Words, Proofs and Diagrams

Dave Barker-Plummer
David I. Beaver
Johan van Benthem
Patrick Scotto di Luzio

ISBN: 1-57586-406-1

Contents

Contributors

MARCO AIELLO
Institute for Logic, Language and Computation, Room P.326 – Plantage Muidergracht 24,1018 TV, Amsterdam, The Netherlands.
aiellom@ieee.org

MARIA ALONI
Faculteit der Geesteswetenschappen, Leerstoelgroep Taalfilosofie, Universiteit van Amsterdam, Nieuwe Doelenstraat 15, 1012 CP Amsterdam, The Netherlands.
aloni@hum.uva.nl

DAVID I. BEAVER
Department of Linguistics, Margaret Jacks Hall, Stanford University, Stanford, CA 94305-2150, USA.
dib@csli.stanford.edu

JOHAN VAN BENTHEM
Institute for Logic, Language and Computation (ILLC), University of Amsterdam, Plantage Muidergracht 24, 1018 TV Amsterdam, The Netherlands; *and*
Department of Philosophy, Stanford University, Stanford, CA 94305-2155, USA.
johan@science.uva.nl

KATHI FISLER
Department of Computer Science, Worcester Polytechnic Institute, 100 Institute Road, Worcester, MA 01609-2280, USA.
kfisler@cs.wpi.edu

PETER GRÜNWALD
EURANDOM, P.O. Box 513, 5600 MB The Netherlands.
pdg@cwi.nl

Alexander Koller
Department of Computational Linguistics, Universität des Saarlandes, Geb. 17.2, Room 1.03, 66041 Saarbrücken, Germany.
koller@coli.uni-sb.de

Oliver Lemon
Center for the Study of Language and Information, Stanford University, Cordura Hall, 210 Panama Street, Stanford, CA 94305-4115, USA.
lemon@csli.stanford.edu

Joachim Niehren
Universität des Saarlandes, Postfach 15 11 50, 66041 Saarbrücken, Germany.
niehren@ps.uni-sb.de

Valeria de Paiva
Xerox PARC, 3333 Coyote Hill Road, Palo Alto CA 94304, USA.
paiva@parc.xerox.com

Martin Stokhof
ILLC/Department of Philosophy, Faculty of Humanities, Universiteit van Amsterdam, Nieuwe Doelenstraat 15, NL-1012 CP Amsterdam, The Netherlands.
stokhof@hum.uva.nl

Hans-Joerg Tiede
Illinois Wesleyan University, Department of Mathematics and Computer Science, P.O. Box 2900, Bloomington, IL, 61702-2900, USA.
htiede@iwu.edu

Dag Westerståhl
Dept. of Philosophy, Goteborg University, Box 200, 405 30 Goteborg, Sweden.
dag.westerstahl@phil.gu.se

Foreword

Academic boundaries can be very permeable, and they sometimes seem almost arbitrary. The last twenty years have witnessed remarkably lively encounters between computer scientists, logicians, linguists, philosophers, psychologists, and economists. This is not because of a sudden urge to be 'interdisciplinary', but researchers from these disciplines often find that, when all is said and done, they are studying the same problems! On this common agenda are ancient questions concerning how humans communicate, reason and decide, and modern questions about how computers should communicate, reason and decide — and how communication is best organized between these two species. Issues of information structure lie at the heart of this, as well as the algorithmic structure of the processes that convey and transform information.

The resulting field does not yet have an established label, and various terms are current, including package names like 'Logic, Language and Computation', or recent neologisms like 'Informatics'. Likewise, there is no single agenda, expressed in some convenient manifesto. What constitutes the field are the collaborative efforts of many researchers and teachers, the cooperative spirit of a number of research institutes spread across the world, and a number of international conferences and summer schools.

The Center for the Study of Language and Information at Stanford University plays a key role in this international effort, both through its many activities and through the facilities it offers as a meeting place for intellectual border crossers. In particular, it hosts a series of annual international workshops on Logic, Language and Computation. The closest thing to an up-to-date agenda for the field is what you will see in the contents of the proceedings of workshops like this.

The Eighth CSLI Workshop on Logic, Language and Computation was a good illustration. It covered a good deal of active research areas at the interface of logic, computer science, and linguistics - and readers will find 'old favorites' like process logics, or issues in formal semantics, or language processing. In addition, the Workshop highlighted one particular new area where all three disciplines meet, namely, the study of images and graphics as information carriers and the diagrammatic reasoning supported by them. Well-known issues of expressive power and computational complexity return here in a new and intriguing guise.

The remainder of this volume is divided into three parts: Diagrammatic Reasoning, Computation, and Logic and Language. Each of these parts is headed by an editorial introduction that maps out the relation of the papers to each other and to the wider field. Each chapter is well-worth reading. But it is their interconnections that give 'the total picture' of the field today.

We are grateful to everyone who participated in the making of this volume, with special thanks going to the reviewers, David Israel, Stefan Kaufmann, Dikran Karagueuzian, and the team at CSLI Publications. The Eighth CSLI Workshop on Logic, Language, and Computation was made possible by the generous support of Vaughan Pratt, the Department of Linguistics, the Department of Philosophy, the School of Humanities and Sciences (all at Stanford University), the Spinoza project at ILLC (University of Amsterdam), and CSLI.

Dave Barker-Plummer
David I. Beaver
Johan van Benthem
Patrick Scotto di Luzio

Diagrammatic Reasoning

Editorial Introduction

Dave Barker-Plummer

In the papers here, logic is unleashed on the study of reasoning with diagrammatic representations. All three of the papers that follow apply techniques from logic to the study of questions concerning the expressiveness and efficacy of diagrammatic notations for reasoning, with quite different emphases and intentions.

Aiello and van Benthem revive Tarski's topological interpretation of modal operators to provide a game-theoretic account of how topologically different, or similar, spaces can be seen as such. This theory is then extended to stronger kinds of equivalence than topology, with geometric equivalence being the strongest.

The focus of this work is on the structure of diagrams and on identifying techniques for measuring similarities between diagrams. The work reported here could, for example, be used in service of retrieval of diagrams from a database on the basis of their similarity to some target diagram. A variety of questions arise from this work, most notably the degree to which the varying notions of equivalence distinguish between similar diagrams and whether these notions of similarity correspond to those that we perceive. Which notion of equivalence is appropriate when we ask an automated system to retrieve a diagram from a database that looks "just like" this one?

Another possible application of Aiello and van Benthem's ideas is in evaluating the relative expressivity of diagrammatic languages. If we are given two diagrams of the same situation in different representations, then the non-equivalence of the two diagrams would indicate a difference in expressivity of the representations — one can capture more subtle aspects of the situation than the other.

In contrast to Aiello and van Benthem's focus on diagrammatic structure, the papers by Lemon and Fisler focus on the task of reasoning with information presented in diagrammatic form.

Lemon discusses a variety of limiting results which demonstrate the

Words, Proofs, and Diagrams
D. Barker-Plummer, D. I. Beaver, J. van Benthem, P. Scotto di Luzio (eds.)

difficulties of selecting diagrammatic representations appropriate for particular reasoning tasks. The primary exemplar of this is the use of Euler circles for reasoning about set containment and intersection. Lemon observes that while this representation is adequate, indeed perspicuous, for treatment of the syllogism, extending the same representation to arguments that involve four or more monadic predicates will result in unsoundness. A host of similar limiting theorems and their applications to other representations are also discussed.

The significance of Lemon's paper is not that one must choose representations appropriate to the domain of interest (this, presumably, we understood), but rather the degree to which this can be a subtle and difficult task. The failure of Euler circles to generalize outside of representing syllogisms represents a difference in the complexity of the problem (four monadic predicates participating in the argument rather then three), not in a change of domain. It is as if we could soundly use *modus ponens* in arguments containing three predicates but not four. One could surely look at and accept many examples of problems represented in a new diagrammatic system before noticing that the system contained a limitation of this kind.

Perhaps van Benthem and Aiello's work has some relevance to Lemon's observations. Rather than completely discard Euler circles as a way of representing reasoning with monadic predicates, an alternative approach based on an analysis of the structure of the diagram might be warranted. Even some arguments with more than three predicates can be represented soundly, and so an examination of the structure of the diagram is required to determine which conclusions can be validly drawn from the representation.

In this work, Lemon points the way to a potentially fruitful research question. Which diagrammatic representations suffer from the problems that are discussed here when used as models of which domains?

One obvious extreme is the case of photorealistic images of a represented domain, or more abstractly, diagrammatic representations that are isomorphic to the represented domain. A photograph of a situation must represent the situation faithfully, and therefore not introduce unwarranted inferences about the domain (inferences that would not be warranted when looking at the scene from the point of view of the camera, of course). As Lemon points out, there are many interesting systems of this kind, for example maps. Which kinds of abstractions from this kind of faithful representation lead to the problems that Lemon describes?

Fisler's examination of timing diagrams from a computational perspective also throws light on some interesting questions concerning the use of diagrammatic representations. It is often remarked that diagrammatic representations facilitate reasoning by making some inferences obvious. Different facts have been pointed out as relevant to this phenomenon, ability of diagrams to represent space, Shimojima's work on "free rides", arguments about homomorphisms between the representation and represented domain. None of these appear to tell the complete story.

Fisler's work suggests another reason based on the observation that the timing diagrams that she discusses can capture languages at different levels of the Chomsky hierarchy without capturing all languages at that level. This indicates that there may be expressive advantages to gain in adopting diagrammatic representations that are simply not available in sentential representations.

Fisler remarks that it is an open question whether timing diagrams represents a new class of expressive languages corresponding to a new computational model. I would add to that the question of whether this is a general feature of diagrammatic representations. What features of timing diagrams as Fisler defines them account for the expressive differences between this representation and sentential representations? How can these observations be generalized to build a more general theory of the expressiveness of diagrammatic representations? Fisler's paper indicates an interesting new approach for researchers interested in the inherent differences between diagrammatic and sentential representations.

With tools for capturing equivalences between representations, theorems describing limitations of classes of representations and new observations about the relative power of diagrammatic and sentential representation, the papers here provide new ideas and insight for researchers in diagrammatic reasoning.

Computation

Editorial Introduction

Johan van Benthem

The contributions in this part 'take the pulse' of current research on logic and computation. While addressing a variety of issues, as befits a lively interface, they also display a number of parallel trends, that become visible in the setting of workshops like CSLI's 'Logic, Language, and Computation'. Let us notice a number of these, and see what they mean in general.

First, it is clear that subject boundaries mean very little in these contributions. E.g., techniques from mathematical proof theory and computational automata theory occur naturally in Tiede's analysis of natural language. This illustrates the continued validity of *Montague's Thesis* that there is little difference in principle between the study of natural languages and formal languages in mathematics. It has been clear for some time, of course, that the Thesis can be extended to include programming and specification languages from computer science, and representation languages in AI.

The use of *formal proofs* is striking in the contributions by De Paiva and Tiede. Both advocate the export of mathematical proof theory, one toward syntactic parsing, the other toward general semantics. Moreover, both operate on a platform of resource-sensitive logics, broadly conceived. But these uses at once raise some general questions. Proofs have been used in categorial grammar as a general notion of 'logical form' for a sentence, which does two things: it gives the syntactic construction, and at the same time, it provides effective information about the intended semantic interpretation. Thinking in traditional terms, this seems like a *representational* use. But on the other hand, proofs have a clear *computational* content, which is more the raison d'etre for their use in computer science. How can we put these viewpoints together coherently? One of the difficulties here is the following. Proof is normally associated with assertions, and a constructive proof gives us specific information about *why* or even *how* a given assertion about

Words, Proofs, and Diagrams
D. Barker-Plummer, D. I. Beaver, J. van Benthem, P. Scotto di Luzio (eds.)

mathematical objects or programs is true. But a categorial derivation is not a proof in this sense: it tells us what an ordinary natural language sentence means, not that it is true. In other words, we have a nice technical convergence, but what does it mean?

Such questions can be raised more profitably in the focus of some empirical application. Let us think of *natural language*, much in evidence elsewhere in this Volume, as this is about the most challenging cognitive process to understand. The two faces of proof raise a larger issue. Natural language use involves a spectrum of activities, from subconscious parsing of a sentence to conscious discourse planning. Both involve 'reasoning' of sorts, but some of this is hard-wired (I cannot decide to change how I parse sentences), and some is in the software: I can change my habits of planning and reasoning – at least to some extent. What is the overall architecture of this phenomenon? Is there perhaps one base logic to it, say linear logic, which can be 'parameterized' to deal with different cognitive tasks, more or less conscious?

But even though they involve logical architecture of this sort, the papers in this part work with different designs. The tasks involved in natural language use can be conceptualized in ways that seem quite different from proof and derivation. Koller and Niehren rather think of them as *constraint satisfaction*. Processing spoken or written language involves accumulation of constraints, and making sense of it means finding a suitable model (say a syntax tree with all kinds of other information attached) that satisfies all these constraints. In the logician's terms, the task is now 'model construction', rather than 'proof search' – and these things can have quite different complexities.

It seems fair to say that there is no consensus on how to conceptualize the tasks involved in natural language understanding. To be sure, there certainly are technical analogies between the derivational and the constraint satisfaction approaches to parsing. De Paiva even suggests looking at proof search as a kind of 'model checking' in a category. But the systematic comparison between these two logical perspectives on natural language is far from complete. Now this diversity is no problem per se: it indicates a healthy stock of approaches, logical maintenance of which is as important as maintaining bio-diversity.

But it *would* be good to agree at least on a set of *common empirical challenges* in the field, that can be described without prejudging issues of theoretical framework preference. One example should surely be *ambiguity and underspecification*, as high-lighted by Koller & Niehren, because it is a logical problem that does not seem to hinder us in practice. Logicians have been good in showing things that were thought to be unproblematic *undecidable* or *intractable*. It would be nice to ask them for the converse: why don't your dire predictions come true in this case? It might become be one test of proposed analyses of 'logical form' how they propose to cope with this cognitive ease, or at least, offer a handle on its explanation. Another instance of the same kind is the evident *learnability* of natural languages, which is not very well-understood. It does not figure on the usual list of 'canonical concerns' that make a standard logical paper. But if one proposes a logical calculus with some descriptive pretention, an analysis of its learnability in some reasonable formal sense might be a standard agenda item, in addition to 'soundness' or 'completeness'.

But logic is not only about natural language. Grünwald's paper is interesting as a counterpoint, as it concerns another 'reasoning practice', viz. *statistical modeling and the predictions it generates*. In his case, the 'empirical' challenge is identified, and met head-on. Statistical models can be wrong, or at least hard to interpret—but people get by developing a sense for where they are good predictors, staying away from more problematic cases. He shows how one can make sense of this, and identifies mathematically where this practice is quite safe. These same issues occur in a broad spectrum of cognitive activities. We learn to use formal systems, treading lightly where they either do not make much sense, or threaten to involve us in endless computations. In other words, we can learn to exploit their strengths, while mostly avoiding their weaknesses.

This seems another challenge to designers of logic systems, which injects an entriely new element into the assessment of their behavior. How does linear or intuitionistic or classical logic perform in combination with modeling strategies for their input assertions? It may be that these questions are deeply tied up with *statistical properties*, of the sort which are only coming to light in current computational logic. We are now discovering emergent phenomena like 'phase transitions' in the behaviour of theorem provers or satisfiability checkers over random input of large size. A new world of 'bulk evidence' is coming into view, far removed from the traditional theorist's intuition's about minuscule phenomena, or practical questionaires for judgments by 'ordinary language users' about single sentences. Thus, significant logical phenomena may occur at different aggregation levels. Against this background, the broad juxtaposition of papers on the foundations of statistics, language, and computation in this part of our volume becomes natural, and suggestive.

Logic & Language

Editorial Introduction

David I. Beaver

Semantics ain't what it used to be. The past thirty years have witnessed both a broadening of the scope of linguistic semantics, and a deepening of our understanding of its mathematical basis. The broadening has meant that issues once taken to be other people's problems are now routinely discussed by semanticists. In general, these have been issues that in some way or other have appeared to challenge the Frege-Montague orthodoxy of compositional interpretation. Classic such issues include the analysis of idioms, questions and anaphora.

Regarding idioms, which are discussed at length in Dag Westerståhl's paper, the problem seems almost definitional: an idiom is by its nature non-compositional, it seems, so a simple claim that all language is compositional (meanings of wholes composed from meanings of parts according to predictable rules of combination) must be false. Twenty years ago, a majority of semanticists would probably have reacted to this with a shrug: idioms are lexicalized, and we need to make some trivial alteration to the definition of compositionality to allow for them. There are empirical difficulties with any particular slight alteration to the nature of compositionality, e.g. because, as Westerståhl discusses, idioms like "to pull strings" ($\equiv$ "to use powerful connections to get a result") are used very flexibly. However, Westerståhl bases his contribution on a deeper point: to assert that a particular expression (e.g. idiom) is non-compositional, is fundamentally wrong-headed. Rather, he argues, it is the compositionality of a complete grammar that must be assessed. Thus to understand the challenge posed by idioms, what is needed is a suitable definition of what it means for a grammar to be compositional. This is exactly what Westerståhl presents, utilizing an algebraic framework due to Hodges. What becomes clear is that a natural notion of compositionality, in and of itself, imposes shockingly few constraints on the relation between form and meaning.

Questions and anaphora provide the main empirical domains for the

Words, Proofs, and Diagrams
D. Barker-Plummer, D. I. Beaver, J. van Benthem, P. Scotto di Luzio (eds.)

contributions of Maria Aloni and myself, respectively. In both cases, interest in the phenomenon results from the contextual variability of interpretation.

Aloni analyzes the fact that the appropriateness of an answer to a question is dependent upon what is known by the conversational participants. In particular, the way in which we individuate the entities under discussion in a conversation will affect what is appropriate as an answer. Furthermore, according to the research paradigm initiated by Hamblin and Karttunen, the meaning of a question is given in terms of its answers, so contextual variability of the answerhood relation can be used as evidence for contextual variability of the meaning of the question. Aloni's solution is founded upon the use of what she terms *conceptual covers*, essentially ways of individuating objects. The use of this device will be of interest to anyone interested in the semantics and pragmatics of questions and answers. But it should also be of interest to anyone who has realized that the classic Kripkean analysis of names under metaphysical modalities does not extend to the epistemic domain. To those seeking knowledge (the posers of questions), there is no single absolute way of individuating entities, but rather an infinity of ways dependent on available knowledge and immediate goals. Conceptual covers provide a way of gaining temporary *rigidity* for individuals. Even in this uncertain world, perhaps we can approximate metaphysical truth.

In my own paper, I propose a framework based upon an extension of *dynamic semantics*. What makes a semantics dynamic is that it offers a representation of the information state of conversational participants as it evolves through a text. Dynamic semantics in its most general sense was initiated by work on anaphora and presupposition by Irene Heim, Hans Kamp and Lauri Karttunen in the seventies and early eighties. One can see the different lines of research in dynamic semantics as involving two main approaches. First there is the approach typified by Hans Kamp's Discourse Representation Theory (DRT). Second there is the Amsterdam approach, as initially developed by Jeroen Groenendijk and Martin Stokhof. The DRT approach is generous, using rich representations that hold most of the information that might conceivably be needed by any semantic or pragmatic process. On the other hand, the Amsterdam approach is more miserly, representations being kept initially much leaner, holding only information relevant to the particular phenomenon being analyzed.

The proposal I make is based on an enrichment of Amsterdam-style dynamic semantics. This enrichment consists of modeling uncertainty about the informational transition intended by the speaker. Where

Amsterdam-style work typically models meaning as a deterministic transition from information state to information state, my proposal introduces non-determinacy, such as is more usually found in DRT style analyses. I go on to suggest that various phenomena, including the resolution of anaphora, be analyzed in terms of competing pressures on which transition should be selected. Formally, these selection pressures are modeled as orderings over the set of possible transitions. If my proposals are on the right track, then a range of disparate phenomena, including presupposition accommodation, anaphora resolution and Gricean implicature, might eventually be analyzed in a common theoretical framework.

Many of the issues raised in the papers by Westerståhl, Aloni and myself come under discussion once more in Martin Stokhof's thinkpiece. This paper suggests that the edifice of modern semantics has been built with insufficient attention to well-known philosophical concerns, as if somehow semantics and philosophy of language were entirely independent enterprises. In particular, Stokhof finds that semantic theory does little to address the problems raised by Quinean Radical Interpretation, by the normativity of language, by our ability to introspect consciously on certain parts of the interpretation process, and by the *situatedness* of interpretation. Taking these issues seriously, Stokhof argues, should lead to a dramatic shift in the way interpretation and meaning are analyzed.

While Westerståhl shows us that the principle of compositionality hardly constrains the relation between form and literal meaning, Stokhof suggests that compositional interpretation can only very partially constrain what meaning is arrived at on a particular occasion of interpretation. Rather, contextual factors must play a more central role. To some extent, Stokhof's proposals are sympathetic to proposals like those of Aloni and myself. He certainly sees a place in semantic theory for sophisticated notions of context, such as introduced by Aloni, and he certainly argues for a view of interpretation in which pragmatic factors are given great currency, as in my own model. However, Stokhof's observations are, by his own account, problematic for any theory of meaning currently under consideration, even for the Dynamic Semantic paradigm which he has helped instigate. (He views Dynamic Semantics, not surprisingly, as a step in the right direction, but little more.) At risk of collectively playing the role of Voltaire's Pangloss, we editors take a positive moral from Stokhof's philosophical diatribe: in linguistic semantics, the best is yet to come!