

INFERENCE, ELLIPSIS AND DEACCENTING

DANIEL HARDT

Computational Linguistics
Copenhagen Business School
dh@id.cbs.dk

It is proposed that inference is available for ellipsis interpretation, only when triggered by a semantically visible violation. This captures well-known observations about available inferences in VP ellipsis, and new observations about unavailable inferences in sluicing. Furthermore, it is argued that inference in deaccenting is governed by the same triggering constraint.

1. Introduction

Paul Grice, in a famous pronouncement, enjoins language users to “avoid unnecessary prolixity” (Grice 1975). Necessary or not, redundant material is ubiquitous in natural language, its presence typically signalled by deaccenting or ellipsis. Naturally enough, deaccenting and ellipsis are infelicitous if the relevant material is not understood to be redundant. This is the minimal condition on deaccenting and ellipsis, and, ideally, this requirement would suffice to uniformly characterize its distribution.

One apparent problem with this ideal view involves inference: while inference is clearly involved in determining redundancy with *deaccenting*, this does not appear to be the case with *ellipsis*, as illustrated by the following examples from Rooth 1992 (*italics* indicates deaccenting.):

- (1) First someone told Mary about the budget cuts, then SUE *heard about them*.
- (2) First someone told Mary about the budget cuts, then SUE did.

The Verb Phrase *heard about them* is deaccented felicitously in (1), even though it has not appeared previously. Rooth argues that this relies on the inference *if someone tells Mary about budget cuts then Mary hears about them*. This inference is not available for (2), which does not have the reading *Sue heard about them*. Based on such facts, Rooth argues that deaccenting and ellipsis cannot be unified.

In fact, there is an early argument, due to Webber 1978, that inference is indeed required for ellipsis interpretation. Rooth does not address this argument, which

has been ignored in much of the subsequent theoretical literature. Certain recent authors have echoed Webber's argument (Fox 2000; Merchant 2001). Like Webber, these authors point to examples of ellipsis where the correct interpretation requires an inference.

These authors argue convincingly that there is good reason to take inference seriously in ellipsis. I will argue, however, that what is most intriguing about inference is its *absence* – it has not been previously observed that there are various cases of ellipsis and deaccenting where inference is systematically unavailable. I will argue that this reflects a fundamental, general condition on inference in interpretation: it must be *triggered* by a violation. I will argue, furthermore, that inference in deaccenting is subject to the same constraint. The evidence, then, points to a unified theory of deaccenting and ellipsis; appearances to the contrary can be traced to the simple fact that deaccenting can more easily trigger inference, because it involves overt lexical material.

In what follows, I first present the argument for inference in ellipsis, as given by Webber. Next, I present various cases where inference is unavailable in ellipsis: these primarily involve case-matching effects in sluicing. I then show that inference is also unavailable for deaccenting, in an example involving scope parallelism. I give a general statement of the triggering condition, and I argue that this uniformly accounts for inference in ellipsis and deaccenting.

2. Inference in Ellipsis: the Argument

Webber 1978 argues that inference is sometimes required in ellipsis interpretation, as in the following example of Verb Phrase (VP) Ellipsis:

(3) Irv and Mary wanted to dance together, but Mary couldn't, because her husband was there.

Webber notes that the desired reading is not *Mary couldn't dance together*, but rather, *Mary couldn't dance with Irv*. She argues that this results from the following inference:

(4) Irv and Mary wanted to dance together \Rightarrow Mary wanted to dance with Irv.

One might wonder if inference is indeed required: an alternative explanation is that what is elided is not the entire VP *dance together*, but just the verb *dance*, giving *Irv and Mary wanted to dance together, but Mary couldn't dance*. This is then interpreted as *dance with Irv* because of context.

While this appears to be a plausible alternative account for (4), this will not account for examples like (5), as pointed out in Hardt 1993:

(5) Martha and Irv had planned to nominate each other, but Martha couldn't, because of her political obligations.

Here, the reading is “Martha couldn’t nominate Irv”. Unlike *dance*, the verb *nominate* cannot appear without its object.

Webber notes that “there seem to be no hard and fast rules delimiting the class of productive inferences relative to verb phrase ellipsis.” (p. 4-38) However, there is a crucial factor in example (4) which is not noted by Webber – the interpretation without inference is unacceptable, involving, as it does, an agreement violation between the singular subject *Mary* and the VP *dance together*.

This fact is the key to solving the problem posed by Webber, about how to delimit the class of available inferences in ellipsis: only inferences triggered by violations are possible. To show this, we examine cases where inferences are systematically unavailable.

3. Missing Inferences

We have seen that inference is possible in ellipsis when triggered by a violation. In this section, I will argue that inference is not available in the absence of a violation. Consider the following variant of (4), where the VP ellipsis has a plural subject:

(6) Irv and Martha wanted to dance together, but Tom and Susan didn’t want to.

Here, the default, or non-inference reading is acceptable: *Tom and Susan didn’t want to dance together*. In this case, the inference is no longer available: this variant cannot mean *Tom and Susan didn’t want to dance with Irv*.

We turn now to sluicing. As first observed by Ross 1967 sluicing is subject to a case-matching requirement, as illustrated by the following example:

(7) Er will jemandem schmeicheln, aber sie wissen nicht
He wants someone.DAT flatter, but they know not
wem.
who.DAT.
He wants to flatter someone, but they don’t know whom

(8) Er will jemandem loben, aber sie wissen nicht wen.
He wants someone.ACC flatter, but they know not who.ACC.
He wants to flatter someone, but they don’t know whom

As Ross points out, the wh-word in (7) must take dative case, which is the case assigned by *flatter*. In (8), the wh-word is required to take accusative case, which is the case assigned by *praise*.

This *case matching* constraint in sluicing is stated by Merchant 2001 as follows:

- *The sluiced wh-phrase must bear the case that its correlate bears.*

What has not previously been noticed is that case matching fails to interact with inference. This is illustrated by (9):

(9) John traf jemanden, aber ich weiss nicht wen [John traf x] /*wer [traf John].
 John met someone, but I don't know who [John met x] / who [met John].

Here, the wh-word must have accusative case ("wen"), the nominative "wer" gives rise to ill-formedness. What is interesting here is that "met" is a symmetric predicate: "John met someone" is semantically equivalent to "Someone met John", so on Merchant's semantic condition, the nominative "wer [traf John]" should also be permitted.

I conclude that inference is possible in ellipsis if and only if there is a semantically visible violation. In the following section, I give evidence that the same is true with deaccenting.

4. Missing Inferences in Deaccenting

Consider the following sentences:

(10) A doctor saw every patient. A NURSE *saw every patient, too*.
 (11) A doctor saw every patient. A NURSE *saw many patients, too*.

(*Italics* indicates deaccenting.) It is widely acknowledged that examples like (10) are subject to a scope parallelism constraint – that is, *a doctor* takes wide scope in the first sentence only if *a nurse* also takes wide scope in the second. A similar scope parallelism effect can be observed in (11). Note that in this case, the deaccented material does not match – *every* has changed to *many*. This is the result of the following intervening inference, which I call the *every-many* inference:

(A doctor x) (every patient y) x saw y. → (A doctor x) (many patient y) x saw y.
 (This follows if one assumes that there are many patients.)

The situation here is exactly analogous to the problem with case-matching: the possibility of inference threatens to undermine the scope parallelism constraint, since there is an inferential relation between the two scopes of *exists* and *every*. Consider the following LF representations for (10):

(12) a. (A doctor x) (every patient y) x saw y.
 b. (every patient y) (A NURSE x) x saw y.

The intention of the scope parallelism constraint is to rule out this representation. But parallelism could be satisfied by the following intervening inference: (A doctor x) (every patient y) x saw y → (every patient y) (A doctor x) x saw y. To maintain

scope parallelism, such scope-shifting inferences must be ruled out, although in other cases, like (11), scope parallelism determination must permit intervening inferences like the *every - many* inference. How is the scope-shifting inference ruled out, while the *every-many* inference is permitted?

In previous work Hardt 2005 I have suggested that inferences be restricted by a *submodel* constraint. However, it is not clear that there is evidence for such a constraint. What is crucial in the current example is that the scope shifting inference is not required to satisfy parallelism. Thus the triggering constraint rules it out.

I turn now to a general statement of the triggering constraint.

5. Triggering Constraint

I propose the following general account of inference and interpretation: for a given discourse *D*, we produce a *default LF L*. If *L* violates no *semantically visible* constraints, it is the preferred interpretation. If *L* does violate one or more constraints, inferences can be performed to derive an alternative interpretation *L'*. *L'* is a potential interpretation of *D* if it avoids the constraint violations. If there are several such alternatives, those LF's *closest* to *L* are preferred. The closeness relation is defined with respect to entailment; if *A* entails *B* and *B* entails *C*, then *B* is closer to *A* than *C*.

I will not give any general characterization of *semantically visible* constraints: in this paper, I will rely on what I take to be completely standard conceptions, namely, that the parallelism requirement is semantically visible, as are agreement violations such as observed in (3). On the other hand, the *case-matching* constraints in sluicing are not semantically visible.

The triggering constraint proposed here has roots in the proposal of Fox 2000, where it is proposed that inference in ellipsis interpretation must be triggered. There are important differences, however. Most important is the triggering condition – on my proposal, semantically visible violations trigger inference, while for Fox the triggering condition involves a mismatch of lexical items. That is, Fox claims that inference is triggered by deaccented overt material in the ellipsis clause that is not present in the antecedent clause. This proposal is far less general than the current proposal, which is simply that semantically visible violations trigger inference. In my view, Fox's proposal lacks the intuitive motivation of my proposal. Furthermore, many cases of triggered inferences, including all those discussed in this paper, are beyond the scope of Fox's proposal.

6. Conclusions

Ellipsis and deaccenting signal redundancy, and thus they are subject to a requirement that they apply to material that is understood as redundant. On the one hand, it is natural to suppose that the determination of redundancy proceed with the help

of inferential reasoning. On the other hand, inferences can be computationally expensive operations, and dependence on inference might be difficult to reconcile with the demands of online processing. The evidence presented in this paper leads to the conclusion that inference is generally available, but in a very limited way – it is only available when triggered by a semantically visible violation.

Acknowledgements

This work is supported by a grant from the Danish Research Council for the Humanities.

Bibliography

Chung, S.: 2005, *A Look at Sluicing in Chamorro (and Germanic)*, Handout, University of Santa Cruz

Fox, D.: 2000, *Economy and Semantic Interpretation*, MIT Press

Grice, P.: 1975, Logic and conversation, in P. Cole and J. Morgan (eds.), *Syntax and Semantics 3*, pp 41–58, Academic Press

Hardt, D.: 1993, *Verb Phrase Ellipsis: Form, Meaning, and Processing*, Ph.D. thesis, University of Pennsylvania

Hardt, D.: 2005, Natural language inference as triggered submodel search, in *Proceedings, FLAIRS05*, Clearwater Beach, Florida

Merchant, J.: 2001, *The syntax of silence: Sluicing, islands, and identity in ellipsis*, Oxford

Rooth, M.: 1992, Ellipsis redundancy and reduction redundancy, in *Proceedings of the Stuttgart Ellipsis Workshop*, Stuttgart, Germany

Ross, H.: 1967, *Constraints on Variables in Syntax*, MIT Department of Linguistics and Philosophy

Webber, B. L.: 1978, *A Formal Approach to Discourse Anaphora*, Ph.D. thesis, Harvard University, (Published 1979 by Garland Publishing, New York)