International Series in Modern Applied
Mathematics and Computer Science
Volume 5

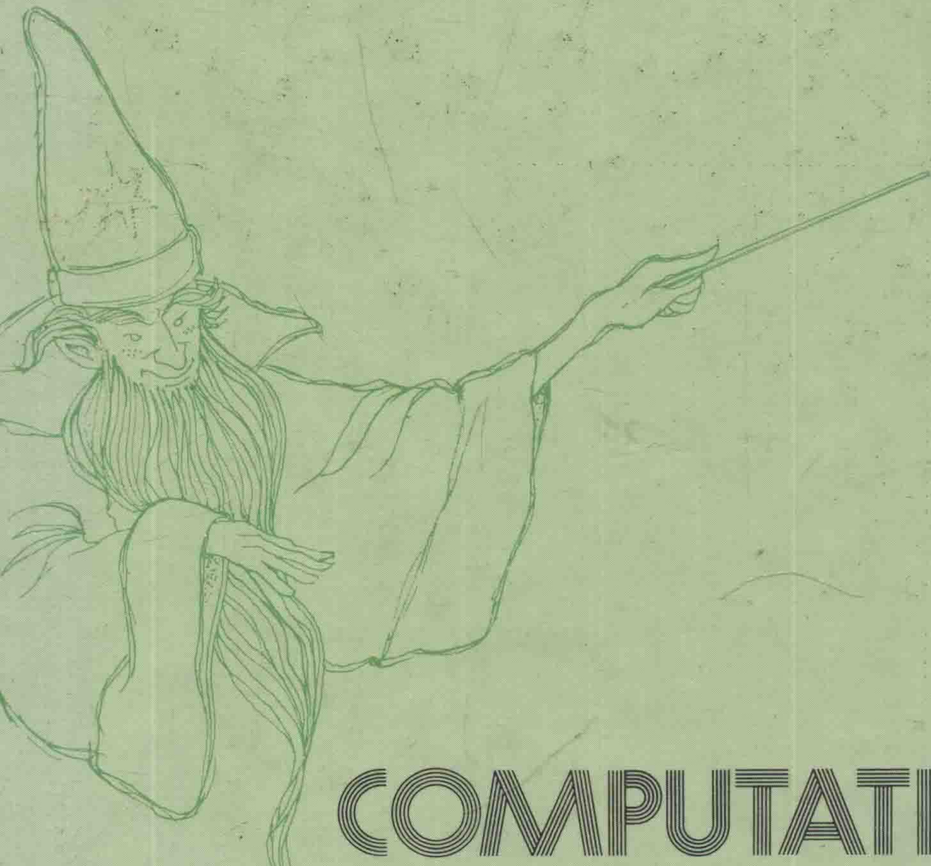

COMPUTATIONAL LINGUISTICS

Editor
NICK J. CERCONE
Simon Fraser University, Canada

PERGAMON PRESS

COMPUTATIONAL LINGUISTICS

Edited by

NICK CERCONE
Computing Science Department, Simon Fraser University, Burnaby, British Columbia, Canada V5A 1S6

PERGAMON PRESS
OXFORD · NEW YORK · TORONTO · SYDNEY · PARIS · FRANKFURT

U.K.	Pergamon Press Ltd., Headington Hill Hall, Oxford OX3 0BW, England
U.S.A.	Pergamon Press Inc., Maxwell House, Fairview Park, Elmsford, New York 10523, U.S.A.
CANADA	Pergamon Press Canada Ltd., Suite 104, 150 Consumers Road, Willowdale, Ontario M2J 1P9, Canada
AUSTRALIA	Pergamon Press (Aust.) Pty. Ltd., P.O. Box 544, Potts Point, N.S.W. 2011, Australia
FRANCE	Pergamon Press SARL, 24 rue des Ecoles, 75240 Paris, Cedex 05, France
FEDERAL REPUBLIC OF GERMANY	Pergamon Press GmbH, Hammerweg 6, D-6242 Kronberg-Taunus, Federal Republic of Germany

First edition 1983

Library of Congress Cataloging in Publication Data

Main entry under title:
Computational linguistics.
(International series in modern applied mathematics
and computer science; v. 5)
"First published in Computers and Mathematics with
Applications, volume 9, number 1"—Verso of t.p.
1. Linguistics—Data processing—Addresses, essays,
lectures. I. Cercone, Nick J. II. Series.
P98.C612 1983 410'.28'54 83-8039

British Library Cataloguing in Publication Data

Computational linguistics.—(International series in
modern applied mathematics and computer science;
v5)
1. Linguistics—Data processing
I. Cercone, Nick J. II. Series
410'.28'54 P98

ISBN 0-08-030253-X

Published as a special issue of the journal *Computers and Mathematics with Applications,* Volume 9, Number 1 and supplied to subscribers as part of their normal subscription. Also available to non-subscribers.

Printed in Great Britain by A. Wheaton & Co. Ltd., Exeter

COMPUTATIONAL LINGUISTICS

International Series in
MODERN APPLIED MATHEMATICS AND COMPUTER SCIENCE
General Editor E. Y. RODIN
Volume 5

PREFACE

This special collection of papers represents a relatively comprehensive overview of the variety of important research in computational linguistics currently taking place in North America. The primary contributors to this edition include computing scientists, philosophers, psychologists and linguists; their professions are indicative of the disparate approaches which each has brought to further research in computational linguistics.

Twenty active and well-known researchers in computational linguistics were invited to contribute to this volume, but unfortunately not all were able to do so; the contents of this volume are not as complete as it might otherwise be. The fifteen papers representing nineteen authors' contributions cover a breadth of computational linguistics consisting of:

(1) Theoretical foundations: the logical foundations of knowledge representation; semantic analyses; and model-theoretical semantics.
(2) Parsing: parsing strategies for natural language; computational aspects of parsing; perspectives on parsing issues.
(3) Discourse processing: psychological and linguistic modelling; discourse analysis.
(4) Text analysis: text and content analysis; text generation.
(5) Natural language understanding and knowledge organisation: memory models; learning; inference techniques.
(6) Programming systems for computational linguistics: knowledge representation languages; special purpose languages.
(7) Programming environments: programming considerations for computational linguistics.
(8) Interactive applications: natural language front-end processors to database systems; the human factors interface.

The narrow approaches to machine translation of the early 1960s pale when compared to the considerable assortment of methodologies available to the modern computational linguist. The growth in the number of publications devoted to computational linguistics parallels a similar increase in computing science literature and is indicative of its rapid development. This impressive maturation has been accompanied by an equally exciting change in the nature of experiments, systems and theoretical speculation. Only a decade or so ago researchers were content to speculate about the results of a program demonstrating limited comprehension in a "micro-world". Contemporary results indicate a broader framework for investigations into the theory and applications of research in computational linguistics.

The challenges of integrating the various approaches to problems faced by computational linguistics were difficult and, at times, frustrating. I owe each contributor a debt of gratitude, first for writing a new and totally original article for this collection and also for rewriting, editing, and reviewing. Each author constructively and patiently reviewed three other authors' submissions which has helped to improve the contents, coherence, and style of the entire issue. I am immeasurably indebted to Ms. Josie Backhouse for her extraordinary organisational efforts and her attention to every aspect concerned with this special collection. The entire volume represents over a year of hard work which significantly extended my original plans for publication. Hopefully, the extra time was well-spent and will make this issue a classic reference for the future.

INTRODUCTION

COMPUTATIONAL LINGUISTICS research should develop a general theory of natural language understanding as a foundation for computer programs which understand natural language. Any theory of natural language understanding must account for the representation, organisation and subsequent utilisation of knowledge (e.g. for making plausible inferences, associating meaningful pieces of discourse, etc.).

Active and well-known researchers in computational linguistics were invited to contribute such that each of the following topics would be given at least two, often disparate, viewpoints: (1) theoretical foundations for computational linguistics; (2) parsing strategies for natural language; (3) modelling of discourse processes; (4) text analysis and generation; (5) natural language understanding systems and knowledge organisation; (6) programming systems for computational linguistics; (7) programming environments and considerations for computational linguists; and (8) interactive applications such as natural language front-ends to database systems. The papers as finalised represent a great diversity of topics, methods, and approaches. Philosophers, linguists, engineers, and computing scientists contributed to this volume divided roughly evenly between Canadian and American institutions.

Theoretically oriented papers include those discussing the adequacy of representational formalisms and their interpretation in natural language understanding computer programs. Most pragmatically oriented papers apply novel techniques to existing systems to enhance their "naturalness".

David Israel postulates that the underlying argument between the logical formulae and the semantic network formalism proponents is not merely one of their precision and computer-interpretable properties as a mathematical notation but of a much wider context. Dr. Israel sketches semantic accounts for at least two kinds of semantic network formalisms. The first account is based on the notion of inheritance, the other is not. Dr. Israel maintains that a critical condition of [representational] adequacy is fidelity to some of the intuitions of the creators of the formalisms.

In an earlier paper Ray Reiter proposed a logic for default reasoning to provide a representation for common sense facts (among other things). In this paper Dr. Reiter and his collaborator Giovanni Criscuolo address some of the representational issues in default reasoning, particularly when anomalous default assumptions are derived. The non-normal default rules required to deal with default interactions lead to a new concept of integrity, distinct from the conventional integrity issues of first order data bases.

John Sowa describes how to generate language from conceptual graphs, a semantic representation that has a direct mapping to natural language. He presents a universal algorithm for scanning the graphs, together with a version of augmented phrase structure grammar for specifying the syntax of particular languages. Since the graphs may allow the generation of multiple surface structures, Dr. Sowa combines them with phrase structure rules to enforce context-sensitive conditions.

Practical results in information retrieval and automatic translation have recently been achieved for naturally-occurring texts in certain narrow technical areas. Richard Kittredge discusses the semantic processing of texts in restricted sublanguages. By way of illustration he outlines a procedure for processing stock market reports into a predicate-argument representation of their content and discusses potential applications of the procedure beyond information retrieval.

Professors Lockman and Klappholz discuss the problems of resolving ambiguities when understanding natural languages. To satisfy this requirement complex inferencing from a large database of world knowledge must take place. Critical to this task, and one for which humans adapt quite readily, is the control of such inferences. Their paper discusses the problems involved in such control of inferencing and an approach to their solution is presented based on determining where each successive sentence "fits" into the text as a whole.

Human conversational participants depend upon the ability of their partners to recognise their intentions, so that their partners may respond appropriately. In such interactions, the speaker encodes his intentions about the hearer's response in a variety of sentence types. Instead of telling the hearer what to do, the speaker may just state his goals, and expect a response that meets these goals at least part way. Candace Sidner presents a new model for recognising the speaker's intended meaning in determining a response. The model shows that this recognition makes use of the speaker's plan, his beliefs about the domain and about the hearer's relevant capacities.

J. G. Meunier and F. Lepage explore new paths for formal semantic and computer text processing of large non-preedited natural language texts. They describe the traditional approaches to this problem, then discuss certain semantic aspects in computer text processing. They use a model theoretic approach embedded in an algebraic language and follow the hypothesis: discourse in a text constitutes a semantic space built of an ordered set of sentences which are of different logical types and which present a specific pattern of coherence expressible in a syntactic manner.

James Allen gives a system overview of ARGOT, a long-term research project that has the ultimate aim of describing a mechanism that can partake in an extended English dialogue on some reasonably well specified range of topics. Dr. Allen describes progress made towards this goal and outlines the current research in which the project is focussed. The underlying current theory and system built according to that theory is outlined and deficiencies in the system presented. The new system, ARGOT, under development is then described.

The implications for natural language generation of description directed control are examined by David McDonald. In description directed control the controlling data structure is the surface-level linguistic description of the very text being generated. This constituent structure tree is itself generated depth-first by the incremental realisation of a hierarchical description of the speaker's communicative goals organised according to the scope and importance of its components. The process of traversing the surface structure gates and constrains the realisation process; all realisations are thus subject to the grammatical constraints that accrue to the surface structure at which they occur, as defined by the grammatical annotation of the surface-structure tree.

Gerald DeJong and David Waltz outline concerns which must be addressed if a machine is to cope with understanding novel languages. They state that systems which operate primarily by pattern matching are less interesting than systems which have general rules which can be used to generate a meaning representation for unanticipated inputs. They discuss a wide variety of types of unanticipated input and give many examples of these types.

The computational approach to fuzzy quantifiers which Professor Zadeh describes in his paper may be viewed as a derivative of fuzzy logic and test-score semantics. In this semantics, the meaning of a semantic entity is represented as a procedure which tests, scores and aggregates the elastic constraints which are induced by the entity in question.

William Havens is concerned with generalising formal recognition methods from parsing theory to schemata knowledge representations. The notion of a schemata as a suitable representation for a variety of artificial intelligence tasks is discussed and a number of contemporary problems with current schemata based recognition systems are presented. Professor Havens shows how to integrate top-down and bottom-up search in schemata representations.

The viewpoint espoused by Gordon McCalla in his paper is that natural language understanding and production is the action of a highly integrated domain-specific specialist. He first describes an object-oriented representation scheme which allows these specialists to be built and then the organisation of these specialists into a four-level goal hierarchy that enables the modelling of natural language conversion. Six specific kinds of recall tasks are outlined in terms of these structures and several dialogues examined.

Nick Cercone, Max Krause and John Boates present a set of tools for the computer lexicographer. They have devised three alternative algorithms which produce minimal and almost minimal perfect hash functions for table sizes suitable for functional computer lexicons. A semi-interactive system has been built based on one algorithm and it has been used to construct natural language lexicons whose size is over 60,000 words.

A natural language interface which gives extended natural language database interactions is presented by Bonnie Webber, Aravind Joshi, Eric Mays and Kathleen McKeown. They discuss two complementary directions for extending natural language interfaces to databases: (1) broadening the range of query types that can be handled; and (2) extending the range of responses that can be provided.

It is my fervent hope that this collection of fifteen papers will both advance knowledge and broaden the experience of present-day practitioners of computational linguistics. For those researchers with only a passing interest in computational linguistics and artificial intelligence this collection should serve as an introduction to the diversity and depth of work in the field.

Vancouver, BC. April, 1983 Nick Cercone

CONTENTS

Comp. & Maths. with Appls. Vol. 9, No. 1, pp. 1–13, 1983
Printed in Great Britain.
0097–4943/83/010001–13$03.00/0
Pergamon Press Ltd.

INTERPRETING NETWORK FORMALISMS†

DAVID J. ISRAEL
Bolt Beranek and Newman Inc., Cambridge, MA 02238, U.S.A.

Abstract—In a recent paper, Reiter and Criscuolo[3] remark "that (semantic) networks are notational variants of logical formulae is by now a truism in Artificial Intelligence circles". Shamelessly exploiting the foregoing quote as a pretext, I attempt to sketch adequate semantic accounts for at least two (kinds of) semantic network formalisms; one, based on the notion of inheritance, one, not. A crucial condition of adequacy to be satisfied is fidelity to some of the intuitions of the creators of the formalisms.

1. IMPRIMATUR

Whatever else we decide to include under A.I., I would like to join those who claim for it the task of developing, maintaining, and using computer-interpretable formalisms for represesenting knowledge.... One might view some of the representational formalisms being studied in A.I. as attempts to create languages that possess the precision and computer-interpretable properties of a mathematical notation, but do so for a much wider range of concepts than those dealt with by classical mathematics[1].

2. DOCTRINAL PREAMBLE

One often hears that modal (or some other) logic is pointless because it can be translated into some simpler language in a first-order way. Take no notice of such arguments. There is no weight to the claim that the original system must therefore be replaced by the new one. What is essential is to single out important concepts and to investigate their properties. The fact that the real numbers can be defined in terms of sets is no argument for being interested in *arbitrary* sets. One must look among the sets for the significant ones and cannot be censured if one finds the intrinsic properties of the reals more interesting than any of their formulations in set theory. Of course if we can argue that set theory provides other significant concepts, then we may find some reason for going beyond the real numbers (and it is not hard to find the reasons!). But clearly this discussion cannot proceed on purely formal grounds alone[2].

AMEN!

3. INTRODUCTION

In a recent paper, Reiter and Criscuolo remark that "the fact that networks are notational variants of logical formulae is by now a truism in Artificial Intelligence circles"[3]. Let us put aside the empirical, sociological claim—about which I am more than willing to defer to Messrs. Reiter and Criscuolo. Let us look rather at the content of the truism itself.

When we do so, we notice a certain ambiguity. Perhaps, Reiter and Criscuolo are supposing that there is at least one logical language for whose formulae the "sentential analogues" of each and every semantic network formalism are notational variants. Or, switching the quantifiers, perhaps the claim is that for every semantic network formalism there is at least one logical language of which it is a notational variant. Now these are two very different claims. (One could easily imagine a third, stronger claim, being presupposed, viz. that all semantic network schemes are notational variants of, in particular, the language and logic of classical, first-order quantification theory. I will, in what follows, address myself to this strongest, most specific, view as well.)

I am not going to speculate on which of the three views Reiter and Criscuolo really had in mind. Indeed, the quote from Reiter and Criscuolo is the merest pretext for the present paper. I intend to sketch quite different semantic accounts for two different (kinds of) semantic network formalisms, one organized around the notion of inheritance, the other, not; the latter, keyed to a certain family of intensional contexts, the former, not. Most important, *the semantic accounts will differ* and, with luck, *the differences will reflect differences in the central intuitions of the*

†This research was supported by the Office of Navy Research under Contract No. N00014-77-C-0371.

semantic net theorists concerned. The aim is, simply, to try to take semantic net theorists at their word; and to show that, in doing so, one must range fairly far and wide beyond the confines of standard first-order logic. I am not going to argue for adopting any particular semantic net formalism; nor am I going to examine any of the intuitions motivating the work to be discussed.

To make clear my intentions, let me note two ways in which one might fail to take the "semantic(s)" in "semantic network" seriously. The first is embodied in the work of Quillian, Collins *et al.*, work which originated the semantic network tradition. By my lights, the structures described in this work are not intended to be languages; rather they are part of a theory or model of a certain range of memory-related psychological phenomena. The nodes in the network might be words, or even sets of sentences, but the accounts of the network are not semantic accounts; they do not constitute a semantic theory. There is no attempt to account for the meaningfulness, or describe the meaning, of the nodes. This, of course, is no criticism of the work; nor am I suggesting that the researchers in question were confused about the present point. So much for the first way of not taking the "semantics" of semantic network formalisms seriously.

As a second way of not taking the semantics of semantic net formalisms seriously, I have in mind the following doctrine: the only way to interpret semantic network formalisms semantically, *no matter what the semantic network theorists may say*, is to treat them as notational variants of standard first-order languages, with their standard "Tarski-style" semantics. This might best be described as a way of not taking semantic network *theorists* seriously. It is against this view, in particular, that I mean to deploy the quotation from Dana Scott with which I started. Of course, for all I know, no one holds this view; in which case, I am arguing only against phantoms of my own fevered imagination. So much the better. I, at any rate, want to try to take some semantic net theorists at their very word. Let the appropriate semantic account fall out however it may.

4. SNePS

I want first to discuss a system in which the notion of inheritance plays no (special) role. It would also be nice to have a case in which the theorists are both explicit about their motivating intuitions and diligent in presenting sufficient detail on which to hang semantic speculation. Sad to say, this narrows the field down quite a bit. Still, there are at least two choices; the (atypically nameless) system of Schubert, Cercone and Goebel[4] and Shapiro's SNePS[5, 6]. I have decided to examine the latter; and this, for two reasons. First, Shapiro (and Shapiro and Maida[7, 8]) presents a fairly explicit "philosophy of semantic networks", as well as an enormous body of detailed description of the workings of the system. I think the philosophy is widely shared and trying to account for the considerations operative in its formulation raises interesting problems. Second, Schubert *et al.* are too explicit for my purposes. It is quite clear that they view their formalism as a Montague-style type-theoretic, intensional system. It would, I think, be illuminating to work out in detail a semantic account for the system described in [4]; but there can anyway be no doubt about the "logical space" within which we would find ourselves. This is not the case, or so I shall claim, for SNePS. A semantic account appropriate to *it* will force us to wander into largely, though not completely unexplored, territory. (And this, I also claim, is no argument against SNePS. Remember the wisdom of the great Scott.)

Shapiro has argued that we should impose the following conditions on semantic networks:

(1) Each node represents a unique concept.

(2) Each concept represented in the network is represented by a node.

(3) Each concept represented in the network is represented by a unique node (the Uniqueness Principle).

(4) Arcs represent non-conceptual (logical?—*D.I.*) binary relations between nodes[7, 8].

In what follows I shall take Shapiro to be simply describing SNePS, thus ignoring his arguments for imposing these conditions on *all* semantic networks. I shall also (largely) ignore the fourth condition and, for that matter, a fifth: "the knowledge represented about each concept is represented by the structure of the entire network connected to the node representing that concept."

On the basis of these conditions, Shapiro (and Shapiro and Maida) contend that "all nodes of a semantic network represent only intentions"[7]. Again, I shall take this as describing SNePS, not as prescribing for all semantic networks.

Now what are we to make of this? Raising the question this way raised the issue as to whether my intensions toward SNePS are honorable. I hope they are. Appearances to the contrary notwithstanding, I am not singling out a few lines, ripped out of context, for malicious attention. First, no malice is intended and I trust no harm is done. Second, and more important, my reconstruction of SNePS *attempts* to embrace a large number of the claims and arguments in the texts. I will not attempt to support my claim in this respect by citation. I hope that any one who has read the papers will agree that I have presented at least *one* way of construing them—not the only way, and perhaps not the way favored by the authors. As for those who have not read the material, I fear they shall have to take me at *my* (immediately preceding) word.

Let's remind ourselves of the project. We are to find *a* language-cum-semantics which can reasonably be taken to be that formal system, formulae of which correspond to the sentential pieces of SNePS. What, for instance, does the logical vocabulary of our target language consist of? Here we get help for [6, 9]. For our purposes, what's crucial is that the logical constants mentioned are generalizations of the familiar truth-functional connectives and quantifiers. But what of the quantifiers; over what do the bound variables range; what kinds of things are assigned to the variables?

We might as well start at the beginning and specify how atomic predicational formulae are to be understood. The standard way of handling intensions in contemporary logic is to treat them as functions from an index set of contexts or possible worlds into some set-theoretic construct on the domain of the model structure. So, to take an important instance, properties—the intensions associated with monadic predicates—are explicated as functions from the index set into subsets of the domain of possible individuals. We might, then, try imagining a model structure consisting of a domain *D* of possible† entities, a non-empty set *I* (the index set or the set of possible worlds), and (optionally) a distinguished element *w* of *I* to represent the real world. Now define an *individual concept* in such a model structure as a function *ic* from *I* into *D*. (Total or partial? The traditional answer has always been total; but it is not clear what answer is appropriate to SNePS.)‡

What, then, can we say about individual terms, about individual variables, individual constants, and definite descriptions of individuals? Given what Shapiro says, it is hard to see how there is any alternative to a uniform intensional treatment. In specifying models for SNePS, all such terms (including, nota bene, individual variables) get interpreted by being associated with individual concepts (not individuals, not members of *D*). A model for SNePS will associate with each individual constant an individual concept (a member of the set of functions from *I* into *D*) and assignments relative to such a model will do the same for individual variables. This means that the modal language-cum-logic is not of the standard variety. I said that I would be assuming the general framework of Kripke–Montague style model-theoretic accounts; however, neither Kripke nor Montague propose semantic accounts in which the individual variables get assigned individual concepts. Dana Scott's advice that one opt for just such a uniform treatment of all individual terms[2], which, as it happens Shapiro seems to be following, has been followed by just about no one else except Aldo Bressan in "A General Interpreted Modal Calculus"[13].

Enough about individual terms; how shall we handle predicate letters? Remember: all nodes are intensional. Individual terms are associated with individual concepts, not with possible individuals; so we can not, in good conscience, assign (e.g.) to one-place predicate letters functions from *I* to subsets of *D*. That is, we can't assign to one-place predicate letters sets of possible individuals. The obvious move *might seem to be* to associate sets of individual concepts with monadic predicate letters; but this does not render "predicate-nodes" truly intensional. The extension of a predicate at a world is no doubt a set of individual concepts; but

†Or to be faithful to Shapiro; conceivable.

‡As we shall see, there are reasons for thinking this effort to reconstrue SNePS-style intentions in terms of Kripke–Montague model-theoretic treatments of intensional contexts slightly misguided. I shall suppress them till the end of my discussion of SNePS. As for such model-theoretic treatments themselves, the classic sources are [10–12].

what is its intension? Surely, it is (as both Bressan and Scott insist) a function from I into the power set of the set of individual concepts. Or what comes to the same thing, a function from individual concepts into propositions, where these are functions from I into $\{T, F\}$. (Shapiro is explicitly committed to "propositional nodes"—nodes for "concepts of the TRUE" and "concepts of the FALSE"[7, 8].) So, predication is an intensional functor.

By the way, this does make it a little hard to understand what Shapiro says about MEMBER and CLASS arcs, and the relation between these and ISA arcs[6]. Arcs represent "structural relations", which, given the examples, must mean binary logical functors. So ISA links represent the predication functor; but this functor must be intensional, i.e. at each world i, the truth value of "Fx" (x is an F) is not a function solely of the extension of "x" in i. In general, the relation between "Fx" and "x is a member of the set $\{y: Fy\}$" is complicated. In particular, it is relatively straightforward only for extensional predicates, predicates which informally meet the following condition: they apply to a given *ic* x (at i) iff they apply to any d such that $d = x(i)$. The truth value at a world of a sentence predicating an extensional predicate to an *ic* does depend only on the extension of that *ic* at that world.†

Skipping lots of nasty details, we are now in a position to wave our hands, with some confidence, over the first order quantifiers. But note the "first-order" in the foregoing. Shapiro explicitly mentions the availability of higher-order constructs in SNePS. Thus, he says we can have nodes representing the second-order concepts of a property being extensional and of a property being intensional[7]. It's not clear how high we can go in this vein, and for reasons of space I herewith demur.‡

Finally, there is the problem of propositional concepts. There are problems, in particular, about sentences embedded in intensional functors, such as "Necessarily . . . " or "S believes that..." About the first, Shapiro doesn't have much to say. Later, I shall suggest that the reason why he doesn't, throws light on the fact that many of the choices we've made in giving a semantic account of SNePS seem ill-motivated by the texts.

The only intensional contexts Shapiro (and Shapiro and Maida) discusses are propositional attitude contexts, those involving verbs such as "know", "believe", etc. These contexts are treated, moreover, as relational; i.e., "know", "believe" (taken in their sentential complement mode) are treated, not as intensional operators, but as two-place predicates whose relata are individual concepts of subjects and either propositions or propositional concepts.§ None of this tells us very much about how they would treat the standard modalities. Why no mention of these modalities, the hard-core intensional operators?

The answer to the mystery of the missing modalities is to be found, I think, in the intuitions behind the Principle of Uniqueness. The crucial motivation is the view that the "nodes [of a semantic network] represent the concepts and beliefs of a thinking, reasoning, language using being (e.g. a human)"[7]. Hence the centrality of the propositional attitudes. But these generate contexts which are arguably not just intensional but hyperintensional (and perhaps to the nth degree). (The slightly garish term is due to Cresswell[14].) Hyperintensional contexts are those in which substitution of logically equivalent sentences or strictly identical terms is not guaranteed to preserve truth. Thus, to take a particularly startling case, it is arguable that from the fact that S believes that P and Q it does not follow that S believes that Q and P. Again, from that the fact that S believes that 4 is even, it does not follow that S believes that the square of 2 is even. (Or that 4 is not odd.) Now Shapiro and Maida certainly seem to view belief contexts as hyperintensional and it is this which gives sense to the Principle of Uniqueness: no two distinct nodes (represent expressions which) are intersubstitutable *salve veritate* in all the contexts the language can generate. That is, every distinction between nodes made purely syntactically by the language is a semantically significant distinction—there is some context which semantically "separates" any two distinct expressions. Thus Shapiro: "No two distinct nodes represent truly equal

†Needless to say, identity between elements of D is world-relative; the primitive notion is strict identity between *ic*'s, i.e. co-extensiveness at all i. For simplicity, we assume totality.

‡I should note, though, that Bressan's is an omega-type intentional system, as are some of the Montague logics.

§As a reminder, a proposition—the intension of a sentence—is a function from I into $\{T, F\}$. A propositional concept would be an individual concept whose extension at a world is a proposition; it is a function from I into functions from I into $\{T, F\}$.

concepts"[7].† Given this, it is clear enough why they would want to treat belief-contexts relationally; for there can be no question of a "logic of belief"—there are no laws to govern the behavior of a logical functor for belief.‡ And, given the very same, it is clear enough why so little is said about the standard modalities, which are merely (not hyper-) intensional.

There is no very happy account of hyperintensionality from within the model-theoretic framework (or elsewhere); although there are attempts. A move first proposed by Carnap in [17] is to specify a finer-grained notion than logical equivalence, to that of *intensional isomorphism*, in terms of the construction trees of complex expressions and the intensions assigned the constituent expressions. For example, take two sentences which are purely truth-functional tautologies, hence which are true in the same (namely, all) possible worlds. They might, however, be formed out of different constituents and these might have different intensions, etc. This move might or might not handle the case of sentential conjunction raised above; and, of course, one most certainly has the option of stipulating that if someone believes that P & Q, s/he believes that Q & P.§

To return to the central point: if one is focussing on propositional attitude, allegedly hyperintensional contexts, it can seem like a waste of time to introduce model-theoretic accounts of intensionality at all. Thus the air of desperation about the foregoing attempt to use such an account (albeit a non-standard one) to explicate a semantic net formalism that is focussed on the propositional attitudes.

(More than) enough has been said, I hope, to support my claim that in taking SNePS seriously, in particular in attempting to present a semantic account which honors some of the intuitions of its creator(s), one is led into rather interesting, if slightly forbidding, logical terrain. Of course, enough has also been said to prove beyond a shadow of a doubt that there are significant open problems to be solved before a fully adequate account can be given; the major one being that of providing a model-theoretic account of propositional attitude contexts which is both formally impeccable and at least a little plausible.¶

5. INHERITANCE

I want now to take a look at those semantic network formalisms in which the notion of inheritance is central. Here, too, the work is motivated by a certain "family" of intuitions. I would not want to try to ascribe these intuitions to any one in particular and perhaps no one believes all of them. I will state them baldly and without comment. (If you'd like, you may imagine me to be making up a position out of whole cloth and then, perversely, imposing upon myself the duty of making formal sense out of it.)

(1) The graph-theoretic nature of semantic networks counts for something above and beyond "ease" of formulating and implementing access and retrieval algorithms.

(2) There is something in principle wrong about the way in which standard semantic accounts separate the language from "theories" expressed in the language—not enough is fixed by the specification of the language itself. In some sense, different languages implicate different theories.

(3) Somehow the central role in thought (and language) of kind terms, in particular of natural kind terms, must be captured—and to do this, one must take seriously the fact that natural kinds come in families on which hierarchical, taxonomic relations are defined.

These three "intuitions" coalesce to form a certain perspective on semantic nets.‖ This

†Note: it is one thing to require that no two primitive terms are co-intensional; it is quite another to argue that no two terms—primitive or not—are such. This latter, though, seems to be the Shapiro–Maida position.

‡It is not necessary to hold this view of the hyperintensionality of propositional attitude contexts to dissuade one from the logical operator position; see Montague[15, 16]. But it sure as heck is sufficient.

§Doubts have been raised about the efficacy of such a move, especially with respect to the iteration of propositional attitudes over different subjects. In particular, try Mates's matrix: Let D and D' be two intensionally isomorphic sentences. Then the following are also intensionally isomorphic: (a) Whoever believes (that) D, believes (that) D. (b) Whoever believes (that) D believes (that) D'. But nobody wonders whether anybody doubts that whoever believes D believes D; but some philosopher may well wonder whether anybody doubts that whoever believes D believes D'[18].

¶There have, of course, been significant attempts in this direction; in the A.I. literature the outstanding candidates are [19] and [20]. Needless to say, there is much, much more that needs to be said; sad to say, it will here go unsaid.

‖Actually, I do not in this paper take natural kinds as seriously as one might; in particular, not as seriously as I do in [21].

perspective often carries along with it a commitment to some theory or other of prototypes and this commitment is often understood, mistakenly, as constituting an essential constraint on any semantic account of such a semantic net formalism. (For more on this, see [22].) In what follows, I shall completely ignore all issues about prototypes.†

Some preliminary comment on (2) is called for. I have often thought that there was a systematic confusion of language and theory evident in the work on semantic networks. One was never given a specification of the "language" neat; rather, what one got were notations for particular sentences and one was supposed to be able to go on from there. But how is one to know how to go on, unless one knows, at the very least, what is fixed and structural, as against what is subject to variation by way (e.g.) of varying meaning assignments?‡ Despite (or because of) my worries about this confusion, I think it worthwhile to attempt a rational reconstruction of the intuition that a partial theory of the world is "directly" embedded in the languages we use in thinking and speaking about that world.

The best way to shed a little light on the second of our central intuitions is by contrast. The standard mode of specifying a quantification language (or language scheme) includes the specification of a typically infinite set of typically infinite sets of predicate letters of all possible arities. To interpret such a language in the classical model-theoretic way, one assigns a set as the domain for the variables and an extension to each predicate letter. In the standard vein, these assignments are to subsets of the n-place Cartesian product of the domain of the variables (for n-place predicates).§

The predicates, individually, are syntactically unstructured and, collectively, are no more than members of various unordered sets. That is to say, from a semantic point of view, there are no constraints on the interpretations assigned to the members of any set of predicate letters beyond that imposed by the arity of the individual predicates. (So, e.g. two-place predicates must be assigned subsets of the set of all ordered pairs of members of the domain, etc.) In fact, though, there is a kind of constraint: there are to be no semantic interdependencies among the extensions of the predicates. That the assignments to predicate letters be independent of one another mandates that the specification of the language cannot impose any relationships among the extensions of the predicates. Hence, the standard scheme imposes a requirement of logical independence among atomic sentences.

For instance, if one has two one-place predicate letters, "P" and "Q", one would not (indeed, on many accounts, should not) assign to "Q" the complement of the set assigned to "P"; that's a job for the negation operator. One can, of course, specify a theory in the language which has "$(x)(Qx <-> -Px)$" as a theorem; but one can also formulate in the same language a theory according to which "Q" and "P" are coextensive.¶

The picture sketched above should be familiar, but it does not seem to fit very well that body of work by researchers in Artificial Intelligence which focuses on the related notions of taxonomic structures and inheritance. I shall now present a mildly non-standard account in which those notions are indeed central.

6. THE SEMANTICS OF INHERITANCE

You may think of the sentences of the language as looking a lot like the sentences of a standard first-order language, with a few wrinkles, of course. There will be a finite number of

†As an aside, the semantic net formalism I have most clearly and fully in mind is KL-ONE, which likewise eschews prototypes (see [23, 24]).

‡Let me remind you of the quotation from [7]: "A semantic network models the knowledge structure of a thinking, reasoning, language using being. In this case, nodes represent the concepts and beliefs such a being would have." This certainly sounds like an identification of a given network with a particular creature's view of its world.

§A point of terminology: logician-types divide on usage here. I tend to speak of one and the same language being susceptible to many different meaning assignments (interpretations). Others speak, rather, of different—applied—languages as instances of the same language scheme. Anything said in one mode can be translated into the other. I leave it to the reader to make the requisite transformations.

¶We couldn't, in the foregoing, talk of the co-extensiveness of "Q" and "not-P"; at least not if we were restricting ourselves to the resources available in standard first-order languages. For complex predicates are not among those resources; so, in *strictu sensu*, there is no predicate available to play the role—solely in virtue of its structure or solely in virtue of the specification of the language and the standard semantic account—of the complement of another predicate. Rather, we must make do with (complex) open sentences.

primitive predicates of not all possible degrees.† The crucial feature of the account resides in the requirement that integral to specifying the semantics of the language is the specification of an algebraic structure of properties (concepts, intensions) by which the assignment of extensions to the predicates is constrained.

One more preliminary point: the elements of the algebraic structures are to be understood as properties. The "linguistic" representatives of these are, in the first instance, lambda-abstracts interpreted as singular terms denoting properties. These are then associated with monadic predicate letters of the language. When we get around to exploiting lambda-abstraction as a complex predicate-forming operator, predicates of the language will look a lot like the singular terms for properties with which they are correlated. This is unfortunate and *could* easily be remedied by choosing a different notation for the complex predicate forming operator. (But it won't be so remedied.) Occasional reminders of the distinction between singular terms for properties, on the one hand, and predicates, on the other, will be sprinkled about.

In the first instance, we shall limit ourselves to properties, properly so-called; i.e. syntactically speaking, to monadic predicates. Even here choices arise. First, is there one most general, all inclusive property? From the graph-theoretic point of view, is the structure a tree, or is it rather a forest (an unrooted or multi-rooted tree)? Second, are there cases in which a primitive property is immediately included in more than one primitive property—has more than one immediate ancestor in the structure? (Are there cases of multiple inheritance among the primitive properties?) This is the question, from a graph-theoretic point of view, as to whether the structure is a tree (rooted or not) or an upper semi-lattice (perhaps rootless).

I need not make a decision on these points; such a choice is up to a user of the scheme I am describing, and his/her decision, in turn, depends on the structure of the domain of application and/or his/her conceptualization of that domain. I should, however, be able to show how an account would be given in each case. So I shall begin with the simplest case, that in which the structure is an honest-to-goodness tree.

The ordering relation which generates the algebraic structure is property inclusion (or porperty entailment), taken as primitive. There is another significant semantic relation: the relation among the immediate descendants of a given node, that is, among siblings. Such properties are taken to mutually exclude one another. (One can then define property independence in terms of inclusion and exclusion.) The intuition, here, is as follows: imagine a portion of the structure which begins with a node for the property of being a mammal. This has, say, 10 immediate descendants marked as mutually exclusive, among which are, e.g. the property of being a cat, the property of being a pig, etc. Each includes the property of being a mammal, and each excludes all the others. Crucially, there is no requirement that the sequence of mutually exclusive immediate descendants exhausts the immediately superior property. So also, in the case at hand, there is no assumption that the language has primitive predicates for each of the mammalian species. Rather, we want to allow for the discovery of new primitive properties (new species), and their introduction into the language, though not perhaps without limit. I shall return to this point in a moment.‡

Given the above, it's fairly easy to see how the structure of properties should constrain the assignments of extensions to the primitive predicates, those predicates which are associated with nodes in the structure. Such predicates are interpreted, first, by assigning them intensions, properties. The structure among these properties, then, generates the relations among the extensions (sets) associated with the predicates. Specifying a model for such a system would go

†We will be able to generate non-primitive predicates of any arity by way of composition, e.g. by forming relational products. But we assume that there are only finitely many primitive properties within the ken of the language, at least to begin with.

‡As should by now have been made clear, and in case it hasn't: "primitive" does not mean "simple". The guiding intuition is that the relation between genus and species is like that between determinables, such as the property of being colored, and determinates, such as the property of being red. The latter includes the former; but the property of being red is not to be analyzed as *consisting* of the property of being colored together with some other property (unless that other property is that of being red). Nor is the property of being colored to be analyzed as an infinite disjunction of its determinates; indeed, it does not *include* any of them; though, to repeat, each of them includes it. In general, more general primitives (determinables) are simpler than the less general primitives (determinates). For a related, but different view, see [25, 26].