

The Panels

I. Words and world representations.

How have these suddenly become more interesting? Do they offer a way through from the old "primitive" dispute, and do they offer a way out from having to separate world and linguistic knowledge? How does what we know about words fit into the language understanding and generation process, and is that different for understanding and generation ?

Don Walker (Bellcore) Chair
Bran Boguraev (Cambridge)
Bob Amsler (Bellcore)
Jerry Hobbs (SRI)
Judy Kegl (Princeton)

II. Unification and the new grammaticism

How far does this really differ from the CFG position of the sixties? Does it yet have any empirical successes in terms of working systems? To what extent are these grammatical formalisms motivated by processing considerations ?

To what extent are these processing claims substantiated ? Are we converging to some class of formalisms that are relevant for processing and, if so, how can this class be characterized in a theoretical manner ?

What are the prospects of these types of formalisms becoming the basis for future natural language processing research ? Has the processing paradigm now really fundamentally influenced linguistics ? Do processing considerations and results show that such systems when implemented can be neutral between analysis and production.

Has everyone really been doing unification for decades and just found out? Is it a real advance or just a Hollywood term?

Fernando Pereira (SRI) Chair
Gerald Gazdar (Sussex)
Steve Pulman (Cambridge)
Aravind Joshi (U.Penn.)
Mitch Marcus (Bell Labs.)
Martin Kay (XEROX-PARC)

III. Connectionist and other parallel approaches to natural language processing

Is NLP inevitably committed to a symbolic form of representation? Can syntactic, semantic or world knowledge be represented in that paradigm if taken seriously? What parts of current CL will fare worst if there turn out to be significant empirical advances with connectionist parsing? Are there any yet (i.e., how far do we trust simulations programmed only on serial machines)?

What new approaches to syntax, semantics or pragmatics will be needed if this approach turns out to be empirically justified? Will it just bring back all the old views associated with associationism, and will they be changed in the journey? Is parallel parsing just a new implementation or a real paradigm shift?

Dave Waltz (Thinking Machines) Chair
Martin Kay (XEROX-PARC)
Gary Cottrell (UCSD)
Gene Charniak (Brown)
Jay McClelland (CMU)
Wendy Lehnert (U.Mass.)

IV. Discourse Theory and Speech Acts.

Is there yet any serious discourse theory with testable computational and empirical consequences? What phenomena ought a processing theory of discourse understanding/generation to address itself that are not already being attended to currently? What aspects of discourse are language problems and which are general AI/KR problems? What makes a theory of discourse a processing theory? Does spoken language affect one's theory of discourse?

Is there any real hope that we will be able to recognize the plans/goals etc. of a speaker? How much of conversation is carried on through the linguistic window anyway? Do current theories of text and dialogue discourse mesh, and should they?

Barbara Grosz (SRI) Chair
Julia Hirschberg (Bell Labs.)
Ray Perrault (SRI)
Bob Wilensky (Berkeley)
Franz Guenther (Tuebingen)

V. Why has theoretical NLP made so little progress?

Has CL advanced in this respect since Tinlap2 in 1978? What can NLP systems do today in the light of what we would have predicted at Tinlap2? Why are we no nearer to a common notation for systems since KRL — would we be helped by CL textbooks geared to particular programming languages (one such is now in preparation)? Is it a case of just cycling through ranges of obscure syntactic and semantic formalisms (and then rediscovering them every 10 years or so)? Are there serious problems about the overall cognitive paradigm being applied to NLP? Are there any serious alternatives to the current paradigms, and what would they imply to NLP research directions and goals?

Roger Schank (Yale) Chair
Norm Sondheimer (USC-ISI)
Larry Birnbaum (Yale)
Ivan Sag (Stanford)
David Israel (SRI)

VI. Formal versus common sense semantics.

What does Montague grammar or situation semantics have to say to CL? Can we distinguish the good parts from what is bad and useless? For what NLP applications might these formalisms be particularly appropriate? What have such theories chosen to ignore, in terms of data or intuitions? How are they to be computed: compositionally, randomly? How well can such formalisms mesh with the rest of language representation processes, e.g., discourse and pragmatic analysis?

Yorick Wilks (NMSU) Chair
David Israel (SRI)
Geoff Nunberg (Stanford/Xerox-PARC)
Wendy Lehnert (U.Mass)
Karen Sparck-Jones (Cambridge)
Susan Stucky (Stanford-CSLI)

VII. Reference: the interaction of language and the world.

When is a noun phrase a referring expression? How does the meaning of a noun phrase contribute to the success of a referring act? How can a "wrong" description be useful for referring? Is there any role for Russell's analysis of descriptions in a pragmatic theory of referring?

What does it mean for a hearer to identify a referent? What is the relationship between knowing who or what something is and referent identification?

Is referring to events and situations inherently different from referring to material objects? What identification criteria are applicable to events and situations?

Doug Appelt (SRI) Chair
Deborah Dahl (SDC Inc.)
Bonnie Webber (U. Penn.)
Amichai Kronfeld (SRI)
Brad Goodman (BBN)

VIII. Metaphor

How relevant are the philosophical, linguistic, and psychological literatures on metaphor? Can any of the recent work in dialogue, planning and speech acts be applied to understanding metaphor? Are existing knowledge formalisms (e.g., conceptual dependency, scripts, semantic networks, KLONE) adequate for metaphor? If not, why not? Given that the recognition of metaphor involves matching together large-scale knowledge structures, are there any existing procedures that do this adequately? How can this matching be done? How might we record the degree of match? Are there additional types of processing necessary for recognising metaphor?

How should metaphor be represented in semantic representations of text? Are there situations when a metaphor should be "resolved", and others when its tension should remain? How can we recognise those situations?

Deirdre Gentner (Illinois) Chair
Andrew Ortony (Illinois)
Ed Plantinga (Toronto)
George Lakoff (UCB)
Geoff Nunberg (Stanford/Xerox-PARC)

IX. Natural Language Generation

Will the demands of language production bring AI, theoretical linguistics (and of course CL) closer together than the demands of comprehension did in the past? Is there anything special about generation?

Does generation constrain problems differently from understanding, in that it would not matter if some high-powered machine could understand things no human could say, but it would matter if the same machine generated them? Are knowledge structures, of the world as much as language, the same or different for understanding and generation? What is the relation between the message the system wants to convey and its lexical, syntactic, etc. abilities to do it.

Aravind Joshi (U. Penn) Chair
Dave MacDonald (U. Mass. Amherst)
Doug Appelt (SRI)
Bill Mann (USC-ISI)
Mitch Marcus (Bell Labs.)
Tony Kroch (U. Penn.)

