

Pragmatics (to a First Approximation)

David Beaver
Department of Linguistics
Stanford University

April 1999

Abstract

In this paper, pragmatics is redefined.

Contents

1 A pain in the butt.	2
2 What do those little numbers mean, and who put them there anyway?	2
3 Little numbers in dynamic semantics.	2
4 Yet more little numbers, and other problems not even amenable to simple numeric indexes.	
5 Doing without the little numbers.	4
6 Why should we care about resolution?	5
7 A strategy for pragmatics.	5
8 Every possible theory of pragmatics (to a first approximation).	6
9 The first approximation.	7
10 Pragmatics to a second approximation and beyond.	8
11 Example: presupposition accommodation.	9
12 Example: anaphora resolution.	10
13 Discussion	11

1 A pain in the butt.

Nobody should be serious about *dynamic semantics* for natural language. Or rather, nobody should be serious in thinking that what goes under the term *dynamic semantics* is just semantics. Interesting dynamic analyses of natural language are interesting precisely because they provide formal analyses of phenomena that have previously been considered both (a) partly pragmatic, and (b) a pain in the butt. Two phenomena that might leap to mind are discourse anaphora and presupposition.

I will introduce a framework for attacking such problems. Let us call it *Transition Preference Pragmatics*. You will soon see why.

I want to begin by considering a process that is a part of the phenomenon of discourse anaphora. Some would say that it is the most important part for many practical purposes. Yet it is a part that has usually been neglected in Amsterdam-style dynamic semantic analyses, presumably because it is considered either too pragmatic, too much of a pain in the butt, or both. The process is that of *anaphora resolution*, whereby the antecedent of an anaphoric expression is identified.

2 What do those little numbers mean, and who put them there anyway?

The level of abstraction at which many of us are accustomed to working would doubtless be staggering to some scholars of the textual form. They might be amazed to find out that some of the most sophisticated analyses of discourse anaphora yet proposed place almost no constraints on the resolution of pronominal expressions within a typical text beyond number/gender agreement, and simple left-right ordering of antecedents and anaphors. The only substantive constraint that is placed on discourse anaphora prevents resolution to antecedents buried under operators such as negations, implications and quantifiers. But just how sophisticated has a theory to be in order to state such a constraint? And given the phenomenon of modal subordination, should we even accept the empirical validity of that constraint?

All this is not to say that dynamic accounts of discourse anaphora have achieved nothing. Crucially, these theories not only place constraints on possible anaphoric relationships, they also provide a powerful theory of what discourses containing anaphoric expressions mean, and they do so without any violation of compositionality. What they do not do is say anything very surprising about the anaphora resolution process or how it is constrained.

3 Little numbers in dynamic semantics.

There is a bizarre assumption made in Groenendijk and Stokhof's Dynamic Predicate Logic (DPL) [GS91a] and Dynamic Montague Grammar (DMG) [GS91b], and Heim's File Change Semantics (FCS) [Hei82], an assumption that is not made in Kamp's Discourse Representation Theory (DRT) [KRe93]. This

is the assumption that anaphor-antecedent relationships are presented ‘gift-wrapped’ to the semantics in the form of pre-indexed NPs.¹ Who or what is supposed to do the wrapping is not made entirely clear.

Perhaps syntax does the job of pre-indexation? Yet syntax seems unlikely to do more than provide a few extra local constraints on resolution via C-command conditions etc.

If not syntax, then presumably pragmatics must do the job. But this would not sit well with what appears to be the architecture of DPL, DMG and FCS. For it seems that pragmatics should either operate after semantic interpretation, as is assumed e.g. in classic work on pragmatics by Grice, Austin and Searle, or at least should operate in tandem with semantics. It seems unattractive to force pragmatics to constrain semantic forms in advance of semantic interpretation. To this, defenders of DPL, DMG and FCS might counter that these models were never intended to describe a model of human sentence/discourse processing, but only of sentence/discourse meaning, so that an actual implementation of the models might somehow postpone indexation, or interleave it with semantic interpretation. However, I doubt whether the original proposers of DPL, DMG and FCS would find such a response attractive: it would certainly sit uncomfortably with the processing oriented rhetoric that one finds in the original papers introducing these theories.

4 Yet more little numbers, and other problems not even amenable to simple numeric indexation.

In numerous articles on linguistics and philosophy, I have seen little numbers attached to words which are intended to indicate which word-sense is intended, as in “Mary went to the bank₁.”

To my way of thinking, word sense markers are not significantly different from anaphor index markers. In both cases, it is obvious that these little numbers are purely artefactual, and not part of the real languages we wish to analyze. In both cases, semanticists routinely act as if syntax will provide semantics with pre-indexed disambiguated forms. And in neither case is this latter assumption really plausible. Pragmatics must surely be a significant factor in identifying word senses, and yet it seems unlikely that we would want to assume that pragmatics acted entirely before semantic interpretation.

But is this a significant problem? After all decorating surface forms with numbers seems such a harmless activity. And the numbers are usually kept very tiny indeed. Yet indexation is only one way that we use to disambiguate language so as to make it more amenable to semantic analysis. In general, disambiguated levels of representation used with semantic theories may step much further from surface form. It was Bertrand Russell’s thesis of misleading form that opened the flood gates. Nowadays, semanticists are happy to posit a wide range of relations between a single surface form and possibly numerous

¹In this respect Musken’s “Compositional DRT” [Mu96] is not like DRT but like DPL/DMG/FCS, since it does assume pre-indexation.

disambiguated logical forms². This disambiguation accounts both for genuinely syntactic phenomena, such as attachment, and for operator scope phenomena for which the assumption of a level of disambiguated LF is more problematic.

The assumption that semantics can operate on a disambiguated LF is not inherently problematic. It is more of a promisory note, something that must be tackled before our theory of interpretation is complete. Yet it is worrying that nearly a century after Russell first introduced semanticists to a convenient way of passing the buck, there is still remarkably little known about how disambiguation occurs.

What I find even more worrying, is that for dynamic semanticists who wish to stick by the processing rhetoric with which they have presented their theories in the past, the assumption of a disambiguated form prior to interpretation seems a simple contradiction.

5 Doing without the little numbers.

DRT does not assume pre-indexation for anaphora, but instead relies upon its *resolution algorithm*. As far as anaphora resolution is concerned, then, this is the dynamic theory which comes closest to living up to the dynamic rhetoric. And Uwe Reyle’s UDRT (and its brethren) incorporate the beginnings of a uniform treatment of all the different forms of ambiguity I have mentioned. I say “beginnings” because, as far as I know, the work has concentrated more on developing strategies for avoiding ambiguity resolution than on choosing the best way of resolution when forced.

Consider anaphora resolution once more. Standard DRT places essentially similar constraints on possible anaphor-antecedent relationships to DPL, DMG and FCS. Or, rather, it places a similar lack of constraints on possible anaphor-antecedent relationships. There is no limit, for instance, on how many sentences apart a pronoun may be from its anaphor. So, while in standard DRT there is clearly a place for a more detailed account of resolution to go, it is not yet obvious just what form such an account should take.³

Amsterdam-style dynamic semantics has an advantage over DRT: it is less committed to any particular form of mental representation. I believe that we are at a stage where it is useful to gain a general formal perspective on ambiguity resolution, that in the long-run this will lead to important results and a clearer picture of what a dynamic model of interpretation is for. For this reason, I will present my thoughts not in DRT, but in the more abstract Amsterdam style. That is to say, I will rely on a model-theoretic characterization of information states rather than an implementation of information states using a particular style of representation. Yet I think it will be clear that the ideas I develop are in the spirit of the analysis of ambiguity resolution that has already occurred

²In Montague Grammar, this becomes a range of alternative semantically unambiguous syntactic derivations.

³Some variants of DRT have been proposed which involve more sophisticated resolution algorithms. I am thinking here of the implementation developed within the ACORD project, and of Sophia Cormack’s PhD dissertation. The current paper is offered in the same spirit as these earlier proposals, albeit in a different formal setting.

in DRT, and that the architecture of the theory mirrors parts of existing DRT treatments in many ways.

6 Why should we care about resolution?

We should care about anaphora resolution because it is an essential part of natural language meaning and inference.

Consider the following discourse, from [GJW83], which is to be understood as a spoken discourse with no special emphasis of the pronouns: “Susan gave Betsy a pet hamster. She reminded her that hamsters were quite shy. And then Susan left.” It follows from this discourse that Susan reminded Betsy of something. It does not follow that Betsy reminded Susan of anything. These are examples of inferences in natural language which stem from general properties of how anaphora is resolved.

I do not see these facts as being different in kind from inferences which stem from general properties of connective or operator meaning. They are facts about what English means, and about which patterns of inference are valid in English. They are facts that a theory of the meaning of English should incorporate. But, unfortunately, they are not facts that are predicted by DPL, DMG, FCS or DRT.

And anaphora resolution is just a special case of ambiguity resolution. We should care about that too. For the same reasons. Consider the word-sense variation of “bank” in the following:

1. Mary had been living it up for some time, and needed some money. She drove over to the nearest bank.
2. Mary went to a financial institution.
3. Mary had been paddling upstream for sometime, and needed a rest. She steered her canoe to the nearest bank.
4. Mary went to the edge of a river.

Clearly, 2 follows from 1 (and not from 3), while 4 follows from 3 (and not 1). These inferences seem to be at least as robust as e.g. Aristotelian syllogisms. Perhaps more so for many speakers. A formal model of language meaning should capture these inferences as well as the inferences beloved of logicians, which largely center around operators and connectives. Theories like DPL, DMG, FCS or DRT do not capture such inferences. Nor do they even pretend to.

In this paper I will not detail a treatment of ambiguity resolution that can account for cases of anaphoric or lexical ambiguity in general. I will however sketch a general framework in which such a treatment could be developed, and I will sketch how this framework can be applied in several example areas.

7 A strategy for pragmatics.

Participants in a discourse have a model of the common ground, including both information about the world exterior to the conversation and information about

the world “interior” to the conversation, i.e. information concerning what is being talked about.

In making an utterance, speakers attempt to convey ways in which they think the common ground should be modified. **But**, and it is a \large “but”, hearers live in ignorance. They cannot be sure what transition the speaker intends. They do not know what ground was assumed initially common by the speaker, and, even if they had known, ambiguities of the utterance mean that they cannot be 100% sure what the speaker takes to be the final common ground.

Here’s the big picture. Or at least *a* big picture. Syntactic analysis and compositional interpretation, yield a set of alternative meanings. Each meaning is itself a set of transitions, i.e. pairs of information states conceived of as inputs and outputs, where an information state is one possible common ground. What do we need pragmatics for? The main reason we need it is to choose the right single transition, the one intended by the speaker, from amongst the set of sets of transitions provided by earlier stages of interpretation.

How then should we create a theory of pragmatics? First, we need a way to represent uncertainty about possible transitions. Second, we need to represent preferences between alternative transitions. Finally, we need to describe those preferences. The first two stages delimit the possible theories of pragmatics. The last stage is the hard one: choosing the right theory of pragmatics from amongst the set of possible theories. In fact the first two stages are so easy that I will suggest a fairly rich solution in the next section, whereas the third is so difficult that I will only make the tiniest of inroads into it in this paper, showing how some existing pragmatic theories fit into the mould.

8 Every possible theory of pragmatics (to a first approximation).

A model of the common ground can be represented with complete generality as a set of possible worlds. Relative to a world, certain constants and predicates have an interpretation in terms a domain of individuals. This domain is assumed for the purposes of this paper to be provided independently.

Note that a possible world must include not only tables and chairs, but also more abstract objects such as subjects of discussion, and facts as to who just said what to whom. For practical purposes, it is often helpful to separate out these different aspects of worlds in our models of the common ground. For instance, following Heim [Hei82], we might use an assignment function (“sequence”) for keeping track of entities that are under discussion, so that a model of the common ground (or, for Heim, a “context”) becomes a set of world assignment pairs. I take it that the main reason for making such a move is convenience: a theory of discourse referents will become long-winded if it includes a complete stipulation of those properties in the world which mean that a certain object is under discussion. It is far easier to represent such facts directly. I will also have cause to diverge from use of worlds alone as the basic building blocks of the theory.

Given that W is the set of possible worlds, a model of the common ground is an element of $CG = \mathcal{P}(W)$. A transition is an element of $T = CG \times CG$. A meaning is a set of such transitions, thus an element of $M = \mathcal{P}(T)$.⁴ All this will be familiar to students of dynamic semantics.

Now we want to move on to the question of how we should represent uncertainty about which transition the speaker intended. Easy. If hearers are uncertain about which was intended, then they must maintain a set of alternative transitions, i.e. an object of the same type as an utterance meaning. The members of this set will be referred to as the *alternative utterance transitions* (of U , for an utterance U).⁵

Next, a way of representing preferences about which transition was intended is needed. I will use what I take to be a non-committal approach: the preferences will be represented as a partial ordering over the elements of the set of alternative utterance transitions. The partial ordering corresponding to U , what I will refer to as the *utterance transition ordering*, will be written \leq_U . This partial ordering will be encoded as a set of pairs of alternative utterance transitions, i.e. a member of $UTO = \mathcal{P}(T \times T)$, modulo some further constraints such as transitivity and reflexivity.

Here, then, is the model of interpretation I wish to propose. Compositional interpretation provides a set of utterance meanings defining alternative transitions. A number of different preferential constraints together induce an utterance transition ordering. These constraints will include (1) what initial common ground is most plausibly being assumed, given previous discourse, (2) what might be useful and coherent contributions to the conversation, and (3) any other constraints on the relative plausibility of different interpretations. On the basis of the utterance transition ordering, the hearer forms a partially ordered set of output common grounds, possibly somehow cutting sufficiently Dis-preferred transitions out of contention. The partially ordered set of output common grounds (1) models the information state of the hearer after processing the utterance, and (2) helps induce the utterance transition ordering for the following utterance, i.e., it serves as the next input.

9 The first approximation.

Doubtless there are many ways in which the model I have proposed is only approximate. The one I have in mind is quite specific. The above model reduces pragmatics to a filtering operation on alternative transitions provided by compositional interpretation. This means that the preferred common grounds after update with an utterance will only contain information that either was taken to be assumed before the utterance was made, or was compositionally derived.

⁴We might well wish to constrain meanings further, e.g. to be partial functions rather than arbitrary relations, or to include only information updates.

⁵Note that here there is an alternative: I could have modeled not the hearer's uncertainty about the intended update, but about the intended meaning. Such uncertainty would obviously be represented as a set of meanings. The issue of whether we really need to reason about alternative transitions or about alternative meanings appears to me to be a complex one.

A more sophisticated model would weaken the link between the compositionally derived possible output common grounds, and the common grounds resulting from the application of pragmatics. In this way, further information could be introduced.

Consider, for example, what are known to Gricean pragmaticists as *clausal* implicatures. Suppose a speaker utters a disjunction, of the form “A or B”. In many contexts, if the speaker had been in possession of evidence that A was the case, then the speaker would have just said “A”. Thus an implicature is triggered that the speaker does not know A to be the case, and, similarly, that the speaker does not know B to be the case. These are *clausal* implicatures, and they depend on an assumption that the speaker is complying with Grice’s *Maxim of Quantity*.

Such implicatures will not be part of the information in the output of the meaning relation provided by compositional interpretation. We must conclude that a distinction is needed, to put it in traditional terms, between sentence meaning and utterance meaning. A sentence meaning is the set of transitions provided by compositional interpretation. The utterance transitions may incorporate additional pragmatic information. If a pair $\langle I, O \rangle$ is in the extension of the sentence meaning, then there will be pairs $\langle I, O' \rangle$ in the utterance meaning, where O' is any subset of O , i.e. O' contains more information than O . Furthermore, the utterance transition ordering will typically involve a preference for output states which include information contributed by a given implicature above output states that do not include this information. So, by default, speakers will prefer outputs including implicatures than those that do not. This is the property that will make (conversational) implicatures defeasible inferences, one of their defining properties.

10 Pragmatics to a second approximation and beyond.

Allowing utterance outputs to contain more information than sentence outputs might be called *pragmatics to a second approximation*. The second approximation is the assumption that utterance meanings involve communication of strictly more information than sentence meanings. However, people may speak figuratively. This means that there are few absolute constraints on the relationship between the compositionally derived sentence transition meaning and the utterance meaning. In the final analysis, almost all such constraints are preferential.

For example, if I say “The cat is on the mat”, I might mean what could have been expressed literally as “Watch out, he’s got a gun.” I might, but I probably do not. The set of utterance transitions corresponding to an utterance of “The cat is on the mat” include transitions which take the output to states where some individual is established to have a gun, but these transitions, especially in contexts where there is a salient cat whose location is of interest, are very low on the utterance transition ordering. In some other contexts, although I find it hard to imagine exactly what they are, such bizarre transitions may be higher

on the ordering.

Allowing sentence outputs to differ from utterance outputs moves us beyond the first approximation. And removing all absolute constraints on the link between sentence and utterance output takes us a level further. But I must caution that while I the framework of itself seems to me to be of great generality, I have not yet determined how the story I have tried to tell in this section and the last might be implemented with any elegance. And, furthermore, the particular distinction made between sentence and utterance meaning makes me feel a little queezy. If you find a cure for my queeziness, be sure to let me know. In the meantime, I would like to suggest that there is plenty of first-approximation work to be done in Transition Preference Pragmatics. I suggest that treatment of implicature and figurative language, although central to the success of the enterprise, be reserved until a later date.

11 Example: presupposition accommodation.

The perspective on pragmatics presented in this paper is a generalization of an approach that I have developed to presupposition accommodation [Bea95].

First, and in this I adapt from work of Stalnaker, Karttunen and Heim, the meaning of a sentence is a partial function from contexts (common grounds) to contexts. Presuppositions place constraints on the input context, and hence produce this partiality. For instance, “Mary didn’t realize that her car was on fire” presupposes that Mary had a car and that it was on fire. Thus the meaning of the sentence contains only transitions in which the input context supports these propositions.

Second, a hearer is typically uncertain about what common ground is being assumed by the speaker, although common sense induces a partial ordering over alternative common grounds. I refer to this ordering as an information ordering. For example, it is possible, although unlikely, that I or any other speaker will assume that we commonly know that Mary’s car was on fire. Given that it is implausible that a speaker will make such an assumption, contexts which support that proposition will be relatively low on a hearer’s information ordering. However, suppose the speaker utters “Mary didn’t realize that her car was on fire”. This will effect a transition on the elements of the information ordering which only maps contexts supporting the proposition that Mary’s car was on fire to new contexts, and eliminates others. The resulting information ordering will only contain contexts supporting the presupposed proposition, hence yielding the effect known as *accommodation*.

There are a number of ways in which an information ordering, $\pi \in \mathcal{P}(\text{CG} \times \text{CG})$, can be used to induce constraints on an utterance transition ordering, $\sigma \in \text{UTO}$ defined over a set of utterance transitions $\tau \in \mathcal{P}(\text{CG} \times \text{CG})$. Here is an obvious one, which I present not as a theory which I wish to justify, but as an indication of the form such a theory might take: *If $\langle I, I' \rangle \notin \pi$, $\langle I, O \rangle \in \tau$, $\langle I', O' \rangle \in \tau$ then $\langle \langle I, O \rangle, \langle I', O' \rangle \rangle \notin \sigma$.*

This says that if two transitions are amongst the alternative utterance transitions, and if the initial common ground of the first transition is not at least as

plausible as the initial common ground of the second transition, then the first transition is not at least as plausible as the second transition.

Using an analysis of this form, it would be possible to duplicate the main results of my analysis of presupposition accommodation in the current setting. Consider, for example, the sentence “When handling a Peruvian cricket, use goggles, and be careful not to touch the red and green spiky bits”. The expression “the red and green spiky bits” presupposes the existence of such bits. But what spiky bits are these, and where are they located? Potentially, the red and green spiky bits might be a permanent feature of the world independent of any particular episode of cricket exploration. Or, these bits might be parts of Peruvian crickets. They might also belong to the goggles, or perhaps to some other entity, such as the plants on which these crickets might hypothetically live. Maybe red and green spiky bits are Peruvian cricket food?

Of all of these possibilities, the thought that comes to mind most readily when hearing the sentence is that the spiky bits are part of the cricket. In the theory I have presented elsewhere, this fact would be represented as a general (although unexplained) property of the information orderings which hearers have. This property would require that (1) there are contexts in the ordering in which every Peruvian cricket has red and green spikes, and (2) some of these contexts are ranked higher in the ordering than any contexts in which goggles have red and green spiky bits, in which the speaker’s world contains some red and green spiky bits, in which Peruvian grasshopper vegetation has red and green spiky bits, or in which red and green spiky bits are Peruvian grasshopper food, etc. Given that information orderings are like this, accommodation of the proposition that Peruvian grasshoppers have red and green spiky bits would be predicted.

In the current setting, the information ordering would not be used directly, but instead used to rank transitions corresponding to meanings of the original sentence. In particular, the resulting ranking would need to have the property that transitions in which the input was a context in which every Peruvian cricket has red and green spikes were ranked higher than all transitions in which the input involved goggles with red and green spiky bits, and so on. The story would be essentially the same as I have told before, but told in a marginally different way.

12 Example: anaphora resolution.

In work in progress, I have been exploring how the Centering theory of anaphora resolution [GJW95, BFP87] can be combined with a dynamic semantics.⁶ The crucial observation here is that well-known accounts of anaphora resolution including [GS86, Si83] in addition to centering, are all in essence dynamic theories.

⁶I first presented a unified account of centering and dynamic semantics at ESSLLI X, Saarbrücken, 1998. When I later discussed the account with my colleague, Stanley Peters, he asked me whether it could be seen as part of any more general theory of interpretation. The answer is the text you are now reading, which relates the account of anaphora resolution to my earlier work on presupposition, and to ambiguity resolution and pragmatic inference in general.

In fact Johan van Benthem and accomplices make this observation with respect to Grosz and Sidner’s model in [vBMV97].

Two properties justify terming NLP models of anaphora resolution *dynamic*. First, there is a notion of a context or information state which is updated from one sentence to the next. Second, the interpretation of each sentence is dependent on the information state resulting from update by the previous sentence.

In the case of Centering, the context is quite simple. It consists of a combination of a value for a single variable, the *backward looking center*, and a sequence of variables, the *forward looking center list*. The *backward looking center* can be thought of as a discourse referent corresponding to the sentence topic. The *forward looking center list* is a list of potential antecedents for anaphors in the next sentence, much like a Heimian *sequence*. To combine a dynamic semantics with Centering, the first step is to define a notion of information state with sufficient structure that both semantic information and centering information can be extracted. This is straightforward, but I will not give definitions here.

In Centering theory, the process of interpreting the anaphors in a sentence consists of a two stage process. Firstly the alternative resolution options are generated. Then a ranking is defined on those options, and the optimal resolution is chosen.

For example, consider the following discourse segment: “Mary was busy yesterday. She took Jane to the park. Then she played football.” After updating with the second sentence, the *backward looking center* will be Mary, and the *forward looking center list* can be informally represented as ⟨Mary, Jane, the park⟩. Ignoring some of the details of Centering, we can conceive of two alternative resolutions for the pronoun in the third sentence. Identifying Mary as the referent would result in the topic (backward looking center) of the third sentence being the same as that of the second. On the other hand, if Jane were the referent, the topic would be different. Centering ranks transitions in which the topic is constant across sentences above transitions in which the topic is shifted. So Mary would be correctly identified as the preferred referent of the third sentence’s pronoun.

I will not detail here a full implementation of Centering in Transition Preference Pragmatics. But I hope it is clear that Centering is a special case of the abstract framework which I have sketched in this paper.

13 Discussion

This paper is programmatic. Some might prefer “superficial”. I have suggested a level at which a wide variety of pragmatic constraints may be stated. But I have not said how we are to combine constraints from differing sources. This is the really important question to answer in future work. In the remainder of this section, I will hint at the scale of the problem.

Suppose that a particular constraint on preferred discourse relation clashes with a constraint derived from our common-sense knowledge on what is a plausible fact about the world. For instance, consider the following discourse seg-

ment: “Mary went for a walk with her puppy, Claudia. There was a bone in her mouth.” Centering theory predicts unambiguously that Mary had a bone in her mouth, whereas common sense tells us it was Claudia.

This might lead us to hypothesize that hearers always prefer the meaning which is *a priori* most probable, regardless of what consequences this might have for text coherence. But such a hypothesis cannot be correct. Consider, for example, the following discourse segment: “Mary recently went water-skiing with her mother, and is a keen follower of professional wrestling. She is 84 years old.” Clearly, common sense knowledge of the world makes it *a priori* more likely that Mary’s mother is 84 years old than that Mary is. However, we interpret the text to mean that Mary is 84 years old. This illustrates that sometimes text-coherence constraints win out over common sense in our interpretation of a text.

How do we know, in general, which sort of constraint is most important? Given that it is not the case that one type of constraint is always stronger than the other, should a clash of preferences between these sources result in an absence of any overall preference? I do not believe that we can expect to find general answers to such questions: they may be the wrong questions. Or perhaps they are just the wrong constraints.

The constraints on text coherence which are embodied in centering theory say things like “prefer the topic to be constant across sentences”. Such constraints are too low-level to ever be compared directly with constraints on how the world is. Perhaps we should instead aim to derive topic-constancy constraints from some higher level statement of Grice’s maxim of manner. It might make sense to weigh up the chances of the speaker failing to follow the maxim of manner against the chances of Mary being 84. At least, this might make more sense than weighing up the chances of the speaker failing to maintain topics across sentences against the chances of Mary being 84.

Whether or not the maxim of manner is the right sort of constraint to have in our system, I think there is a more general point to be made. Many general properties of discourse interpretation can be found. It would be naive to gather such properties into a great heap, and expect some marvelous mechanical constraint solver to use them to classify texts into good and bad. For general properties are not necessarily defining properties. They may often be merely symptomatic of deeper generalizations. The result is that generalizations drawn from differing sources often appear incommensurable.

When we find the right level to state generalizations about text coherence, the result, I optimistically believe, will be constraints which are commensurable with constraints on what the world is like. We will have a theory in which common sense knowledge and knowledge of discourse may be blended seamlessly together. The result, in other words, will be a pragmatic theory in which different sources of information define a single set of transition preferences.

References

- [Bea95] Beaver, D., 1995. *Presupposition and Assertion in Dynamic Semantics*, PhD Dissertation, University of Edinburgh. To appear, *Studies in Logic, Language and Information*, CSLI, Stanford/CUP Cambridge.
- [vBMV97] van Benthem, J., R. Muskens and A. Visser, 1997. “Dynamics”, in van Benthem, J. and ter Meulen, A., *The Handbook of Logic and Language*, Elsevier.
- [BFP87] Brennan, S., M. Friedman and C. Pollard, 1987, “A Centering Approach to Pronouns”. In *Proceedings of the 25th Annual Meeting of the Association for Computational Linguistics*, Cambridge, Mass.
- [GS91a] Groenendijk J. and M. Stokhof, 1991. “Dynamic Predicate Logic”, *Linguistics and Philosophy* 14, pp. 39–100.
- [GS91b] Groenendijk J. and M. Stokhof, 1991. *Dynamic Montague Grammar*, in J. Groenendijk, M. Stokhof and D. Beaver (eds.), DYANA deliverable R2.2a *Quantification and Anaphora I*, Center for Cognitive Science, University of Edinburgh.
- [GJW83] Grosz, B., A. Joshi and S. Weinstein, 1983. “Providing a Unified Account of Definite Noun Phrases in Discourse”. In *Proceedings of the 21st Annual Meeting of the Association for Computational Linguistics*, Cambridge, Mass, pp. 44–49.
- [GJW95] Grosz, B., A. Joshi and S. Weinstein, 1995. “Centering: A Framework for Modeling the Local Coherence of Discourse”, *Computational Linguistics* 21:2, pp. 203–226.
- [GS86] Grosz, B. and C. Sidner, 1986. “Attention, Intentions and the Structure of Discourse”. In *Computational Linguistics* 12, pp. 275–204.
- [Hei82] Heim, I., 1982. *On the semantics of Definite and Indefinite Noun Phrases*, PhD dissertation, University of Amherst.
- [KRe93] Kamp, H. and U. Reyle, 1993. *From Discourse to Logic*, Dordrecht, Kluwer.
- [Mu96] Muskens, R., 1996. *Combining Montague Semantics and Discourse Representation Theory*, *Linguistics and Philosophy* 19, pp. 143–186.
- [Si83] Sidner, C., 1983. “Focusing in the Comprehension of Definite Anaphora”. In Brady, M. and R. Merwick (eds.), *Computational Models of Discourse*, MIT Press.