



UvA-DARE (Digital Academic Repository)

Discourse-based lexical anticipation : the nature and contextual basis of predictions in language comprehension

Otten, M.

Publication date
2008

[Link to publication](#)

Citation for published version (APA):

Otten, M. (2008). *Discourse-based lexical anticipation : the nature and contextual basis of predictions in language comprehension*.

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.



Chapter 1

Introduction

“Each word is doubly awakened; once from without by the lips of the talker, but already before that from within by the premonitory processes irradiating from the previous words, and by the dim arousal of all processes that are connected with the 'topic' of the talk.”

[William James, *The Principles of Psychology*, 1890]

Prediction in language comprehension

Language comprehension is highly incremental: when we hear or read a sentence or a combination of sentences, we immediately and continuously extract meaning. Concurrent use of the extracted meaning to anticipate how the sentence or story will continue could benefit the comprehender. If the new, continuously incoming information matches the expectancy, comprehension can be limited to rudimentary processing. If the new information does not match the prediction, this can provide the language comprehension system with an early ‘flag’ so that attention can immediately be directed to processing this unexpected part of the sentence.

But can we indeed make such predictions? Our own every-day experience shows that we can sometimes finish other peoples sentences, if our conversational partner falters mid-sentence for some reason. But, do these sentence-completions also occur when time is limited, concurrent with our normal comprehension process? After reading a sentence like “*The UPS man delivered the...*” the word *package* is processed more easily than the less expected word *box*: it requires shorter naming latencies (Hess, Foss, & Carroll, 1995; Traxler & Foss, 2000) and is read faster (Cook & Myers, 2004; Morris, 1994). These effects could result from the fact that the reader or listener has pre-activated the word *package*, based on the constraining sentence that they have just read. However, these effects could also occur simply because the meaning of *package*, once unlocked by having seen the word, is more easily integrated in the context than *box*, for example because *package* fits the “UPS delivery”-scenario better.

Using innovative experimental designs, several researchers have shown in recent years that people sometimes indeed pre-activate predictable words that are likely to follow. In an experiment by Kamide, Altmann and Haywood (2003) participants were looking at a scene such as the one depicted in Figure 1.1. When they heard the sentence “*The woman will spread ...*” their gaze shifted to the only spreadable object in the display, namely the butter. Importantly, this happened before they had actually heard the word butter. In



Figure 1.1 Example of a display used in (Kamide, Altmann, & Haywood, 2003).

a less constraining context (“*The woman will slide ...*”) no such preference for a specific object was detected. This indicates that the listeners were thinking about butter before the word “butter” had actually been pronounced, i.e. that they predicted that the word butter would follow. Of course, the visual display already limited the possibilities for the words that could follow. So, even though the results are certainly suggestive, this experiment does not provide definite evidence that prediction occurs in normal language comprehension.

More direct evidence for online prediction in language comprehension was provided by experiments by Van Berkum and colleagues (2005). Most of you who read this introduction are native speakers of Dutch. It could be that after you have read

*De inbreker wist precies waar hij de geheime familiekluis moest zoeken.
Deze bevond zich namelijk achter een...*

you are thinking about the word *schilderij* (see Figure 1.2 for an approximate English translation of this story). At least, an offline cloze test indicates that the majority of readers finishes this story with this word. Van Berkum et al. looked for prediction-related ERP effects not on the noun, but on an adjective preceding the noun. In Dutch, adjectives in indefinite singular noun phrases have a suffix that depends on the arbitrary, lexically memorized gender (Van Berkum, 1996, Ch. 2) of the noun they precede. Adjectives that modify a common-gender noun carry an *-e* suffix (e.g., “*oude boekenkast*”), whereas adjectives modifying a neuter-gender noun are not inflected (e.g., “*oud schilderij*”). Van Berkum et al. (2005) reasoned that to the extent that listeners strongly anticipated a specific noun (*schilderij*), an adjective with a mismatching gender suffix (“*oude*”) would come as an unpleasant surprise, and might as such elicit a differential ERP effect relative to a prediction-consistent control (“*oud*”). Adjectives with a prediction-inconsistent inflection indeed elicited a differential event-related potential (ERP, see Box 1): a small but significant positivity emerging in the 50-250 ms after acoustic onset of the

Box 1: On-line measures of language comprehension

Studying the way listeners finish other peoples faltering sentences during conversation, or how readers finish an incomplete sentence in a pencil-and-paper cloze test seems helpful to examine whether people actually pre-activate a word while they are reading or listening. However, in both cases people are prompted, either by instruction or by social customs, to think about a possible continuation, and they have relatively much time to do so. To evaluate the *ongoing* process of language comprehension and the predictive processes involved, these processes need to be monitored *on-line*. For the studies reported in this thesis we have used two different methods.

Self-paced reading

In this paradigm, participants read sentences or stories word for word (or phrase by phrase, or sentence by sentence). The participant can signal that he/she has finished reading a word by pressing a button, so that the word on the screen is replaced by the next word. This provides a reading time for each individual word. These reading times are thought to reflect the underlying cognitive processes associated with the processing of that word.

Event-related potentials

The electrical activity of the brain can be recorded with electrodes that are placed on the scalp. This record of fluctuating voltage across time is the electroencephalogram (EEG). The event-related potential or ERP is that part of the signal that is related to an external event. At the scalp, such an ERP is much smaller (5-10 μV) than the background EEG (50-100 μV). By averaging several ERPs that are elicited by similar stimuli, most of the background EEG and environmental noise can be averaged out. The rationale is that the event-related activity is time-locked to the appearance of the stimulus, but the background EEG and environmental noise are not. In the averaging process, unrelated noise present from the individual trials will cancel each other. The resulting average ERP waveform reflects the sensory and cognitive brain-processes related to the stimulus, up to the millisecond.

The burglar had no trouble whatsoever to locate the secret family safe. Of course, it was situated behind a...

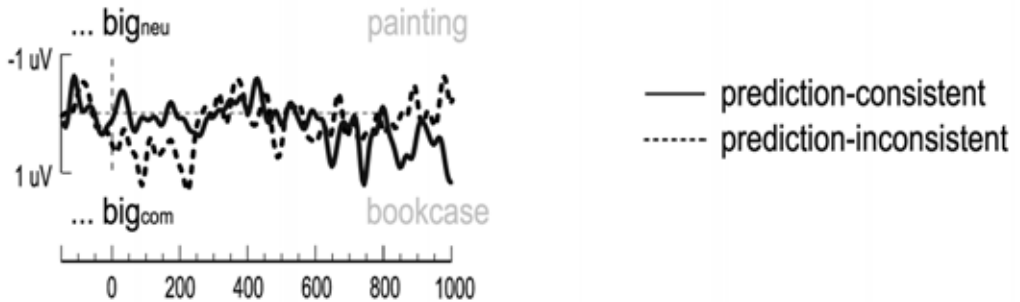


Figure 1.2 ERPs to spoken adjectives with an inflection that was in line (solid line) or inconsistent (dotted line) with the gender of the predictable noun, timelocked to the onset of that inflection (based on data from Van Berkum, Brown, Zwitserlood, Kooijman, & Hagoort, 2005).

inconsistent inflection (see Figure 1.2). This ERP effect occurred before any noun had been presented, at a point in time where both gender inflections were equally correct. Van Berkum et al. therefore took this effect as evidence that their listeners had indeed pre-activated a specific word and its lexical features, like gender, based on discourse information. In a follow-up self-paced-reading study (Van Berkum et al., 2005, Experiment 3), prediction-inconsistent adjectives also slowed down the reading process.

A series of studies by Wicha and colleagues (Wicha, Bates, Moreno, & Kutas, 2003; Wicha, Moreno, & Kutas, 2003, 2004) with Spanish single sentence stimuli, showed that specific lexical prediction is not limited to longer stories or spoken materials. Prediction-inconsistent gender-marked determiners elicited an enhanced negativity between 300-500 ms (spoken sentences: Wicha, Bates et al., 2003; written sentences: Wicha, Moreno et al., 2003). In these two experiments the expected and unexpected articles were followed by line drawings of the expected or an unexpected concept. When all the stimuli were presented in writing, unexpected articles evoked an enhanced positivity between 500 and 700 ms (Wicha et al., 2004). In an experiment that used the fact that in English indefinite determiners differ depending on the initial phoneme of the word that follows (a/ an), DeLong, Urbach and Kutas (2005) found that determiners that were not in line with the expected word elicited a negative shift between 300 and 500 ms compared to the prediction-consistent determiners.

Taken together, these results suggest that we are able to make very specific predictions when we are processing incoming language. The experiments collected in this dissertation explore the nature and extent of these specific lexical predictions.

The contextual basis of prediction

Lexical and scenario-based association

The results reviewed above show that people pre-activate lexical information of words that they deem likely to follow. One relevant process underlying these predictions could be automatic activation of lexical-semantic information by words in the discourse (as already suggested by William James, as you can see in the quote at the beginning of this chapter). According to the spreading activation account (Collins & Loftus, 1975; Dell, 1986; Meyer & Schvaneveldt, 1976; Swinney, 1979) activation feeding from one representation in the lexicon to other representations that share a connection causes priming (i.e. facilitated processing of the related word). *UPS* can thus prime *package* because the two words co-occur frequently, and this is captured in the strength of the connection between them. The greater the strength of the connection, the more activation is fed from one representation to the other when one becomes activated.

There is extensive evidence that lexical association facilitates processing. Words are processed faster and more accurately when they follow a related prime (see Neely, 1991 for a review) and the N400 is attenuated for these related words (cf. Bentin, McCarthy, & Wood, 1985). The priming effects of a related word are also visible when the word pairs are presented as part of a coherent sentence (Camblin, Gordon, & Swaab, 2007; Carroll & Slowiaczek, 1986; Hoeks, Stowe, & Doedens, 2004; Van Petten, 1993; Van Petten, Weckerly, McIsaac, & Kutas, 1997).

Automatic pre-activation is not necessarily based purely on lexical priming. Models of text comprehension and memory suggest that the words in a text can also provide semantic constraints via the activation of related information stored in long term memory (Kintsch, 1988; McKoon & Ratcliff, 1992; Sanford, 1990). When you read about a UPS delivery, this could activate not only those words in your lexicon that are related to the word UPS, but also the world-knowledge about deliveries that is stored in your long-term memory (i.e. the man from UPS comes over to your house in his brown van, rings the bell and asks you to sign for your package). According to the 'resonance model' (Cook, Halleran, & O'Brien, 1998; Myers & O'Brien, 1998;

Myers, O'Brien, Albrecht, & Mason, 1994), individual concepts from the linguistic input send out a signal to long term memory. Concepts in memory then resonate as a function of their relatedness to the input, based on the overlap between the semantic and contextual features of the concepts involved. Eventually, those concepts that have the highest level of activation enter working memory. The resonance process is assumed to be fast-acting and autonomous (or “dumb” (Myers & O' Brien, 1998))

Message-level processing vs. automatic activation

Comprehending the actual meaning of a sentence is much more than extracting the scenario-relevant gist, or a simple addition of individual words. We are able to understand that “*The UPS man delivered the ...*” is fundamentally different from “*The man delivered UPS the ...*” even though the same words are present. It is indisputable that we are able to extract the complete and correct meaning of both these utterances. But is this complete message-level representation, recomputed with each new piece of information, also what drives the prediction of upcoming words? Or are these anticipatory processes completely based on the abovementioned automatic activation? The evidence for the role of message-based processing in early language comprehension is mixed.

Several experiments have shown that the message of an utterance can overrule the facilitating effects of the primes that are present in that sentence or story (Morris, 1994; Traxler, Foss, Seely, Kaup, & Morris, 2000; Van Petten, Coulson, Rubin, Plante, & Parks, 1999). For example, Morris (1994) showed his participants sentences that contained several primes. When the message of the entire sentence was in line with the individual primes, such as in “*The gardener talked as the barber trimmed the moustache.*”, the related word *moustache* was read faster. In sentences where the actual message of the sentence did not support the appearance of the primed word, i.e. “*The gardener talked to the barber and trimmed the moustache*”, this facilitatory effect disappeared.

However, other experiments have shown that scenario-mediated or lexical association can play a role in discourse comprehension even when the activated information is irrelevant to, or at odds with, the actual message of that discourse (Duffy, Henderson, & Morris, 1989; Garrod & Terras, 2000; O'Seaghdha, 1997). Garrod and Terras (2000), for example, showed that the word “*pen*” is initially just as effectively integrated when presented in a sentence following “*The teacher wrote a letter*” as it is after the sentence “*The teacher wrote the exercise on the blackboard*”. Only in regression path analysis and second pass reading times a significant difference was observed between the

appropriate and inappropriate contextual message. This indicates that participants did not at first notice the message-level incongruence of “pen”, which in turn suggests that in some cases (scenario-mediated) automatic activation can overrule the actual message of the discourse.

It thus seems that both message-level processing and lexical/ scenario-based association play a role when we try to understand and integrate incoming linguistic information. This suggests that both levels of analysis can play a part in predictive processing. In the following chapters, we explore whether prediction is purely based on automatic activation by lexical associations or related scenarios, or whether the actual message of the discourse also plays a role.

Outline of this dissertation

The central question addressed in this dissertation concerns the contextual basis of specific lexical predictions: are these predictions a simple by-product of relatively “dumb” automatic activation processes, or are they rooted in a more comprehensive representation of the predictive discourse?

In chapter 2, two ERP experiments are reported that approach this question from two different angles. The first experiment tests whether unexpected and incoherent words are interpreted as more unexpected and incoherent when a reader has made a (message-based) prediction. In the second experiment inflected adjectives that were not in line with the gender of the predicted word were used to test for the presence of predictions (as in Van Berkum et al., 2005). These prediction-probes were presented in stories that were either message-level constraining or that contained the same potential prime words, but were not message-level predictive. A replication of this second experiment with spoken materials is reported in chapter 3. The ERP experiment reported in Chapter 4 focuses on possible differences between people with larger and smaller working memory capacity in their ability to a) predict upcoming words and b) make message-based predictions.

Previous experiments have shown that specific lexical predictions involve semantic and lexical features of the upcoming word (Van Berkum et al., 2005; Wicha et al., 2004) and perhaps also acoustic properties (DeLong et al., 2005). Chapter 5 explores, in a self-paced reading paradigm, whether specific linguistic predictions can also include the exact visual manifestation of the predicted word.

The ERP study reported in chapter 6 investigates whether the classical difference in N400 for expected and unexpected words is completely based on (scenario-based) priming, or whether message-based processes are relevant as well.

Finally, in chapter 7 all findings are briefly summarized and discussed.

