Prosodic prominence shifts are anaphoric

Jeffrey Klassen [1], Michael Wagner *

Department of Linguistics, McGill University, 1085 Doctor Penfield Avenue, Montréal, Québec H3A 1A7, Canada

ABSTRACT

This paper presents evidence that shifts in prosodic prominence are anaphoric and require a contextually salient antecedent, similar to pronouns. The argument is based on a series of experiments looking at prosodic optionality in dialogues in which there are multiple potential antecedents embedded in a contextually salient coordinated structure. By looking at the interaction with adverbs that restrict the choice of antecedent, we show that the observed prosodic variability reveals different anaphoric choices, and hence different speaker intentions. The results are incompatible with the hypothesis that prominence shifts can be explained purely in reference to low-level facilitation due to repetition of the linguistic structure or accessibility of it referent, and are not reducible to existing accounts of prominence in terms of predictability.

© 2016 Elsevier Inc. All rights reserved.

Prominence shifts and coordinated antecedents

In English, the prosodic prominence pattern of a sentence is affected in systematic ways by discourse context. Consider the following dialogues:

(1) a. A: What happened while the kids went for a swim?
    B: JOLENE pitched the TENT.
   b. A: Who pitched the tent?
    B: JOLENE pitched the tent.

Whereas in the first discourse, both the words *Jolene* and *tent* typically carry an accent and receive prosodic prominence, in the second discourse, only the word *Jolene* is likely to be accented, and the material in the VP *pitched the tent* will likely be prosodically reduced and perceived as being less prominent than the subject.

Such contextual effects on prosodic prominence reflect what part of the information in a sentence constitutes new and given information respectively. Earlier studies have shown that placing prominence on given material as well as failing to place prominence on new material can incur a processing cost: Both accenting given material and deaccenting new material lead to increased processing time in a variety of languages (Arnold, 2008; Birch & Clifton, 1995, 2002; Bock & Mazzella, 1983; Braun & Tagliapietra, 2010; Cutler, 1990; Cutler, Dahan, & van Doneselaar, 1997; Dahan, Tanenhaus, & Chambers, 2002; Nooteboom & Kruyt, 1987; Terken, 1984; Terken & Nooteboom, 1987).[2] Failing to deaccent or accenting old information also has been shown to have measurable electrophysiological repercussions in English, French, German, and Japanese (Baumann & Schumacher, 2011; Heim & Alter, 2006, 2007; Hruska & Alter, 2004; Ito & Garnsey,

[2] Arnold (e.g. 2008) finds evidence that at least in certain circumstances, there is an asymmetry such that deaccenting new information leads to infelicity while accenting given information is relatively more acceptable. One reason for this asymmetry may be that it might be easier to accommodate a contrasting alternative in the case where a contextually given constituent is accented, compared to accommodating an antecedent for marking as given a constituent that is not actually salient in the context.

2004; Johnson, Breen, Clifton, & Morris, 2003; Magne et al., 2005; Schumacher & Baumann, 2010; Toepel, Pannekamp, & Alter, 2007; Wang, Bastiaansen, Yang, & Hagoort, 2011).

The prosodic rendition of a particular constituent and its prominence relative to the other constituents in the same sentence therefore encodes information about the discourse status of a constituent, which lead Fowler and Housum (1987, & many others) to conclude that the prosodic rendition of an utterance is *anaphoric* to the context in a very broad sense of the word, where any utterance that reflects in some way what is salient in the context counts as anaphoric.

This much is uncontroversial, but what is controversial is the proper characterization of the mechanism responsible for such contextual effects on prominence. On the one hand, activation-based views of prosodic reduction look upon the relationship between context and prosodic prominence as a direct consequence of the prior activation of the relevant linguistic constituents or their referents in discourse: If a linguistic expression or its referent is already activated or 'accessible,' then its prosodic prominence will be reduced (Arnold & Watson, 2015; Terken, 1984; Terken & Hirschberg, 1994; Watson, 2010, i.a.). Semantic accounts, on the other hand, treat prominence shifts on par with pronominal anaphors, as a way to encode a presupposition about what information is salient in the context. This paper provides evidence that an account purely in terms of activation, predictability, or accessibility cannot capture the distribution of prosodic prominence, and a notion of anaphoric contrast is needed.

Accessibility-based accounts

One group of theories on prosodic reduction maintain that prominence can be explained by prior activation of either constituents or their referent/denotation. The underlying idea is that if processing a constituent is 'easy' or accessible, it will be reduced, and if it's 'hard', it will be not be, and might be hyper-articulated (Clopper & Turnbull, submitted for publication, for review). One dimension on which accessibility-based accounts vary is whether reduction reveals information about accessibility to the speaker (speaker-oriented accounts), or whether the speaker reduces constituents based on their assumptions about what will be easy or difficult for the listener (listener-oriented accounts). Another dimension along which they vary is in the assumed mechanism underlying the activation effect; that is, based on why exactly a constituent is supposed to be 'easy' or 'hard'. In the following, we review the most central hypotheses about what types of contextual accessibility can cause prosodic reduction.

One way in which a linguistic expression can be contextually accessible is if it itself was recently used: We will call this 'repetition-accessibility'. Linguistic reuse has been shown to be an important factor in the allocation of prominence in Fowler and Housum (1987), Terken (1984, 1985), and many other studies. Lam and Watson (2010) argue that prior articulation of a word by a speaker will have a strong reductive effect on a repeated instance of the same word. Arnold, Kahn, and Pancani (2012) and Kahn and Arnold (2015) presented evidence that prior audition of a word is sufficient to induce prosodic reduction of a subsequent articulation, but the mechanism to which they attribute this reduction is also production facilitation. Such a speaker-oriented account of repetition effects is supported by findings that suggest that speakers' accent placement does not take into account the listener's knowledge state (Bard et al., 2000). See Arnold and Watson (2015) for a recent summary of relevant research. Under this low-level production-facilitation-based view, the prominence shift observed in (1-b) would be seen as a direct consequence of the previous mention of the constituent *pitched the tent*.

Another class of contextual accessibility factors discussed in the literature involve prior activation at the level of meaning, and we will refer to this as 'referential-accessibility'. For example, linguistic expressions that refer to discourse referents that are already contextually salient have been argued to be prosodically reduced in Chafe (1974), as well as in centering theory (Grosz, Joshi, & Weinstein, 1983; Terken, 1984; Terken & Hirschberg, 1994). The idea is that material that refers to individuals or predicates that are already introduced as referents into the current discourse are more likely to be realized in reduced form (e.g., as a personal pronoun or the predicate pronoun *did* involved in verb phrase ellipsis, or through prosodic reduction). According to this line of work, the prominence shift observed in (1-b) could be seen as a consequence of the prior activation of the meaning of the constituent *pitched the tent*.

Many recent studies on phonetic reduction have focused on various notions of contextual predictability (Arnold, 1998; Aylett & Turk, 2004; Bell, Brenier, Gregory, Girand, & Jurafsky, 2009; Calhoun, 2010; Jaeger, 2010; Lam & Watson, 2010; Pluymaekers, Ernestus, & Baayen, 2005; Schuppler, van Dommelen, Koreman, & Ernestus, 2012; Terken, 1984, i.a.). Context can affect the likelihood of a certain referent or meaning to occur (predictability at the semantic level) and the likelihood that a certain linguistic constituent will be used (predictability at the form level). Proxy-measures for predictability used in the literature include previous mention (Aylett & Turk, 2004); bigram or n-gram frequency or conditional probability of syllables (Aylett & Turk, 2006) or words (Bell et al., 2009, 2003; Hale, 2003; Jurafsky, Bell, Gregory, & Raymond, 2001; Levy & Jaeger, 2006; Pan & Hirschberg, 2000; Pluymaekers et al., 2005); predictability of syntactic constructions (Gahl & Garnsey, 2004); or cloze-probability (Clopper & Pierrehumbert, 2008; Lieberman, 1963; Hunnicutt, 1987). There have also been studies that directly manipulate probabilities within an experimental design (Lam & Watson, 2010; Terken, 1984), although in those studies predictability effects are small when compared to the effects of repetition.

Accounts that try to explain such predictability effects vary in the mechanism they take to be responsible for these effects: Predictable constituents might be faster and easier to process for the speaker, which, downstream, could lead to phonetic reduction effects; or speakers could adjust the level of signal they produce in response to how difficult they estimate it to be for listeners to retrieve (see Clopper & Turnbull, submitted for publication, for an overview).