Project proposal for January 17-31: Production and Interpretation of Natural Language

Henk Zeevat

1 Background

I am currently about halfway writing a book under the (provisional) title of this project. The project would be a one-week course going through the whole book including the unfinished chapters, rounded off by a short essay on an aspect of the material to be completed in the second week. The course would take place in the first week, the essay in the second week and the whole thing in the second half of January.

The book tries to make sense of a position not unlike the one of Clark and Hecht (1983) in which both processes, production and interpretation are taken as fundamentally different processes and the coordination between speaker and hearer on the message only occurs at the highest level. The arguments of Clark and Hecht come from language acquisition and involve especially the gap between production and interpretation: language users at any level of acquisition understand far more than they can produce and the gap still exists in adult grammar. But it should not need any special argument that the two processes are not the same or simply each other's inverse from each other: we do not speak with our ears or listen with our tongue, the deeper level in production would be most comparable to planning, and in comprehension most comparable to perception.

From this perspective, standard grammar models giving a formal characterisation of the relation between forms and meanings taken as a model of both psychological processes typically fall short on the following observations.

1. production and interpretation have linear time complexity
2. utterances can be surprisingly incorrect without impairing understanding
3. the underdetermination of meaning by form does not seem to impair a high level of successful coordination on meanings

Parsing and generation with a rule based grammar are search problems not solvable in linear time (1). If ambiguity is as significant as is predicted on theoretical grounds that means that the chance of picking the intended meaning for an utterance is minute, so that convergence would a very rare (3) event. (2) is not understandable without assuming an additional mechanism.

While the standard response to these problems is to postulate additional mechanisms that bridge the gap between competence and performance, it does not

seem necessary to postulate these additional mechanisms when production and interpretation are separated. Production can be done as a semi-deterministic transformation of the message into the form and interpretation can be seen as the stochastically constrained integration of the concepts cued by the words, the main mechanism in high level vision for the concepts cued by the results of low level vision.

A problem in accounts that separate production and interpretation is that they seem to require separate grammars. That has a problem of its own: one would expect errors in communication that are due to production and interpretation grammar to be out of sync and such errors have not been shown.

The solution that is advocated in the book is that this can be done by assuming the involvement of both production and interpretation in the other process (in the spirit of bidirectional optimality theory) and assuming only a minimal overlap between the knowledge that is required in each process. Interpreters simulate production in the attempt to find the intended interpretation (to the degree that they can) in order to increase the probability that they are right. Similarly, speakers inhibit productions which they can see would lead to misinterpretation.

Syntactic and grammatical generalisations will be primarily explained by letting these be constraints on the production process: a combination of paradigmatic mapping, word order constraints and morphological constraints. Constructions, word order and morphology are after all properties of the utterances that speakers produce and an account of production needs to account these properties. They have an influence on interpretation only by the assumption of simulation: different interpretations will often lead to different utterances.

The interpretation component on the other hand will be constrained by conceptual knowledge. After all concepts —this could be an approximative definition of concepts— are what is cued by words and an interpretation is an integration of concepts. Concepts come with various sorts of knowledge, but an important part is information of how to connect them to other concepts. Just assuming that kind of information alows a formalisation of conceptual combination in the style of a probabilistic categorial grammar.

An interesting aspect of the overall model is that one can conceive of both processes as consisting giving models of natural language that in isolation are only bad approximations to language as produced and interpreted: the isolated syntax also produces unacceptable utterances, the isolated interpretation also produces interpretations that the utterance could not even have. But the conditioning of the processes on each other would let both processes approach human performance.

2 Chapters

Chapter 1 **Parity** presents the argument from coordination under underdetermination for the overall model by showing that the model —if properly

developed— can explain the fact that verbal communication between people usually leads to convergence on the intended meaning even when the utterances by themselves are highly ambiguous. Competing models do not allow speaker and hearer strategies that make this more than a rare event. Probabilistic theories of interpretation only give part of the explanation: coordination is still unlikely if the speaker has not done her part in promoting it.

Chapter 2 **Syntax** is an investigation of the computational and linguistic content of optimality theoretic syntax which is taken as the starting point for a production model. The example development will be an account of Dutch clause word order and Dutch morphology and of some aspects of lexicalisation and sentence planning. The chapter moreover introduces a general way of implementing OT syntax as a ordered set of procedures which try to add more information to an underspecified sign in the HPSG tradition. The algorithm can be interpreted as one in which the activation of a partial sign given by its semantics cause the activation of a full sign and this interpretation can explain why production is fast (linear). The chapter also shows that the algorithm can be seen as an incremental generator that can produce initial sentence fragments from a initial fragment of the semantic specification in the context. This last result is crucial for letting the production grammar simulate production in the incremental version of the interpretation mechanism of chapter 6 **Interpretation**.

Chapter 3 **Automatic Self-monitoring** gives the linguistic evidence for automatic self-monitoring (the speaker inhibits productions that are expected to lead to misunderstanding), based on a treatment of hard syntax as in chapter 2. It is successful in giving a general OT style formalisation of the process and in giving an linguistic empirical characterisation of the phenomenon. But it centrally tries to show that an account of syntactic correctness is incomplete without bringing in self-monitoring since syntax alone cannot deal with quite a number of central properties of utterances. The phenomena covered are particle insertion, word order freezing, differential case marking and NP selection, and less conclusively the French silent h. The chapter ends with a brief discussion of the relevance of self-monitoring for the evolution of languages.

Chapter 4 **Pragmatics** is about the foundations of pragmatics. It comes out of an attempt to take the general assumptions in the presupposition theory of Heim (1983) and van der Sandt (1992) as defeasible constraints in an OT constraint system. The chapter shows that the small resulting constraint system is in fact a complete account of general pragmatics, capturing more than e.g. Grice's cooperation-based pragmatics and closely related to Relevance Theory (Sperber and Wilson (1995)), minimal models approaches (e.g. Schulz (2007)) as well as Hobbs' interpretation by abduction Hobbs et al. (1990). The system uses RELEVANCE (interpret the utterance as settling activated questions when the utterance matches one such question) as the weakest constraint, below a coherence constraint (*NEW), a constraint that maximises the probability of the message (PLAUSIBLE) and a constraint demanding obedience to syntax (FAITH).

As in Hobbs' pragmatics, the system can be interpreted as finding the most probable explanation of the utterance as a communicative act of the speaker in the context. In this respect, the constraint FAITH makes the interpretation an explanation of the utterance and PLAUSIBLE selects the best explanation. The addition of *NEW and RELEVANCE can be seen as other aspects of maximising the probability, base don the expectation that the speaker is coherent and cooperative.

The chapter runs through a set of examples of pragmatic enrichment and shows how they can be derived in the system. It ends with a discussion of semantics versus pragmatics in which argues against strict separation, while defending at the same time that NL semantics has an important and largely unaccomplished programme, that of doing the logical analysis of the basic building blocks of the primitive concepts in which interpretation must be carried out.

Rhetorical Structure is a further application of the approach to pragmatic enrichment to the area known as rhetorical relations, discourse relations, discourse grammar and otherwise. The current approach gives a simple explanation of the basic assumptions of this field.

Chapter 5 **Semantics** develops a concept-based variant of discourse representation structures Kamp and Reyle (1993) that formalises a number of central notions in the book: interpretational heuristics, semantic interpretation, contexts, semantic input for production, and mental representation. The claim is that the structure proposed is essentially needed for the heuristic algorithm and for production, but there is a further claim that the structure is also an improvement from a psychological and philosophical perspective, as well as a simplification of DRT. The chapter is foundational for many of the aspects of the central thesis and is largely logical in nature.

Chapter 6 **Interpretation** proposes basic heuristic interpretation and full interpretation and so closes the circle of the argument. Basic interpretation is a probabilistic free categorial grammar combining ambiguous concepts into larger conceptual units, reducing unsaturatedness and ambiguity in the process. Full interpretation is the constraint system of chapter 4 operating on the output of basic interpretation and can be seen as probability guided inference. An important thesis is that full interpretation cannot be reduced to basic interpretation.

3 Why go to the course?

It should be clear that you would be doing me a favour. I need as much feedback as I can get in order to finish the writing. A very important aspect is also to see to what extent the content comes across to people who have not seen it develop.

I think the content is special for its broad disciplinary perspective. The aim — parity— is philosophical and the means —low complexity incremental accounts of production and interpretation— cognitive science for a psycholinguistic goal: how can two people coordinate with something as vague and ambiguous as

natural languages. Chapter 2 and 6 are formal syntax and semantics of the kind that is relevant for computational linguistics, as it turns out a field that is directly relevant for the philosophy of language and cognition: questions such as the role of probability in interpretation get very principled answers. Chapters 4 and 5 fall within formal semantics and pragmatics and make contributions to dynamical interpretation. Chapters 2 and 3 can be read as contributions to natural language syntax and deals with a wide range of problems that are central in current discussions, including the foundations of historical linguistics. The bidirectional aspect of all the whole model is directly relevant for accounts of perception, understanding by simulation, mirror neurons and others.

The course will be rounded off by an essay which has to be written in one week. It is likely that you will at least make it to a credit in the book, but it is not unlikely that you can produce some ideas which would be publishable.

4   Practical (to be expanded)

Monday to Friday 11-1: lecture
optional: working lunch: 1-2
rest of day: read the material and think about essay topic

reading material: MS, follow up references therefrom (to be selected according to your needs)

References

Clark, E. and Hecht, B. F. (1983). Comprehension, production, and language acquisition. *Annual Review of Psychology*, 34:325–349.

Heim, I. (1983). On the projection problem for presuppositions. In Barlow, M., Flickinger, D., and Westcoat, M., editors, *Second Annual West Coast Conference on Formal Linguistics*, pages 114–126. Stanford University.

Hobbs, J., Stickel, M., Appelt, D., and Martin, P. (1990). Interpretation as abduction. Technical Report 499, SRI International, Menlo Park, California.

Kamp, H. and Reyle, U. (1993). *From Discourse to Logic*. Kluwer, Dordrecht.

Schulz, K. (2007). *Minimal Models in Semantics and Pragmatics: Free Choice, Exhaustivity, and Conditionals*. PhD thesis, ILLC, University of Amsterdam.

Sperber, D. and Wilson, D. (1995). *Relevance: Communication and cognition (2nd ed.)*. Blackwell.

van der Sandt, R. (1992). Presupposition projection as anaphora resolution. *Journal of Semantics*, 9:333–377.