

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Logical and Diagrammatic Reasoning: the Complexity of Conceptual Space

Permalink

<https://escholarship.org/uc/item/96k4c7j0>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 19(0)

Authors

Lemon, Oliver

Pratt, Ian

Publication Date

1997

Peer reviewed

Logical and Diagrammatic Reasoning: the Complexity of Conceptual Space

Oliver Lemon (LEMONOJ@CS.MAN.AC.UK)

Ian Pratt (IPRATT@CS.MAN.AC.UK)

Department of Computer Science, University of Manchester
Oxford Road, Manchester M13 9PL, UK

Abstract

Researchers currently seek to explain the observed tractability of diagrammatic reasoning (DR) via the notions of “limited abstraction” and inexpressivity (Stenning and Oberlander, 1995; Stenning and Inder, 1995). We point out that these explanations are inadequate, in that they assume that each structure to be represented (i.e. each model) has a corresponding diagram. We show that *inefficacy* (in the sense of incorrectness) arises in DR because some (logically possible) models fail to have corresponding diagrams, due to non-trivial spatial constraints. Further, there are good explanations of why certain restricted languages are tractable, and we look to complexity theory to establish such results. The idea is that graphical representation systems may be fruitfully analysed as certain restricted quantifier fragments of first-order logic, similar to modal logics and *vivid* knowledge bases (Levesque, 1986; Levesque, 1988). This focus raises some problems for the expressive power of graphical systems, related to their topological and geometrical properties. A simple case study is carried out, which pinpoints the inexpressiveness of Euler’s Circles and its variants. We conclude that there is little mileage in spatial (i.e. diagrammatic) approaches to abstract reasoning, except perhaps in relation to studies of human performance. Moreover, these results have ramifications for certain claims about mental representations, and the recent trend in cognitive semantics, where “meanings” and “concepts” are to be explicated spatially. We show that there should be combinations of “concepts” or “meanings” which are prohibited by the structure of the spaces they supposedly inhabit. The formal results thus suggest an empirical programme.

Introduction

Recent years have seen much effort in the explication of human information processing which employs diagrammatic representations (see Glasgow, Narayanan and Chandrasekaran (1995) for example). The much-lauded *efficacy* (Larkin and Simon, 1987) of reasoning with diagrammatic representations (DRs) has been explored both experimentally and theoretically (Stenning and Oberlander, 1995; Stenning and Inder, 1995), but there are as yet few formal results concerning the efficiency and expressive power of diagrammatic representation schemes. Further, recent logical analyses of diagrammatic representations, eg: (Hammer, 1995), do not account for the ways in which *spatial relations* are employed in representation. For example, overlap between regions may be used to represent set intersection, relative size of points may represent relative populations of cities, and so on. As

we shall show, neglecting the analysis of spatial representing relations leads to serious misapprehensions about the effectiveness of DR. In fact, we shall be careful to distinguish two types of efficacy of diagrammatic reasoning: representational efficacy (can the system represent all that it is required to?) and computational efficacy (what is the complexity of inference with the diagrams?)

Indeed, it is argued that current formal analyses fail on two counts (Lemon, 1997; Stenning and Lemon, 1997); (a) to do justice to the rich representational behaviour of DRs (in particular, their exploitation of spatial and mereological relations permitted by a medium of representation), (b) to account for their efficacy in computational terms.

We propose to remedy this situation as follows: first we consider some natural diagrammatic systems for logical reasoning, and some topological restrictions on them. The resulting lack of expressive power, and its potential computational pay-offs, are then explored. All in all, we shall consider three classes of language and their inter-relationships: the language S of set-inclusion and intersection statements, the diagrammatic system \mathcal{EC} of Euler’s Circles (and some weaker variants), and relational fragments of first-order logic.

Reasoning with convex regions

Suppose, then, that as in \mathcal{EC} a reasoner solves logical problems in the monadic predicate calculus by drawing regions of the plane representing the extensions of various predicates. For the present, let us suppose that these “blobs” representing the atomic properties are *convex*, and hence connected (but not necessarily circular). The idea is that n such blobs divide the plane into (at most 2^n) regions, and that each region represents a possible type of individual.

Determining whether a finite set of formulae in the monadic predicate calculus is consistent (which is of course equivalent to determining whether one formula follows from a finite set of formulae) can then be seen as a matter of determining whether there is an arrangement of convex regions in which all of these formulae come out true. Drawing “blobs” is a natural way to reason about combinations of properties. But there is just one snag: it doesn’t work, in general. To see this, consider the following reasoning task.

The examinations problem: A number of university stu-




Figure 1: The Helly constraint for two dimensions

dents have enrolled for end-of-term examinations in various combinations of subjects. There is no limit to the number of examinations a single student may enrol for. Four of the examination subjects are A(lgebra), B(iology), C(chemistry) and D(ivinity). Some students have enrolled for (at least) A, B and C; some students have enrolled for (at least) B, C and D; some students have enrolled for (at least) A, C and D; and some students have enrolled for (at least) A, B and D. (In other words, any trio of subjects is taken by some students.) What follows?

Answer: not much. Trivially, some students take A and B; some students take A, and so on. But it should be clear that nothing else of substance follows. However, that is not the answer that you would get if you represented the state of affairs by drawing convex regions on paper. And the reason for this is the following standard result:

Theorem 1 (Helly's theorem) *Let X_1, \dots, X_N be convex regions in n -dimensional Euclidean space, $N \geq n + 1$, such that any $n + 1$ -membered collection of the X_1, \dots, X_N has a nonempty intersection. Then X_1, \dots, X_N has a nonempty intersection.*

For instance, let $N = 4$ (for the 4 regions A, B, C, D), and since the regions are plane, $n = 2$. Then if every trio of the regions has an intersection, all four of them must intersect too. See figure 1, where there is no way to add a new convex region (eg: the dotted ellipse) overlapping each pairwise intersection (the regions denoted by "X") of the others, without also producing a quadruple intersection (in the region marked "O"). Regarding the above examinations problem, then, no matter how you draw the blobs, provided they are convex, the intersection of A, B, C and D will unavoidably turn out non-empty. This forces any would-be diagrammatic reasoner to draw the *unwarranted conclusion* that some students take all four examinations. The upshot of the examinations problem is that spatial representations isomorphic to convex blobs would be a bad idea for reasoning in the monadic predicate

calculus¹, even if that reasoning were restricted to relatively simple problems involving up to four terms. (For three or more dimensions, note that increasing the number of examinations leads to analogous errors.)

The important point to notice about the examinations problem is the extent to which it relies on the spatial nature of the representations involved. Helly's theorem is non-trivial where $n > N + 1$. That is, it identifies a constraint on the representational system of convex blobs which does not arise from logic alone. That is why this representation scheme is bound to yield incorrect inferences, because it cannot represent some logically possible situations. In other words, as we have seen above, the spatial representation scheme may "force" representations which are geometrically, rather than logically, necessary. The representation scheme is "over-specific" or "information enforcing" to use the terminology of Shimojima (1996) or Stenning and Inder (1995). If people really reasoned by drawing something like "pictures in their heads", the inference that some students take all examinations is just the sort of error one would expect. Far from being merely an intuitive characterisation of representational inefficacy (as is common in the DR literature) we can show in this instance precisely where "mistakes" with DR might spring from. Before beginning an analysis of the expressive power and complexity of certain diagrammatic systems, we shall consider some other topological constraints on diagrams.

Reasoning with connected (non-convex) regions

To obviate the problem described above, a spatial representation scheme for the monadic predicate calculus must use at least *non-convex* blobs. Again, the system using non-convex connected plane regions is a less constrained version of the convex "Euler regions" system presented above (which itself is a less constrained version of Euler's Circles). Thus the results derived for non-circular and non-convex plane regions certainly apply to the standard system \mathcal{EC} .

Let us suppose that the blobs representing the atomic properties are *connected* (i.e. "one piece" regions). Otherwise, let their interpretation be as for the convex regions. But might not a similar problem arise here? Might there not be a non-trivial property of connected regions in 2D which renders them similarly unsuitable as a representation scheme? Again, the use of connected plane regions to reason about properties strikes one as quite natural. And again, it doesn't work, in general. To see this, consider the following reasoning task.

The musicians problem: Nine musicians, A(lison), B(rian), C(ornelia), ... and I(an) play various pieces of music in all sorts of combinations. Some pieces involve at least the following players (possibly others as well):
ABC, DEF, GHI, ADG, BEH, CF, CI, FI.

In addition, no pieces involve any two musicians not grouped together in the above list. (For example, A and

¹There is a suggestion that an analogous result exists, due to Martin Gardner, though as yet we have been unable to confirm it.

I are not grouped together, so no pieces involved both Alison and Ian.) What follows?

Again, not much, apart from the trivial inferences that some pieces involved both Alison and Brian, no pieces involved Alison, Brian and Ian, and so on. But that is not the answer that you would get if you represented the state of affairs by drawing connected (but not necessarily convex) blobs on paper. It turns out to be impossible to realize the above arrangement so as to make C, F and I overlap, corresponding to the conclusion that no piece involved Cornelia, Fiona and Ian. Again, this conclusion is unwarranted, because there could have been six trios ABC, DEF, GHI, ADG, BEH, CFI, consistent with the statement of the problem.

The reason for this is that the situation describes a non-planar graph (Kuratowski, 1930), if the relation of playing-music-together is represented as overlap of regions. We shall see this property of limited representability for overlap relation in the plane work out in detail in the proof of theorem 2 below.

Complexity of fragments of FOL

The primary motivation for investigating diagrammatic languages as spatially restricted logical languages, is computational. We wish to establish the noted computational *efficacy* (meaning tractability) of DRs by way of the complexity properties of the logics to which they correspond.

It is well known that certain fragments of FOL are decidable, and enjoy polynomial satisfiability. For example, satisfiability of the Horn fragment of FOL is in P (Papadimitriou, 1994). Other interesting fragments include the purely universal fragment, the purely existential fragment, the monadic predicate fragment, and restricted quantifier fragments (corresponding to modal logics). In connection with spatial logics, Bennett's intuitionistic logic (Bennett, 1994) for qualitative spatial reasoning has been shown to be a polynomial time fragment (Nebel, 1995) (actually, it is in NC; efficiently solvable on parallel machines). We shall begin to ask similar questions about logical fragments which correspond to diagrammatic systems.

Constraints and expressive power

Here we consider the simple diagrammatic language of non-convex connected regions of the plane for representing and reasoning about \mathcal{S} , the language of set inclusion and intersection. This example illustrates the proposed analysis of diagrammatic efficacy in terms of expressive power and complexity theory.

Consider the simple diagrammatic language where connected regions of the plane are interpreted as sets. Set s properly includes set s' if and only if the connected region $r_{s'}$ representing set s' is *entirely contained* in the region r_s representing set s . Similarly, set s has a non-empty intersection with set s' if and only if the region $r_{s'}$ representing set s' *overlaps* the region r_s representing set s . Thus, the diagrammatic system allows us to express facts about set inclusion and intersection.

Indeed it is trivial to note that the system exhibits *self-consistency*: every diagram may be interpreted as consistent set of set-inclusion and intersection statements. This consistency is guaranteed simply by the properties of regions of the plane; inclusion of regions is transitive (as is set inclusion), overlap of regions is symmetric (as is set intersection), and so on. However, although the structural restrictions (transitivity, symmetry) on the set-theoretic operations are preserved in the spatial restrictions on inclusion and overlap of connected plane regions, it turns out that this relation of overlap obeys more constraints than that of set intersection.

Constraint mis-matching

"Constraints", to use the terminology of Barwise and Shimojima (1995) and Shimojima (1996), are the restrictions inherent in a class of structures (for example, that collections of convex regions obey Helly's theorem, or that set intersection is symmetric.) Problems of representational efficacy in DR arise when there is a mis-match between the constraints of the diagrammatic system, and those of the system it is supposed to represent.

Consider the systems above. The set-inclusion relation forms a strict partial order. Similarly, the diagrammatic relation of one region being inside another also forms a strict partial order. However, set *intersection* relations form a symmetric structure, while *overlap* relations in \mathcal{EC} form a symmetric structure *in the plane*. Thus, in some cases, \mathcal{EC} is *more constrained than* \mathcal{S} . In fact, we may identify these problematic cases as those corresponding to the non-planar graphs, leading to the following result.

Limited representability of \mathcal{EC} : non-planarity

Let *full representability* of a diagrammatic language \mathcal{G} with respect to a set of models \mathcal{M} be the property that every model $m \in \mathcal{M}$ may be represented by a diagram of \mathcal{G} . Thus, for the full representability of \mathcal{EC} , we need to show that any logically possible collection of sets may be drawn as a diagram of that system. Sadly, as we shall show, this is not generally true. Just as it is impossible to use \mathcal{EC} to represent a contradiction, it turns out to be impossible for \mathcal{EC} to represent some logically possible models too.

Theorem 2 *There are consistent sets of set intersection statements which cannot be represented by any diagram of \mathcal{EC} .*

Proof: Consider the system where sets are represented by (non-convex) connected regions of the plane. Let there be 5 sets $v_1 \dots v_5$, and 10 sets e_{ij} , $1 \leq i < j \leq 5$, such that the following constraints hold;

1. $v_i \cap e_{ij} \neq \emptyset$
2. $v_j \cap e_{ij} \neq \emptyset$
3. $e_{ij} \cap e_{i'j'} = \emptyset$ if $i \neq i'$ or $j \neq j'$
4. $v_i \cap v_j = \emptyset$ if $i \neq j$
5. $v_k \cap e_{ij} = \emptyset$ if $k \neq i$ and $k \neq j$

This situation is logically possible, but this is not the conclusion you would reach if you were to represent the sets by way of connected regions of the plane. To see why, let the sets be regions of the plane. Since regions v_i are connected and $\forall j \neq i, e_{ij} \cap v_i \neq \emptyset$, we can create nodes $V_i \in v_i$ and $V_{ij} \subset v_i \cap e_{ij}$, and edges $\langle V_i, V_{ij} \rangle$ which intersect only at V_i . Now add edges $\langle V_{ij}, V_{ji} \rangle$ for all $1 \leq i < j \leq 5$, drawn entirely within e_{ij} , since the e_{ij} are connected. The resulting graph is Γ_5 , which is non-planar (Kuratowski, 1930). Thus the \mathcal{EC} representation forces there to be an overlap violating the above constraints. As regards the “musician’s problem” (stated above); it relies on the fact that a similar construction may be carried out for the non-planar graph $\Gamma_{3,3}$, using 9 regions.

Obviously, if we restrict the number of sets which we wish to reason about, so that the planarity problems cannot arise, then the (non-convex) diagrammatic system exhibits full representability of the restricted problem domain. To summarize, the following restrictions apply, for full representability, in the sense that (as far as we know) spatial difficulties do not arise if restricted to reasoning about the following numbers of regions.

Convex 2D \mathcal{EC} : 3 sets.

Convex 3D \mathcal{EC} : 4 sets.

Non-convex connected 2D \mathcal{EC} : 8 sets.

Constraints for convex regions arise from the Helly property, and those for connected plane regions arise from planarity considerations. As far as we know, there are no restrictions on representability for representation systems which employ 3D non-convex representations, or for those which employ non-connected regions. However, such systems, although perhaps nominally spatial, are so unconstrained as to fail to be diagrammatic in any contentful sense.

Given that effective diagrammatic systems exploit spatial properties of the medium, how might a logical analysis proceed? We now describe a class of first-order languages which allow us to investigate diagrammatic systems more formally.

Logics and Diagrammatic systems

We identify a certain class of languages, defined over $\mathcal{L}_r = \{\wedge, \neg, =, Cons, \mathcal{P}\}$, where $Cons$ is a set of distinct constants, \mathcal{P} is a set of distinguished predicates, and negation only applies to atomic formulae. These fragments of FOL are further constrained to exhibit only restricted quantification for their distinguished predicates (in the sense used in correspondence theory for modal logics), are implication-free and disjunction-free, and thus avoid expressing indeterminacy. We propose that some DR systems exhibit a formal correspondence to certain of these languages, where distinct constants stand for distinct connected regions of the plane, and relations between constants are constrained so as to capture the structure of these plane regions. Further, the languages in \mathcal{L}_r operate under the following conditions;

1. (Unique Names/Specificity)
all constants are non-identical (ie: regions are distinct).
 $\forall c_i, c_j \in Cons, \text{ if } i \neq j, \text{ then } c_i \neq c_j$
2. (Closed World Assumption)
For all consistent sets of sentences $T \subset \mathcal{L}_r$,
 $T \not\vdash \phi \Rightarrow T \vdash \neg\phi$

Note that these languages are similar to Levesque’s notion of a *vivid knowledge base* (Levesque, 1986; Levesque, 1988); a first-order language which contains only ground atomic sentences, inequalities between all constants, universally quantified sentences over the domain, and whose predicates obey the closed-world assumption. Determining entailment in such knowledge bases is known to be tractable. Indeed, Levesque speculates further that,

“perhaps the main source of vividly represented knowledge is *pictorial information*.” (Levesque, 1986).

Similar claims may be found in literature on the logical form of pictures as mental representations (Howell, 1976; Sober, 1976).

Complexity of spatial inference

Recall our motivation for investigating diagrammatic languages as spatially restricted logical languages. We wish to establish the computational *efficacy* of DRs by way of the complexity properties of the logics which they embody. We have seen that reasoning with \mathcal{EC} (and its variants) is only correct for a small number of sets. When restricted in this way, reasoning with \mathcal{EC} is clearly polynomial (it reduces to table look-up).

Two cheers for DR?

But what of the complexity of the unrestricted systems (where non-planar representations may occur)? We know that this version of \mathcal{EC} is incorrect for set-theory, and worse still, that its complexity is that of reasoning about overlap and inclusion for connected regions of the plane – recently shown to be NP hard (Grigni et al., 1995). Furthermore, the use of circles (as opposed to simply connected regions) may impose even more geometrical constraints on the representations than those which we have considered.

More positively, note that although the spatial restrictions on possible diagrammatic representations lead to incorrect inferences in the cases presented here, they *are* effective in the representation of similarly constrained structures (most obviously, those which themselves obey spatial constraints). This fact, for example, makes cartography a successful venture.

Such considerations point to the conclusion that there is little mileage in spatial (i.e. diagrammatic) approaches to abstract reasoning, unless one is fortunate enough to be able to prove that the problem domain obeys all the (topological and geometrical) restrictions inherent in the chosen diagrammatic system. Of course, this leaves open the possibility that diagrammatic reasoning is interesting from the point of view of a model of human performance.

Implications for spatial modelling

We conclude with some wider considerations for cognitive scientists generally, who might be tempted to use spatial structures in their cognitive models.

Some mileage has been found, in recent years, in the claim that there is a spatial structure to concepts and word meanings, eg: (Gärdenfors, 1996; Gärdenfors, 1995). Cognitive semantics often employs the idea that meanings and concepts have a spatial nature; that our mechanisms for reasoning about and representing space are also made use of in our representations of other, more abstract, properties, entities, and relations. Thus, categories or concepts, as well as word meanings, are supposed to have a spatial structure². The spatial version of the idea that meanings are mental entities holds that they exist as positions or regions within “conceptual spaces”. In particular, it seems that an analysis of *metaphor* and of *spatial prepositions* have given rise to this view. For example:

“What characterizes a metaphor is that it expresses a similarity in topological or metrical structure between different quality dimensions.” (Gärdenfors, 1995)

Thus a conceptual space is taken to be determined by a number of quality dimensions. *Regions* of conceptual space are then to be understood as spatial entities, with respect to the topology and metric of that space. For example, we might talk of the “opposition” of colours red and green, or of political parties, of “lengths” of time, and so on. These are not just ways of speaking, it is claimed, but they rely on the structure of “conceptual spaces” which really do have some particular spatial structure and properties.

As we have seen above, such spatial representational structures must be constrained in important ways (i.e. their topologies and metrics must be specified), or else they are structures of some other sort, and “spatial” in name only. If internal representations are to be interestingly spatial they must satisfy some particular structural properties. So we ask how seriously this “spatial turn” in cognitive science can be taken. The results presented above show that topology and geometry must constrain possible spatial representations – that “conceptual spaces” cannot represent every combination of their constituents. That means that certain combinations of concepts or meanings should be impossible in certain conceptual spaces. These “impossible representations” should be empirically detectable. Perhaps this is the kind of evidence cognitive semanticists should seek. Otherwise the claims of the “conceptual space” version of cognitive semantics are in danger of being unfalsifiable.

An Empirical Programme

Regarding the formal results presented above, we currently have only informal evidence that people typically do not make the kinds of inferential mistakes that the results predict for the use of convex or connected 2D regions. For instance, given the examinations problem, subjects typically do

²This structure is argued to support non-monotonicity and metaphor.

not make the inference that some student takes all four examinations, but this is the mistake that they would be forced to make (via Helly’s theorem) if they happened to represent the problem by way of convex plane regions in \mathcal{EC} . Similarly, we found that subjects are not tempted to make mistakes predicted by the planarity constraints in the musician’s problem. However, since 9 regions are involved in this problem, it seems unlikely that any would-be diagrammatic reasoner would attempt the problem without pencil and paper. So it may be that “diagrams in-the-head” are only used where small numbers of regions are required. Again, experimental work remains to be done on these issues.

The general empirical issue arising is that of to what extent the “spatial” structures posited in cognitive science meet requirements on interestingly spatial representations. If they do not obey contentful spatial constraints, we claim that they are merely structural descriptions, as “spatial” as any other theory. The concern is that this spatial metaphor, as it stands in cognitive science, may be an empty one. Our proposal is that, empirically and formally, “spatial” hypotheses in cognitive science ought to be investigated by way of the structural constraints imposed by the use of space as a representational medium, and its potential computational payoffs.

For example, by the topological results applied above, there ought to be configurations of regions in every “conceptual space” which are not possible within that space: that is, combinations of representations (“concepts” or “meanings”) which are prohibited by the structure of the space they inhabit. There could be empirical studies exploring such phenomena, as well as formal results establishing the computational properties of such spatially restricted representation languages. But we are unaware of any such studies to date.

Conclusion

The representational power of diagrammatic systems such as Euler’s Circles is investigated with respect to their use in solving simple logical conundrums in set theory. Topological results expressing constraints on possible diagrams are used to show that certain logically possible configurations (models) cannot be represented diagrammatically (or, in general, spatially). We conclude that diagrammatic reasoning is only effective for a certain tightly constrained set of problems, and is only interesting as a potential model of human performance. Even there we are unaware of any study which investigates the impact of spatial constraints on possible representation and reasoning strategies.

It is proposed that the efficacy (representational and computational) of diagrammatic systems be explicated via the expressive power and computational properties of the restricted languages of which they make use. We believe that the computational properties of language fragments may be used to explain the efficiency of diagrammatic reasoning in more detail than the “limited abstraction” hypothesis of (Stenning and Oberlander, 1995). We identify an interesting class of (restricted first-order) languages in this regard. This first attempt to apply the concepts of complexity theory to systems

of diagrammatic representation is a necessity in the evaluation of non-psychological claims about their tractability. Furthermore, ramifications of these results for a current "spatial" trend in certain branches of cognitive science, are discussed.

A potentially fruitful ground for collaboration has been prepared; that between spatial logic, formal semantics, complexity theory, and the analysis of diagrammatic representation systems.

Acknowledgements

With thanks to Jeff Paris. This research is supported by the Leverhulme Trust (grant number F/120/AQ) and the British Council's British-German academic research collaboration programme, project number 720.

References

- Barwise, J. and Shimojima, A. (1995). Surrogate Reasoning. *Cognitive Studies: Bulletin of Japanese Cognitive Science Society*, 4(2).
- Bennett, B. (1994). Spatial Reasoning with Propositional Logics. In Doyle, J., Sandewall, E., and Torasso, P., editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the 4th International Conference (KR '94)*, San Francisco, CA. Morgan Kaufmann Publishers.
- Gärdenfors, P. (1995). Meanings as Conceptual Structures. Technical Report 40, Lund University Cognitive Science.
- Gärdenfors, P. (1996). Mental Representation, conceptual spaces, and metaphors. *Synthese*, 106:21-47.
- Glasgow, J., Narayanan, N. H., and Chandrasekaran, B., editors (1995). *Diagrammatic Reasoning: Cognitive and Computational Perspectives*. AAAI Press / The MIT Press, Cambridge, Mass.
- Grigni, M., Papadias, D., and Papadimitriou, C. (1995). Topological Inference. In *International Joint Conference on Artificial Intelligence (IJCAI '95)*. AAAI Press.
- Hammer, E. M. (1995). *Logic and Visual Information*. Studies in Logic, Language, and Computation. CSLI Publications and FoLLI, Stanford.
- Howell, R. (1976). Ordinary Pictures, Mental Representations, and Logical Forms. *Synthese*, 33:149 - 174.
- Kuratowski, G. (1930). Sur le probleme des courbes gauches en topologie. *Fund. Math.*, 15:271-283.
- Larkin, J. and Simon, H. (1987). Why a Diagram is (Sometimes) Worth 10,000 Words. *Cognitive Science*, 11:65-99.
- Lemon, O. (1996). Semantical Foundations of Spatial Logics. In Aiello, L. C., Doyle, J., and Shapiro, S. C., editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Fifth International Conference (KR '96)*, pages 212 - 219, San Francisco, CA. Morgan Kaufmann Publishers.
- Lemon, O. (1997). Review of "Logic and Visual Information" by E. M. Hammer. *Journal of Logic, Language, and Information*. (forthcoming).
- Lemon, O. and Pratt, I. (1997). On the incompleteness of modal logics of space: advancing complete modal logics of place. In Kracht, M., de Rijke, M., Wansing, H., and Zakharyashev, M., editors, *Advances in Modal Logic*. CSLI Publications, Stanford. (forthcoming).
- Levesque, H. J. (1986). Making Believers out of Computers. *Artificial Intelligence*, 30:81 - 108.
- Levesque, H. J. (1988). Logic and the Complexity of Reasoning. *Journal of Philosophical Logic*, 17:355-389.
- Nebel, B. (1995). Computational Properties of Qualitative Spatial Reasoning: First Results. In Wachsmuth, I., Rollinger, C.-R., and Brauer, W., editors, *KI-95: Advances in Artificial Intelligence*. Springer-Verlag, Berlin. (19th German Conference on Artificial Intelligence).
- Papadimitriou, C. (1994). *Computational Complexity*. AddisonWesley, New York.
- Shimojima, A. (1996). *On the Efficacy of Representation*. PhD thesis, Indiana University.
- Sober, E. (1976). Mental Representations. *Synthese*, 33:101 - 148.
- Stenning, K. and Inder, R. (1995). Applying Semantic Concepts to Analyzing Media and Modalities. In Glasgow, J., Narayanan, N. H., and Chandrasekaran, B., editors, *Diagrammatic Reasoning: Cognitive and Computational Perspectives*. AAAI Press / The MIT Press, Cambridge, Mass.
- Stenning, K. and Lemon, O. (1997). Diagrams and human reasoning: aligning logical and psychological perspectives. In Blackwell, A., editor, *Thinking with Diagrams: discussion papers*. Portsmouth. (under revision).
- Stenning, K. and Oberlander, J. (1995). A Cognitive Theory of Graphical and Linguistic Reasoning: Logic and Implementation. *Cognitive Science*, 19(1):97 - 140.