AI Magazine Volume 5 Number 2 (1984) (© AAAI)

Research in Progress

Center for the Study of Language and Information Research Program on Situated Language

Stanford University
Stanford, CA 94305

FOUNDED EARLY IN 1983, the Center for the Study of Language and Information [CSLI] at Stanford University grew out of a long-standing collaboration between scientists at research laboratories in the Palo Alto area and the faculty and students of several Stanford University departments and out of a need for an institutional focus for this work on natural and computer languages. At present, CSLI has 17 senior members and about as many associate members, from SRI International, Xerox PARC, Fairchild, and the Departments of Computer Science, Linguistics, and Philosophy at Stanford. Since the Center's research will overlap with the work of other researchers around the world, an important goal of CSLI is to initiate a major outreach, whereby members of CSLI both inform themselves of work done elsewhere and share their own results with others.

As its first major research program, CSLI is undertaking a study of situated language, Program SL, that is based on three insights:

1 Language use is fundamentally computational in that it is used by finite agents with limited resources to process, store, and communicate information;

2 Computational practice is fundamentally linguistic in that computers are used by humans under the assumption that the symbols and processes of computers are about entities in the world; and

3 Understanding linguistic activity in any real situation requires theories based on solid semantic foundations, connecting computation and language with information about the world.

The major goals of the program, dictated by the current state of affairs in research on languages, are:

1 To extend the study of natural language to include the active, situated agent;

2 To extend the study of computer languages to consider information content and the embedding world;

3 To merge the insights and theories of these two traditions into an integrated whole, based on solid philosophical and mathematical foundations.

Program SL, outlined here, comprises 16 projects in three areas:

1. Area NL—the evaluation and development of syntax, morphology, and phonology, the study of their computational aspects and their relation to language use, and the development of theories of natural-language use from the perspective of language users as finite processors of information derived from the world;
2. Area CL—the understanding of computing languages and architectures supporting the thesis of computation as linguistic activity; and
3. Area F—the development of the philosophical and mathematical foundations needed to support the theories on situated language use.

Questions about CSLI or Program SL should be addressed to Elizabeth Macken, Assistant Director, CSLI, Ventura Hall, Stanford University, Stanford, CA 94305.

Theories of Situated Natural Language

This collection of projects aims at developing scientific theories of natural-language use consonant with our basic perspective on language users as finite information processors. It will also evaluate and further develop areas of traditional linguistics—syntax, phonology, and morphology—in the light of our commitment to computation, semantics, and language use.

Extensions of Semantic Theories

Human languages contain a variety of expressive devices whose meaning is poorly understood because existing semantic theories are inadequately developed in crucial respects. Some of these constructions are well understood at the syntactic level, but either no semantical analysis of them exists or the semantical analyses that have been provided give no insight as to how a finite agent with bounded computational resources can understand them. Many key constructions are related to our understanding of foundational issues; the treatment of conditionals, for example, turns on the treatment of constraints, a key notion in the new theory of information emerging from situation semantics.

Of particular interest is how the interpretation of an utterance is related to its syntactic form. The principle of compositionality is the claim that the meaning of an expression is a function of the meanings of its constituent parts. This principle has served as a guide in Montague grammar and in most truth-conditional accounts of meaning. However, such accounts have usually ignored the difference between the meaning of an expression and its interpretation in a given discourse, so the principle has become very confused. Also, it is virtually a theorem that one can always make meanings compositional if one makes them complicated enough. The problem is to capture the intuition behind the principle while keeping the theory computationally tractable. Another aspect of the problem is the application of the principle to languages with freer word order than English, such as German, Hebrew, and Walpiri.

We aim to develop computationally tractable semantical analyses for a range of natural-language constructs not currently well understood. These will include tense, aspect, and time adverbials; comparatives; reflexive pronouns and reciprocal noun phrases; mass terms, plurals, and other forms of collective reference; modal verbs and conditional sentences; locative prepositions and adverbials; and verbal modifiers in general. Cases testing the principle of compositionality include wide-scope phenomena, extraposition, and topicalization.

Semantics of Sentences About Mental States

Of all types of expressions in natural language, sentences about mental states, such as belief, desire, and intention, play a uniquely important role in the development of semantical theories. The reason for this is that sentences about mental states typically involve embedding a sentence or sentencelike expression in a way that depends critically on the information content of the embedded sentence. It has proved to be extremely difficult to develop a theory of natural-language semantics that systematically assigns meanings to sentences in such a way that the facts come out right when those sentences are embedded in contexts like "John thinks that ..." Many semantical theories have foundered on exactly this issue.

Furthermore, understanding the semantics of sentences about mental states provides essential support for understanding the nature of mental states themselves, which is of critical importance to a comprehensive theory of language use. New semantical theories have cast all of these problems into a new light, in ways that emphasize the informational focus of sentences about mental states. These theories need to be developed and tested by their ability to guide the development of overall theories of natural-language semantics and computational theories of mental states as part of a comprehensive theory of linguistic communication.

Project Manager: John Perry.

Integrated Syntactic and Semantic Accounts of Discourse

Everyday language is rife with ellipses, sentence fragments, and other anaphoric devices. At present, we only partially understand these phenomena and hence are unable to explain a central facet of situated language. Work in linguistics has developed taxonomies of anaphoric devices, isolating distinctive properties of various types. Research in artificial intelligence has isolated extralinguistic factors that influence the interpretation of anaphoric elements, ellipses, and fragments. And some psycholinguistic research has suggested specific hypotheses about the way incomplete sentences of various kinds are processed. We will attempt to isolate the levels of linguistic structure that serve as the basis for the interpretation of elliptical, fragmentary, and anaphoric sentences, and the computational processes by which these structures are manipulated in language use.

Project Manager: Ivan Sag.

Integration of Semantical and Computational Accounts of Discourse

Research in new semantic formalisms (*e.g.,* by Barwise and Perry) and semantical theories of discourse (*e.g.,* by Kamp) have provided formal foundations and notions such as "role" and "value loading" that are essential for the explanation of the use of pronouns and other referring expressions within a theory of discourse. However, this work so far has not considered extended sequences of utterances and deals with discourse only statically. Research in artificial intelligence (*e.g.,* by Grosz and Sidner) has led to computational mechanisms such as focusing and centering to address the same issues. This work has dealt with extended sequences of utterances and with the dynamics of discourse but has been based on inadequate semantics. This project aims to integrate these two treatments to provide a firm foundation for the computational work and to extend the semantic work beyond a static view of discourse.

Project Manager: Barbara Grosz.

Communication

The study of language, not just as an abstract structure but as a medium through which agents exchange, store, and process information, should place linguistic behavior within the larger framework of a language-independent theory of reasoning and action. A theory of discourse, and of linguistic actions in general, should tie actions to mental states, identify the effects of utterances, and specify how mental states constrain what acts can be performed. Starting from work in artificial intelligence that extends plan construction and recognition systems to include linguistic actions, we shall attempt to provide a semantics for linguistic actions inspired by semantic theories of programming languages. In conjunction with the project on reasoning and planning (F.2), this should lead to algorithms for the planning and recognition of complex actions, including both linguistic and nonlinguistic subparts.

An important test for this theory will be to provide a uniform account of questions, imperatives, and declaratives that shows, for example, relations between questions and requests to assert. Making use of developments in the project on integrating semantical and computational accounts of discourse, we will extend the treatment of communicative acts to the discourse level, including pronouns and referring expressions, and then study discourse elements that function as indications of how various parts of the discourse are to be related.

Project Manager: Raymond Perrault.

Phonology, Morphology, and Syntax

New theories of lexical phonology and morphology have stressed the natural decomposition of sound-mapping rules into word-internal and phrasal types, each having distinctive properties. Independently of this, new syntactic theories have given evidence that many processes previously thought to be syntactic are morphological. To explore further the relation between phonology and morphosyntax, and their contribution to meaning, this project will study word formation and syntax in a representative sample of morphologically varying language types. The results of these studies will be integrated with the current theory of morphologically governed sound structure. We will look into whether there is a natural computational interpretation for the theoretically motivated rules and representations and whether the natural phonological and morphological algorithms interact properly with higher level syntactic and functional computations.

We will investigate syntactic structure as it relates to semantic interpretation and sentence generation. The study of syntax is particularly relevant in the context of topics such as the extensions of semantic theories

To develop adequate semantic treatments for many of the constructions listed there, we need better descriptions of their syntactic composition. We will try to integrate the insights of current theoretical approaches to syntax, such as lexical-functional grammar, generalized phrase structure grammar, and government binding theory. Another topic with close connections to syntax is the computational properties of parsing algorithms.

Project Manager: Lauri Karttunen

The Effect of Syntax and Phonology on Discourse Structure

Our current theories of phonology and syntax are largely restricted to the phrase and sentence levels, respectively Theories of discourse are largely unconnected with the phonological and syntactic levels, yet several cases are known where theoretical accounts at one of these levels depend on factors at another. Phrase-level phonological grammars have advanced our knowledge of sound variation and intonation, yet they do not show how these phenomena depend on syntactic and discourse boundaries Conversely, intonation influences illocutionary force, as does the discourse situation. Certain sentence forms (*e.g.,* clefts and topicalization) influence focusing and centering and thus constrain the referents of pronouns. This project will investigate the relation between the phonological, syntactic, and discourse levels and the computational problems involved in doing recognition and generation across them.

Project Manager: Joan Bresnan.

Strategies and Tactics in the Processing of Utterances

Many sentences have multiple meanings, though in any given context one is typically preferred. To model effective communication, we need to develop algorithms that produce the preferred reading. Semantic and pragmatic factors strongly influence preferences, but there is also evidence for independent lexical and syntactic effects. Explicit computational models of how lexical and syntactic properties interact during the comprehension process have also been developed.

To simplify the development of the theory of lexicosyntactic interactions, possible semantic and pragmatic influences were treated as constant boundary conditions that could be ignored, at least temporarily. As our computational theories of discourse and semantics improve to the point at which coherent characterizations of the situational context can be given, we expect to be able to extend our theory of ambiguity resolution to account for the interacting influence of factors at all levels of analysis. The result of this project, then, will be a full-fledged model of sentence interpretation that is sensitive to both structural and contextual factors.

Project Manager: Ronald Kaplan.

Computational Properties of Parsing Algorithms

Principals of CSLI have developed and are currently investigating several formalisms for encoding syntactic descriptions of natural languages: the generalized phrase structure grammar [GPSG], the lexical-functional grammar [LFG], the phrase-linking grammar, and PATR-2. These theories are noteworthy in that they all admit simple and direct recognition and generation algorithms and thus can be incorporated naturally into realistic models of a finite agent's language processing. Some parsing models have already been constructed, and it seems that despite their superficial differences, all of the theories depend on a combination of algorithms for context-free analysis and unification.

This project will study and implement alternative ways of performing the mappings that the syntactic theories describe, by extending current parsing technology and also by constructing new strategies for language generation. We will also strive to understand the computational consequences of various descriptive devices that our formalisms include and perhaps develop a computational rationale for choosing among them.

Project Manager: Stanley Peters.

Theories of Situated Computer Languages

The object of the projects of this section is to understand current theories and practices of computing within the perspective of computation as linguistic activity.

Semantics of Computer Languages

The aim of this project is to develop a semantic account of computation, rich enough to account for current computational practice, that can lead into the kind of theory developed in Area F dealing with the foundations underlying a unified view of language. In particular, we plan to develop a semantic account that encompasses both natural and programming languages. Because of the similarities between the information systems in domain theory and the informational relations in situation semantics, one of our first aims will be to connect with the substantial body of work that has been done on programming languages by Dana Scott and others working with Scott's domain theory. To construct a theory that is adequate to deal with the structures of current computation theory, we will try to clarify the interrelation of three things:

1 The higher order formalisms of computation theory;
2. The formalisms normally used in programming languages, specification languages, verification systems, and analysis systems; and
3 The new structures being developed in situation semantics

Also, we will investigate the semantics of the reflective lambda calculus, leading into the semantics of 3-LISP, and the semantics of computational and reflective logics. Finally, we will attempt to develop a general theory of the syntactic and structural operations on which all formal systems are based (procedure calls, structural unifications, schema instantiation, etc.).

Acting Project Manager: Jon Barwise

The Analysis and Design of Linguistically Coherent Computer Languages

Traditional programming languages have been restricted primarily to commands; their syntax was essentially trivial, although not their semantics. Now they are becoming extremely complex—far more so, for example, than the artificial languages typically studied in mathematics. There are linguistic structures that deal especially with temporality (sequencing, including the interweaving of concurrent sequences), plurality (sets and other such collections), metadescription (constructs that refer to the program text or some aspect of its interpretation, rather than to the computation it describes), and information grouping (grouping into modules whose contents are in some sense isolated from those of other modules). The aim of this project is to put some theoretical order into this rather diverse set of practices, partly to develop even more powerful structural protocols and partly to understand a level of information complexity intermediate between natural and traditional formal languages. We will concentrate on languages with declarative structures (*e.g.*, ALEPH and 4-LISP), on the relation between description and control, and on the relation between input/output and communication.

Project Manager: Terry Winograd.

Computational Architectures for Reasoning

Under the assumption that computational processes are radically idealized language users, it is natural to develop specific architectures that use and reason with language in just the ways that our theories describe. The idea is to design a simple calculus that explicitly reveals the important aspects of the emerging theory, a role played by the first-order predicate calculus in the development of logic and model theory. We will develop a computational architecture, called MANTIQ, that serves this role in our developing theories of language in use.

MANTIQ will be based on essential insights from both computational and natural languages. From the computational side, it will draw on both procedural languages (like LISP) and descriptive languages (such as specification languages and knowledge-representation languages in artificial intelligence). It will also be based on the theories of inference and reasoning that grow out of our studies of natural language and out of our foundational studies on information. Technically, it rests on two important developments: a full theory of what we call reflection and an internal notion of structure based on information content. By reflection, very briefly, we mean the ability of an agent to reason effectively about its own operations, structures, and behavior. Structurally, MANTIQ will be based directly on the theories of intensional identity developed in the foundational semantic parts of our overall program. A goal of the MANTIQ design is to have internal structural identity directly encode semantic identity, which will make it possible to describe the architecture entirely in terms of content, in line with the informational orientation of our entire research program.

MANTIQ holds out the promise of modeling the intentional aspects of human language use, which are an essential part of a theory of action. It will also serve both as a test bed and as a forcing function on theories of belief, planning, and so forth. Formally, it will rest on semantic accounts of concurrent architectures and on theories of procedure (both mentioned in the project about the semantics of computer language); its implementation may rely on the unification procedures described in the project about the computational properties of parsing algorithms.

Project Manager: Brian Smith.

Foundations Underlying a Unified View of Language

This area aims at developing the philosophical and mathematical foundations needed to support the theories on situated language.

Computation, Information, and Logic

This project is intended to contribute to the mathematical development of those parts of logic relevant to our overall program on language, computation, and information. In particular, we plan to study the mathematical properties of (semantic) information, relating it to work in model theory and generalized recursion theory as well as to the measure of information used in Shannon's communication theory. Long-term goals of this project include:

- Providing theoretical concepts with which to analyze the interaction of computation and information in situated language users, both human and computer;
- Providing useful measures of the semantic content of a message relative to certain background information and constraints;
- Providing more semantically relevant notions of computational tractability; and
- Contributing to the logic of both human and computing languages

Project Manager: Jon Barwise.

Reasoning and Planning

This project will investigate the computational processes necessary to perform the reasoning required by the use of language. It will be divided into three interrelated parts: reasoning about the external world, reasoning about mental states and actions, and planning.

In its first part, we will develop a formal, computational theory of commonsense reasoning precise enough to permit a direct, efficient implementation. We will limit ourselves to those aspects of the world studied in the project on the commonsense world. The second part on reasoning about mental states is the computational complement of the project on mind and action and an extension of the work already done by Moore, Appelt, and Konolige to include desires and intentions. Finally, "planning," which denotes the mental processes by which intentions are established and revised, is a form of reasoning that differs from general reasoning in that it has a concrete aim: finding (executable descriptions of) actions that achieve the agent's desires. While several computational models of the planning process have been suggested, there are still considerable difficulties in handling more subtle notions of desire (*e.g.*, graded preferences), in smoothly integrating the monitoring and replanning processes, and in reasoning about complex future actions that cannot easily be thought of as sequences of abstract operations.

We also intend to explore whether an integrated view of planning and reasoning was abandoned prematurely. One possible end result of our research would be the disappearance of planning as a separate subject of study altogether, subsumed in a more inclusive and deeper theory of general reasoning.

Project Manager: Stanley Rosenschein

Mind and Action

This project will attempt to bridge the gap between computational theory and practice, on the one hand, and philosophical insight, on the other, by using formal methods to bring intuitive theories of mind and action into a computational frame of reference. This will involve building a common technical vocabulary, possibly based on work in theoretical computer science on formalizing the relation between levels of abstraction in the description of complex computational processes (*e.g.*, work on abstract data structures and the semantics of high-level languages). The ultimate goal is either a computationally meaningful reinterpretation of much of the intuitive terminology from the philosophy of mind and practical reasoning or a more radical revision of our ideas on how to describe mental structure and process. One important question we will try to answer is whether our model should include as a separate component each of the many attitudes that our language names (*e.g.*, believe, want, intend, fear). If not, what criteria should be used to collapse them? We will look at the relationship between an objective "observer's" theory of mind and action and the commonsense "participant's" theory that we apply to each other in everyday life. Further, we will examine whether there is a systematic method for abstracting the latter type of theory from the semantics of propositional attitudes, how we can account computationally for the way that rational deliberation results in the causation of action, and what mechanisms lead to "changes of mind."

Project Manager: John Perry.

The Commonsense World

Generating and interpreting fluent natural language requires considerable abilities to do commonsense reasoning, which in turn presupposes an explicit elaboration of our commonsense theories of the world. Such theories are also needed for extending semantical theories of natural language, since the semantics of our language and our commonsense view of the world are inextricably intertwined. We will focus on a handful of commonsense theories that are so basic to our view of the world that they arise in some form in almost any domain of discourse, for example, the commonsense theory of space and motion. We will also choose areas in which natural language has evolved special mechanisms for expressing information, so that a commonsense theory in such an area is almost essential to carrying out the semantical analysis of that part of language. The commonsense theory of time, for instance, must be understood in order to explicate adequately the semantics of tense and aspect.

Project Manager: Robert C. Moore.

Current Research Members of CSLI

John Barwise	Director of CSLI Stanford University
Stanley Peters	Associate Director of CSLI Stanford University
Joseph Almog	Stanford University
Douglas Appelt	SRI International
Joan Bresnan	Stanford University Xerox PARC
Marsha Bush	Fairchild
Philip Cohen	Fairchild
James des Rivieres	Xerox PARC
Marilyn Ford	Stanford University
Joseph Goguen	SRI International
Barbara J. Grosz	SRI International
Kris Halvorsen	Xerox PARC
Jerry Hobbs	Stanford University
Pentti Kanerva	Stanford University
Ronald Kaplan	Xerox PARC
Lauri Karttunen	SRI International
Martin Kay	Xerox PARC
Paul Kiparsky	(principal researcher, visiting)
Kurt Konolige	SRI International
John McCarthy	Stanford University
Jose Meseguer	SRI International
Robert C. Moore	SRI International
Helen Nissenbaum	Stanford University
Fernando Pereira	SRI International
C. Raymond Perrault	SRI International
John Perry	Stanford University
Carl Pollard	Stanford University
Geoffrey Pullum	UC Santa Cruz
Stanley J. Rosenschein	SRI International
Ivan Sag	Stanford University
Stuart Shieber	SRI International
Mark Stickel	SRI International
Susan Stucky	Stanford University
Patrick Suppes	Stanford University
Brian Cantwell Smith	Xerox PARC
Hans Uskoreit	SRI International
Thomas Wasow	Stanford University
Terry Winograd	Stanford University
Meg Withgott	Stanford University