (in fact, of concentric regular polygons) which has at least one minimal region representing every one of 2^N possible types. So for any finite number of properties, a maximal diagram is available. Sun-Joo Shin 1991, describes an extension of Venn's system

which uses a system of Xs joined by arcs to link regions into disjunctions of regions expresses a much larger fragment of monadic predicate calculus.

But for the cognitive theorist, Euler's method is richer because it does exploit the expressive limitations of graphical representations. It is altogether plausible that human perceptual and memory systems have developed mechanisms for implementing reasoning about the topological relations between circles. And that these mechanisms might be used to implement reasoning on the basis of analogies to other than spatial domains. If Euler's method relies on the topology of circles to express logical constraints, what do these constraints amount to in logical terms[3]?

Of course the first questions that arise are 'how can the graphical algorithm be interpreted as a logic?' and 'what extensions or generalisations are 'natural?'. The algorithm specifies how to express syllogistic premisses each as a diagram; how to unify pairs of diagrams; how to decide whether there are any valid conclusions of certain forms; and how to draw any valid sentential conclusions from the product diagrams. As such it embodies a theorem prover as well as a logic. A distinctive property of this graphical theorem prover is that it provides a test of consistency for any set of premisses. Inconsistent sets of premisses cannot be represented.

Three levels of generalisation of the syllogistic system are of interest. Firstly we consider the case where we permit repeated additions of information to an accumulating diagram but only by adding two-set diagrams of the sort we have seen represent syllogistic premisses. This is effectively generalising to many-premiss syllogisms.

Retaining the limitation on how information can be introduced to an accumulating diagram to the addition of syllogistic premisses restricts which diagrams can be uniquely specified. We can strengthen the permitted additions of information by allowing the marking of any sub-region as a critical region and then ask what logic results. And thirdly we can permit 'negative' marking of sub-regions—the addition of notations indicating emptiness of sub-regions. This latter addition takes us back to Venn's method because it removes any topological constraints on the expression of combinations of sub-regions. So the first two generalisations are what will concern us here.

The first extension, to multi-premiss reasoning raises questions about whether our graphical mechanisms can keep pace with a sentential extension. The second extension, to allow arbitrary marking of regions, focusses attention on the sentential logical fragment expressed by the topological constraints exploited by Euler's Circles.

[3] In what follows we consider only the expressive power of *single* diagram systems. Systems which, for example, allow arbitrary disjunctions of diagrams are of far greater expressiveness, but lack cognitive plausibility.

3.1 Generalising to many-premiss reasoning

Circles are much less expressive than even concentric regular polygons. A diagram with sub-regions for each of the 2^N combinations of inclusion and exclusion of N circles cannot be drawn for N > 3, and even for N = 3, about half the combinations of sub-regions cannot be diagrammed. So generalising our graphical algorithm to multi-premiss syllogisms immediately raises the question whether can we diagram the right maximal and minimal models for multi-premiss reasoning? As the number of predicates involved increases, the proportion of models which are diagrammable falls exponentially. Do the topological constraints which Euler's Circles exploit allow extension of the graphical method to multi-premiss reasoning despite this diminishing coverage? The answer turns out to be positive, the reason being that only a proportion of sets of premisses have any valid implications. And the maximal models of the sets that do licence conclusions can be represented by circle diagrams. Although there are consistent sets of premisses which cannot be diagrammed, there are no valid conclusions expressible in the syllogistic fragment which follow from these undiagrammable sets of premisses.

The reason that the graphical method keeps pace with multipremiss reasoning is that only unbroken 'chains' of universal premisses lead to valid inferences dependent on more than two premisses. Positive universals connected by a continuous string of middle terms in alternating grammatical position (*All A are B. All B are C. ...*) give rise to sets of nested circles. Thus hierarchies of predicates can be defined by sets of simple two-predicate premisses. So the extension of our graphical method to multi-premiss syllogisms does succeed in capturing the rather strictly circumscribed sets of premisses which licence long distance inferences. There is a coincidence between what is expressible in this diagrammatic reasoning system and the properties of this weak logical fragment.

3.2 Extending Euler's system by allowing marking of any minimal region

We now turn to the second extension of the graphical system by allowing the marking of any minimal region in a diagram with an x to mark it as a critical region. This extension immediately requires a specification of what the topological constraints on representation of models are, and what partitionings of models the free marking of regions can express.

3.2.1 Geometrical constraints defining Euler diagrams

As outlined above, the essential property of Euler's method is that it uses topological constraints to exclude types of individual inconsistent with the premisses represented. What models—combinations of maximal types—cannot be represented by reason of topology?

Some regions of diagrams can be eliminated (or added) without eliminating or adding any other areas. These are *solely eliminable* regions. If all regions in all diagrams were solely eliminable, then all possible combinations of regions would have corresponding circle diagrams. Examples of regions which are not *solely eliminable* appear in Figure 5 marked 'nse'.

There is a loose association between numbers of arcs bordering a region and its eliminability—the non-eliminable regions are bordered by three arcs (if one counts the infinitessimal borders with circle C in Diagram vi) of Figure 5). But there are 3-arc bordered regions which are solely eliminable (e.g. all regions in the 7-region diagram). The insightful generalisation about which sub-regions are solely eliminable has to do with the degree of geometrical constraint

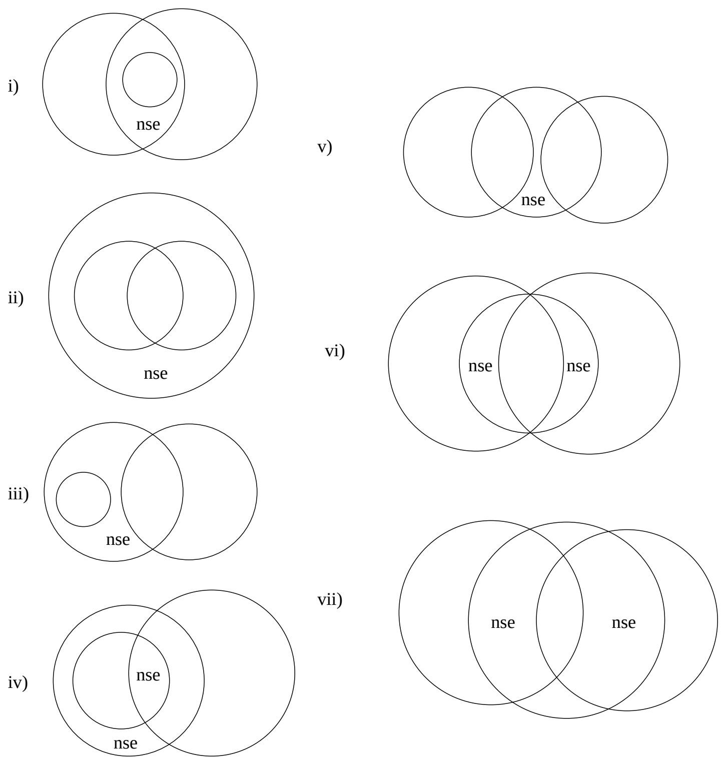

Figure 5: Ineliminable regions of Euler Circle diagrams

defining regions. We can specify all circle diagrams in terms of constraints. If rX signifies radius of circle X. The four premisses of the syllogism are constraints on the geometry of two circles:

- All A are B — rB > AB + rA
- Some A are B — rA > AB - rB
- No A are B — rA < AB - rB
- Some A are not B — rA < AB + rB

The existential premiss diagrams have additional constraints which ensure that they represent the maximal model, thus:

- Some A are B — rA < AB + rB
- Some A are not B —rA > AB - rB

To take one example of a full 3-circle diagram, Diagram vi) from Figure 5 can be specified by the following constraints:

- rA > AB - rB
- the centres of A, B and C are colinear
- AC = $F\_position(rA, rB, AB)$

- rC = $F\_size(rA, rB, AB)$

The position and size of C are wholly determined by the diameters of, and distance between, A and B and this explains why the marked regions are not solely eliminable.

Each not-solely eliminable region of a diagram gives rise to two models which cannot be distinguished by the graphical system, one containing the type corresponding to the region and one not containing that type. Diagrams containing N such regions correspond to 2^N graphically indistinguishable models. So, for example, Diagram vi) of Figure 5 has two not-solely eliminable regions and corresponds to 4 models which cannot be uniquely identified by combinations of regions in circle diagrams.

If not-solely eliminable regions define the combinations of regions which can and cannot be represented in circle diagrams, what partitionings of the space of 128 models can be expressed by marking arbitrary regions and combinations of regions as non-empty? Such diagrams can uniquely specify any model which contains elements corresponding to each of the regions in a diagram, simply by marking all its regions. But what other sets of models can be identified by diagrams? Each region left unmarked adds a disjunction of two possibilities—either there is an element of that type, or not. The representational system is like the tabular system mentioned above in which empty cells represent disjunctions of binary possibilities, but no contingencies between cells' contents. To go beyond this expressive power requires marking of empty regions (which removes all topological constraints) and beyond that, linking of regions into arbitrary disjunctions (as in Shin's Venn Diagram extension).

4 The general theory revisited

Erickson's interpretation of Euler's system is a MARS. Euler's system properly interpreted is a LARS, and Sun-Joo Shin's extended Venn system is an UARS. Although much remains to be done in the way of controlled comparative studies of these systems' cognitive efficacies, there is good empirical evidence that Euler's circles are useful didactic tools. Grossen & Carnine (1990) for example demonstrate that

embedded in a suitable teaching regime, Euler's Circles can actually raise the performance of learning impaired highschool students on several syllogistic reasoning tasks to the level of 'gifted' students. The controlled comparison of Euler's Circles with the more expressive Venn system has not, to our knowledge been carried out, but our theory would predict that the LARS middle ground is the cognitively useful one. There is strong evidence against the MARS interpretation of the system being useful—viz the combinatoric argument. There is less direct evidence that usefulness declines as the LARS system is augmented towards a UARS system such as Sun-Joo Shin's but there is at least some anecdotal evidence from teachers' experiences with these systems (Shin, personal communication).

Exploring what logics can be implemented within various topological and geometrical constraints is of potential cognitive interest because of the likelihood that special purpose mechanisms have developed for inferring the outcomes of transactions with the spatial world. Our aim should be to learn more about the abilities of human working memory to store and transform these topological patterns and about the exploitation of these abilities in general reasoning. Our study of Euler's circles brings this strategy of pursuing the study of these weak systems into sharp constrast with the other logical impulse—to generalise. Why circles the generaliser asks? Why not generalise to Euler's ellipses, or closed convex curves, or closed curves or .... The point is that the narrower the 'syntactic' restriction on the diagrams, the weaker the systems that results, the more likely that the constraint might be exploited in some implementation of the system. It is this weakness which we see as the particular cognitive virtue of graphical methods.

Acknowledgements

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/J586129 from SERC/DTI. I thank Richard Tobin for sorting out some of my errors.

References

Barker-Plummer, D. and S. C. Bailin (1992) Graphical theorem proving: an approach to reasoning with the help of diagrams. In *ECAI-92*, pp. 55–59.

Barwise, J. and J. Etchemendy (1990) Visual information and valid reasoning. In W. Zimmerman, ed., *Visualization in Mathematics.* Washington, D.C.: Mathematical Association of America.

Berkeley, G. (1937) *The Principles of Human Knowledge.* London: A. Brown. Text of the first ed. (1710) inc. variants in the 2nd (1734) and in an autograph ms. Ed. T. E. Jessop.

Cleveland, W. S. (1985) *The Elements of Graphing Data.* Monterey, Ca.: Wadsworth.

Coutaz, J. (1992) Multimedia and multimodal user interfaces: a taxonomy for software engineering research issues. In *EWCHCI*, pp. 229–240.

Erickson, J. R. (1975) A set analysis theory of behaviour in formal syllogistic reasoning tasks. In R. L. Solso, ed., *Information Processing and Cognition: The Loyola Symposium.* Hillsdale, N.J.: Lawrence Erlbaum Associates.

Grossen, G. and D. Carnine (1990) Diagramming a logic strategy: effects on difficult problem types and transfer. *Learning Disability Quarterly* **13**, 168–182.

Hartley, J. (1987) Designing electronic text: the role of print-based research. *Educational Research & Technology Journal* **35**(1), 3–17.

Hartley, J., ed. (1992) *Technology and Writing: Readings in the Psychology of Written Communication.* London: Jessica Kingsley.

Hovy, E. and Y. Arens (1990) When is a picture worth a thousand words? —allocation of modalities in multimedia communication. In *AAAI Symposium on HCI.*

Johnson-Laird, P. N. (1983) *Mental Models.* Cambridge: Cambridge University Press.

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

Levesque, H. J. and R. J. Brachman (1985) A fundamental tradeoff in knowledge representation and reasoning (revised version). In R. J. Brachman and H. J. Levesque, eds., *Readings in Knowledge Representation*, pp. 41–70. Los Altos, Ca.: Morgan Kaufmann.

Meyers, B. A. (1990) Taxonomies of visual programming and program visualization. *Journal of Visual Languages and Computing* **1**(1), 97–125.

Poythress, V. S. and H. S. Sun (1972) A method to construct convex connected Venn diagrams for any finite number of sets. *Pentagon* 80–82.

Rips, L. J. (1986) Mental muddles. In M. Brand and R. M. Harnish, eds., *Problems in the Representations of Knowledge and Belief.* Tucson: University of Arizona.

Salvendy, G. (1987) *Handbook of Human Factors.* New York: John Wiley and Sons.

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., 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 (1991) 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 (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.

Tufte, E. R. (1983) *The Visual Display of Quantitative Information.* Cheshire, Connecticut: Graphics Press.

Twyman, M. (1979) A schema for the study of graphical language. In P. A. Kolers, M. E. Wrolstad and H. Bouma, eds., *Processing Visible Language*, Vol. 1, pp. 117–150. New York: Plenum Press.

Yule, P. and K. Stenning (1992) The figural effect and a graphical algorithm for syllogistic reasoning. In *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*, pp. 1170–1175. Distributed by Lawrence Erlbaum Associates, Hillsdale, N.J.