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

Otten, M.

Publication date
2008

Chapter 7

General discussion and summary

Throughout the studies collected in this dissertation, we find that people are capable of predicting how a story will continue while they are hearing or reading that story. These predictions are highly specific, incorporating meaning (chapter 2), lexical features such as the grammatical gender of the upcoming word (chapter 2, 3 and 4), and possibly also the visual characteristics of the expected word in writing (chapter 5). Importantly, the results show that specific lexical predictions are not simply based on automatic word- or scenario-based activation. The experiments presented in chapter 2, 3 and 4 only showed evidence for the presence of specific lexical predictions when the actual meaning of the discourse supported such an inference. These predictions were not present for stories that contained the same words as the original predictive stories. This shows that people can anticipate upcoming words on the basis of what the discourse is really about.

Within the range of predictive processes examined here, limited cognitive or attentional resources (as indicated by the working memory capacity of the reader) do not seem to influence either the ability to make specific linguistic predictions or the contextual basis of these predictions (chapter 4). Readers with low and with high working memory capacity both make message-based predictions. Interestingly, readers with low working memory capacity show additional activation when confronted with information that is not in line with their prediction. This suggests that a lack of resources leads to increased processing of unexpected information.

Although the effect of prediction is consistently only present in predictive stories and not in prime control stories, the expected and unexpected nouns that follow the critical determiners and adjectives show a different pattern. The difference in N400 amplitude for expected and unexpected nouns that is present in the predictive context is also present in the prime control context. As we by now know, the message of the prime control stories did not support any specific predictions, and based on this there should thus not be a difference between the expectedness of the 'expected' and the 'unexpected' noun in a prime control story. The similarity of the effect in the predictive and the prime control story thus could be taken to suggest that this N400 effect does not reflect message-based expectancy, but scenario- or prime-based automatic activation. The nouns in chapter 2, 3 and 4 are, however, always preceded by gender-marked adjectives or determiners, which are likely to influence how a reader or listeners expects the story to continue. Indeed, in an

additional cloze test which included gender-inflected adjectives, reported in chapter 2, there was a difference in cloze value between expected and unexpected words in the prime control condition, which was not present in the original cloze test. This indicates that the difference in N400 amplitude observed for the expected and unexpected nouns in the prime control condition does not necessarily indicate that N400 amplitude reflects automatic activation. In chapter 6, we used a 'cleaner' design in which the gender-marked articles or adjectives were removed from the stories, to explore the influence of scenario-mediated priming and message-level meaning on the amplitude of the N400. In the prime control context, where the cloze values were very much comparable for the previously expected and unexpected nouns, unexpected nouns still evoked a more negative ERP in the N400 time-window (300-500). However, this effect had a different distribution compared to the N400 effect observed in the predictive context. These results show that the discourse dependent N400 effect cannot simply be completely reduced to (scenario-mediated) priming. The N400 effect reflects the rapid use of precise message-level constraints in comprehension. However, these results also show that, at least in the absence of strong message-level constraints, scenario-mediated priming can rapidly affect comprehension.

The role of discourse-message and lexical association in predictive processes

The results reported in chapters 2, 3 and 4 show that the message of the discourse lies at the basis of the on-line lexical predictions that readers and listeners make when confronted with a constraining story. This does not mean, however, that lexical association can play no role in the predictive process. There is quite a lot of evidence that the close proximity of a strongly related prime influences the processing of a related word, even when the message of the discourse is actually unrelated (Camblin et al., 2007; Ditman, Holcomb, & Kuperberg, 2007). The stories used in chapter 2, 3 and 4, however, were not designed to specifically include strong primes, and if they did, the distance between the strong prime and the critical region (inflected adjective and (un)expected noun) was relatively large. It is thus impossible to evaluate the exact role of lexical association from the present results. What the present results unequivocally show, though, is that lexical association is definitely not the only contextual factor that plays a role in the development of linguistic predictions. What the text is actually about plays an important role as well.

Implications for current models of language comprehension

Several models of lexical activation in language comprehension state that simply reading words activates other, related words (Collins & Loftus, 1975; Cook et al., 1998; Myers & O' Brien, 1998; Myers et al., 1994). Such models could thus explain how a reader or a listener can pre-activate words that are likely to follow in that sentence or story. DeLong and colleagues (2005) indeed suggested that words or combinations of words that are present in the preceding sentence activate related words and world-knowledge, which leads to the pre-activation of specific words. However, as discussed before, the results of the experiments in chapter 2, 3 and 4 show that this cannot be the whole story.

So, what are the other options? We suggest that what underlies the specific prediction observed here is the result of convergent predictions being made at several levels of unfolding structure. It is well known that language comprehenders compute the syntactic and conceptual analysis of the incoming language *incrementally* and in parallel (see Jackendoff, 2002; Jackendoff, 2007 for an overall framework and; Vosse & Kempen, 2000 for an explicit computational model of the syntactic side of things). As a consequence, at any point in an unfolding sentence, readers and listeners have at their disposal a partial syntactic and conceptual analysis of the preceding sentence fragment. Each of these partial representations can by itself suggest what might come next. Although these predictions arise at different levels of representation, it is not difficult to see how they might come together and converge onto a specific word. As laid out by Jackendoff (2002; 2007; see also Kempen & Huijbers, 1983; Levelt, 1989), an individual lexical item consists of bits of orthographic, phonological, syntactic, and conceptual information, bundled together into a single multi-leveled structure. If people actually read or hear a word the associated fragments of syntactic and conceptual structure are activated via their orthography or phonology, and merged ('unified') with the syntactic and conceptual analyses constructed for the language input so far. However, within the same framework, the preceding syntactic and conceptual context can, if sufficiently constraining, also each *pre-activate* the relevant bits of structure, resulting in the prediction of the related lexical item. In this case, it is the convergent pull of syntactic and message-based conceptual constraints that activates a particular word, and not the orthographic or phonological input.

Electrophysiological consequences of disconfirming a lexical prediction

In three experiments that explored on-line lexical predictions in language processing (chapter 2, 3 and 4) we consistently found that the ERPs to prediction-consistent and prediction-inconsistent information differ in the predictive context. However, the nature of this ERP effect of prediction is not as consistent over experiments. When gender-inflected adjectives were used as prediction-probes (chapter 2, experiment 1B), the (visually presented) adjectives with an inconsistent inflection evoked a late negative deflection, from 900 to 1200 ms after stimulus onset with a right-frontal distribution, compared to consistent adjectives. In a replication of this experiment with spoken stimuli (chapter 3), the inconsistent adjectives also evoked a negative deflection with a right-frontal distribution, but much earlier, from 200-600 ms after the acoustic onset of the adjective. Finally, when gender-marked articles were used as the critical probes in written stories (chapter 4), prediction-inconsistent articles evoked a negative deflection between 200 and 600 ms over right frontal electrodes, which was followed, for readers with a low working memory capacity only, by a later more centrally distributed negativity between 900 and 1500 ms.

It is important to note that, irrespective of the variation in the nature of the prediction effects observed in the predictive context, we have found a consistent difference between predictive and prime control contexts: significant effects of prediction-mismatch are observed in predictive stories, whereas the prime control stories show no differential activation. This consistent pattern, that is related to our contextual predictiveness manipulation, suggests that the effects observed in the predictive condition are not simply 'false alarms'. If this were the case, then the same noisy processes that underlie the effects in the predictive context should have been present in the prime control context, leading to ERP 'effects' in both types of discourse. Therefore, the ERP effects observed in chapters 2, 3 and 4 can thus not simply be disposed of as noise.

Furthermore, there are important similarities between the prediction-related ERP effects described in this dissertation and previous experiments. The early negative deflection observed in chapters 3 and 4 show important similarities with other ERP effects found with the prediction-probe paradigm. Three previous studies have observed an early negative inflection with a widespread, central distribution, between 300 and 500 ms after the onset of the prediction-inconsistent information (in this case determiners with prediction-inconsistent features in English (DeLong et al., 2005) and Spanish (Wicha, Bates et al., 2003; Wicha, Moreno et al., 2003)). The time course and

polarity of these effects seems comparable to the findings reported in chapter 3 and 4, although the scalp distribution differs slightly (visual inspection of the scalp distribution of the prediction effect in Wicha, Bates et al., 2003 however does seem to suggest a somewhat right-frontal distribution).

Additionally, the late negativity observed in chapter 2 is somewhat similar to the additional late negativity that is displayed by low WMC readers in chapter 4. This resemblance could be taken to suggest that these two effects are actually the same. Such a conclusion, however, immediately raises the question why the late negativity observed in chapter 2 is not preceded by the early negative component that is present for both low and high span readers in chapter 4. One important difference between the two experiments was the type of prediction-probe used (adjectives with gender-dependent inflections in chapter 2 and gender-marked determiners in chapter 4). Indeed, one could argue that an inflected adjective provides a somewhat more subtle clue about the gender of the upcoming noun than a gender marked determiner, since determiners are very common and very short. The unexpected gender of the determiner are thus likely to be perceived at a glance (i.e. compare "de" vs. "het") whereas the subtle difference in adjective inflection, which usually relies on the presence or absence of a single "-e" is less conspicuous (i.e. compare "klein" vs "kleine"). The fact that the inconsistency of the adjective inflection is less striking than the inconsistency of the determiner could have led to a delay in noticing the inconsistency or to a different, less profound response for the participants in chapter 2. However, if this would be the case, then we should have observed the same pattern of results in the experiment reported in chapter 3, where, just as in chapter 2, inflected adjectives were used as the critical probes. As stated above, the spoken adjectives with a prediction-inconsistent inflection from chapter 3 showed an early negative shift that is comparable to the early negative shift evoked by the prediction-inconsistent determiners in chapter 4. It could thus be that perhaps other experimental features, such as individual differences of participants, or specific features of the experimental design such as the type of fillers, influence the way in which a reader or a listener makes predictions, or processes prediction-inconsistent materials.

Taken together, we must conclude that we can not draw any final conclusion about the nature of the prediction effect. One of the core goals of future research should thus be to determine what factors underlie this variability. One important factor that deserves to be explored is the influence of strategic processes in prediction: can we as readers unconsciously 'decide' to rely more strongly on predictive processes, for example in a noisy

environment (see Pickering & Garrod, 2007), or, on the contrary, diminish the influence of prediction, for example when everything that we read follows unexpected and seemingly incongruent paths? It could be that such broad contextual factors influence the presence or nature of the predictions that are made, and with that, perhaps can explain the variability in observed ERP effects.