NeuroImage 147 (2017) 447–460

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/neuroimage

Language in context: Characterizing the comprehension of referential expressions with MEG

Christian Brodbeck[a,b,*], Liina Pylkkänen[a,b,c]

[a] Department of Psychology, New York University, New York, NY 10003, USA
[b] NYU Abu Dhabi Institute, Abu Dhabi, United Arab Emirates
[c] Department of Linguistics, New York University, New York, NY 10003, USA

ARTICLE INFO

Keywords:
MEG
EEG
Sentence comprehension
Reference resolution
Visual short-term memory

ABSTRACT

A critical component of comprehending language in context is identifying the entities that individual linguistic expressions refer to. While previous research has shown that language comprehenders resolve reference quickly and incrementally, little is currently known about the neural basis of successful reference resolution. Using source localized MEG, we provide evidence across 3 experiments and 2 languages that successful reference resolution in simple visual displays is associated with increased activation in the medial parietal lobe. In each trial, participants saw a simple visual display containing three objects which constituted the referential domain. Target referential expressions were embedded in questions about the displays. By varying the displays, we manipulated referential status while keeping the linguistic expressions constant. Follow-up experiments addressed potential interactions of reference resolution with linguistic predictiveness and pragmatic plausibility. Notably, we replicated the effect in Arabic, a language that differs in a structurally informative way from English while keeping referential aspects parallel to our two English studies. Distributed minimum norm estimates of MEG data consistently indicated that reference resolution is associated with increased activity in the medial parietal lobe. With one exception, the timing of the onset of the medial parietal response fell into a mid-latency time-window at 350–500 ms after the onset of the resolving word. Through concurrent EEG recordings on a subset of subjects we also describe the EEG topography of the effect of reference resolution, which makes the result available for comparison with a larger existing literature. Our results extend previous reports that medial parietal lobe is involved in referential language processing, indicating that it is relevant for reference resolution to individual referents, and suggests avenues for future research.

1. Introduction

When language is used for communication, new information is not presented in a vacuum but is connected to information that is already known. Consequently, a fundamental device of language are expressions that invoke entities that are already known to the addressee (Lambrecht, 1994). For example, use of the definite noun phrase *the revolution* signals that the addressee should be able to identify which revolution in particular the speaker is talking about. The interpretation of referring expressions is thus a very fundamental process in language comprehension, allowing new meaning to be constructed on the base of known background information (see also van Deemter (2016)). And yet, the neural correlates of successful reference resolution remain largely uncharacterized. This is apparent from the absence of referential processing from recent models of the neural basis of sentence level language comprehension (e.g. Friederici, 2011; Hagoort and Indefrey, 2014). Here we report on work that introduces a paradigm to investigate the neural basis of reference resolution. We provide evidence across three experiments and two languages that reference resolution in simple visual referential domains involves medial parietal cortex.

Previous research has shown that referential language processing is fast and incremental (Tanenhaus et al., 1995), takes into account a wide array of extra-linguistic sources of information (Chambers et al., 2002; Altmann and Kamide, 1999; Kamide et al., 2003; Altmann and Kamide, 2007) and can even affect syntactic parsing decisions (Tanenhaus et al., 1995; Spivey et al., 2002). Most of this evidence has come from studies using the so-called visual world paradigm, in which participants' eye movements are recorded while they follow instructions to perform various tasks with objects laid out in front of them (for reviews see Tanenhaus and BrownSchmidt (2008) and Huettig et al. (2011)). Studies that focused directly on reference

* Correspondence to: Institute for Systems Research, University of Maryland, College Park, MD 20742, USA.
E-mail address: christianbrodbeck@nyu.edu (C. Brodbeck).

http://dx.doi.org/10.1016/j.neuroimage.2016.12.006
Received 9 September 2016; Accepted 4 December 2016
Available online 16 December 2016
1053-8119/

resolution found that people typically move their eyes to the referent of an expression as soon as they have sufficient information to identify it (Eberhard et al., 1995; Sedivy et al., 1999). For example, when participants were asked to *Touch the starred yellow square* in a context with only one starred item, they moved their eyes to the relevant item on average 250 ms after the end of the word *starred*, significantly earlier than in contexts where more than one object had stars. However, eye movements constitute an indirect measure and do not necessarily reflect reference resolution to a unique item. When the first syllable of an utterance is compatible with two different continuations (e.g., *the beetle* vs *the beaker*), addressees distribute their fixations over both items (Allopenna et al., 1998). This suggests that eye movements reflect attentional processes, guided by even partial word information, and before reference can be positively resolved.

Neural measures can provide complementary information to improve our understanding of the computational stages of reference resolution. In particular, electrophysiological measures can provide temporally precise signals, making it possible to measure the response to individual words in coherent language stimuli. A repeated finding from event related potential (ERP) studies is an N400 reduction to expressions resolving reference to an entity introduced in the previous sentence in a non-marked way (Burkhardt, 2006; Ledoux et al., 2007). Referential context can influence this response, suggesting that referential processing is occurring in the relevant time window between 250 and 500 ms (Ledoux et al., 2007). A later positive component has been associated with retrieval and updating when a prior referent is mentioned again (Van Petten et al., 1991). However, this observation is complicated by the fact that a similar response has also been associated with the introduction of a new discourse referent, when compared with reference to an existing one (Burkhardt, 2006, 2007). Yet another study that manipulated referential status through the article (*Kathy sat nervously in the cab on her way to the airport. A/The cab…*) found no late component in either direction (Anderson and Holcomb, 2005). Together these results suggest that the late ERP component is sensitive to multiple factors and not yet interpretable as a direct measure of referential status per se. Another brain measure, which has achieved a higher degree of functional specificity, is a component related to referential ambiguity. A group of EEG studies have found that referentially ambiguous expressions evoke a sustained frontal negative-going event-related potential when compared to unambiguous controls (reviewed by Nieuwland and Van Berkum (2008)). This includes determiner-noun phrases like *the girl* after a story context that introduced two girls vs one girl (Van Berkum et al., 1999, 2003; Nieuwland et al., 2007a; Boudewyn et al., 2015) as well as pronouns matching two vs one previously introduced characters (e.g., *he* in *Ronald told Frank that he…* can refer to *Ronald* as well as *Frank*, whereas *he* in *Ronald told Emily that he…* can only refer to *Ronald*) (Nieuwland and Van Berkum, 2006).

An fMRI study localized the response to referentially ambiguous pronouns to multiple prefrontal and parietal cortices, while unambiguous pronouns were associated with higher activation in the inferior frontal gyrus bilaterally (Nieuwland et al., 2007b). While referential ambiguity involves a direct manipulation of reference resolution, other cognitive processes might be involved. The failure to find a referent for a pronoun could be associated with an activity increase in the basic referential search processes (searching harder) or a decrease (giving up), as well as activity in other regions recruited to deal with the ambiguity (Nieuwland and Van Berkum, 2008). A subsequent study that focused on the decision making process associated with assigning pronoun referents based on different cues such as gender and verb bias confirmed that this component of the task involves a broad network of frontal, parietal and temporal regions (McMillan et al., 2012). While these results thus indicate a broad set of regions that could be relevant for reference resolution, they might also include regions involved in higher order decision making processes.

Our investigation had the goal of testing for activity in these regions that is immediately associated with successful reference resolution when higher order decision making requirements are minimal. A challenge for studying successful reference resolution is finding a suitable control condition that does not involve creation of new discourse referents while at the same time not introducing another contrast like referential failure or referential ambiguity. To this end we adapted the visual world paradigm that has been used for studying incremental reference resolution (Eberhard et al., 1995, see above) for MEG/EEG in three experiments, each with the same basic trial structure (see Fig. 3 for Experiment 1). In each trial, participants first saw a simple visual world display and then read a question about that display, presented word by word. Participants were required to answer each question with a yes/no button press, encouraging them to process the questions naturally with the goal of comprehending them and without drawing undue attention to the referential aspect of the task. We analyzed the neural response to simple referential expressions like *the blue heart*, comparing the same expressions in different referential contexts. In the example in Fig. 3 the adjective *blue* could resolve reference in a context with one blue item, but not in a context with two blue items. This contrast isolates reference resolution with minimal overt ambiguity. Based on EEG data collected concurrently with the MEG data reported here for Experiment 2, we previously reported that evoked potentials to reference resolving words reflect the location of the referent on the display around 333 ms after adjective onset, confirming that the manipulation is effective in engaging referential processing at the adjective (Brodbeck et al., 2015). MEG allowed us to measure brain activity associated with the processing of individual words with high temporal precision. This allowed us to track neural activity occurring during the processing of specific words in referential expressions, in contrast to fMRI which does not allow attributing activity temporally to individual words in a sentence.

While the fMRI study on referential ambiguity (Nieuwland et al., 2007b, see above) provides us with a broad set of regions of interest, there are also other relevant prior results. First, another fMRI study found that discourses involving two conjoined referents compared to discourses with one or two singular referents engaged medial and superior/lateral parietal regions (Boiteau et al., 2014). This result suggests that parietal cortex is involved in the representation of discourse referents because it is sensitive to the manner in which multiple discourse entities are introduced. Second, reference resolution is a fundamental component of processing coherent language, because coherence critically depends on repeated reference to the same entities. A meta-analysis found that the processing of coherent language is associated with increased activity in medial parietal, medial frontal and bilateral temporal areas (Ferstl et al., 2008). Together, these results highlight the parietal lobe as possibly relevant for referential language processing, since it is the only region that was reliably affected by all three contrasts. Almor et al. (2007) suggested that parietal involvement in referential processing could reflect recruitment of circuits originally devoted to perceptual organization, tracking multiple objects in space, for keeping track of multiple discourse referents. This connection is particularly relevant for our design, which used visuo-spatial referential domains, and predicts that reference resolution should be associated with parietal activity.

2. Experiment 1

Experiment 1 demonstrated the viability of the paradigm and established the main result. However, in order to keep the paradigm and the task simple, compromises were made which lead to some potential confounds. Those were addressed in Experiments 2 and 3. Each visual world display was composed of three colored shapes, providing the context for simple adjective-noun referential expressions such as *the blue heart*. The primary target stimulus was the color adjective, the first word that differed between trials in whether it