1 Cognitive neuroscience of natural language use: introduction

Roel M. Willems

The cognitive neuroscience of language investigates the neural infrastructure underlying the comprehension and production of language. When we think of language use in our daily lives, the first things that come to mind include colloquial conversations, the small chat on your way to work, reading and writing emails, sending text messages, reading a book. To the surprise of outsiders, the topic of study in the cognitive neuroscience of language is rather far removed from these language-in-use examples. There are well-founded (historical) reasons for this (see below), but the current excitement in the study of language comes from new developments that make studying the neural underpinnings of naturally occurring language much more feasible.

The chapters in this book provide a state of the art overview of current approaches to making the cognitive neuroscience of language more 'natural' in the sense of closer to language as it occurs in real life. Before giving an overview of the book's content, I will briefly introduce two strands of language research, that approaches in this book draw upon.

Two traditions of language research

Two research traditions in the study of language are important for present purposes.

(1) The controlled, simplified stimuli tradition

Research in this tradition is done in the laboratory, under controlled circumstances, presenting participants with carefully selected stimuli.

For instance, participants see single words presented on the screen and they have to decide as fast as possible whether the word is a noun or a verb. Or, participants listen to phonemes ('d' and 'b'), and the phonemes are manipulated to make the one sound more like the other. Task for the participant is to decide whether she hears a 'b' or a 'd'.

In this style of research, the researcher has tight control over experimental factors of interest, and potentially confounding factors ('nuisance

variables') are taken care of. The drawback is that the setting is highly unnatural, and that the language that participants listen to is decontextualized.

This tradition dates back at least to the early days of reaction time experiments, notably the experiments in the laboratories of Wundt and Donders (e.g. Donders, 1869; see Levelt, 2012 for a historical overview). Sophistication in experimental design and analysis, and understanding of the cognitive processes underlying language, has increased enormously since then. However, the basic mode of working in this research tradition has remained the same: Careful isolation of a subprocess of language is studied under highly controlled conditions with decontextualized and simplified language stimuli.

(2) The ecological laboratory tradition

In this line of research, stimuli and situations are closer to language in real life. Research is usually done in a laboratory, but participants are much less restrained in their behaviour as in the other tradition.

For instance, pairs of participants come to the lab and they solve a problem (e.g. matching two shapes, as in a puzzle) together. One individual knows something about the problem that is important for the other one to be able to solve it. The focus of the research could be on how people adapt their gesturing or speaking to the amount of information that is shared between participants. The dependent variable in this line of research gets quantified (e.g. number of gesture elements), and often requires researcher-dependent coding ('What counts as a gesture?'). The flavour is much more towards naturalistic language use, at the expense of a loss of 'experimental control', meaning that there are more degrees of freedom for behaviour of the participant than in the other research tradition.

The first, controlled, tradition has been most popular in experimental psychology and cognitive neuroscience, and, although tremendously productive, its limits are starting to become more seriously considered. Simply put, there is a growing understanding that studying what happens when people read a single word or sentence, does not necessarily scale up to what we do when we understand natural language, say read an e-mail or talk with our neighbour.

Combining the two research traditions in studying the brain basis of language

The usual status quo is that the two research traditions are thought of as mutually exclusive. Increasing ecological validity comes at the cost of

experimental control, and tightening experimental control decreases ecological validity. Hence, combining them is considered impossible. This is even more pertinent for studying the brain basis of language. Experimental constraints reduce the possibility of studying more natural language even further (e.g. participants cannot move in an MR scanner). The main starting point for this book is that there is no reason for such pessimism: recent studies have investigated natural language with cognitive neuroscience methods, and they have obtained much more encouraging results than one may expect from the standard 'impossible' reply.

This book brings together scholars who borrow from both research traditions, in order to come to a cognitive neuroscience of natural language use. The combination of research traditions is more obvious in some contributions than in others, but generally speaking there are two starting points:

First, the language investigated is 'natural language': language as it could be encountered by participants in the real world. This means contextualized, and in larger chunks than usually done (beyond word and single sentence level).

Second, experiments maintain tight experimental control, and use the tools of cognitive neuroscience in a creative manner to enable the study of naturalistic language.

However diverse the topics of the separate chapters are, together they make the case that using more naturalistic settings and stimuli is (a) possible while keeping experimental control, and (b) provides new insights into the cognitive neuroscience of language.

Why study the cognitive neuroscience of natural language use?

A commonly heard reply to pleas for cognitive neuroscience to 'go more natural' is: 'Why?' Why not stick to the syllables and single-word experiments, that we know how to do so well? What is the added value of using more natural stimuli, why make experiments unnecessarily complicated? A simplistic answer is that one should study one's topic of interest ('language') as close to real life as possible. Recent developments in data acquisition and analysis show that it is possible to stay closer to natural language, and still find interesting and interpretable results. This will not satisfy a skeptic who will reply that it is unnecessary to stay close to natural language, because one can extrapolate findings from well-controlled and simplified experiments to more natural language situations.

And here is where the better answer comes in: extrapolation from findings obtained in restrained lab settings to real life may be much less

obvious than we think. In the visual (neuro)sciences, it is becoming more and more clear that some basic rules about visual perception, apply to perception in more natural settings only to a limited degree. An example case is visual search, where recent experiments show that the search rules discovered in controlled lab experiments, should be complemented by additional rules to explain visual search in more complex scenes (Wolfe, Võ, Evans & Greene, 2011). On a similar note, others have argued that the use of simplified and decontextualized stimuli has been an important hindrance in understanding the working of visual brain areas (Olshausen & Field, 2005), and studying the neural basis of visual perception with more natural and complex stimuli is currently an important topic in the visual neurosciences (e.g. Çukur, Nishimoto, Huth & Gallant, 2013; Peelen & Kastner, 2014).

This is not to say that findings obtained with controlled lab experiments, using simplified stimuli, are uninteresting. Quite the contrary, some of these findings will also apply to more natural settings, but to a more limited degree than thought by proponents of the controlled, simplified stimuli approach. As I will stress again in the concluding paragraph of this introduction, it is still too early days to appreciate whether and how studying more natural language use will change the landscape of cognitive neuroscience. But initial results are interesting and push towards a revision of some of our knowledge about language and the brain.

A final answer to the question 'Why study cognitive neuroscience of natural language?' is more obvious. Stepping up from single words and sentences allows to investigate the neural basis of for example narrative comprehension, face-to-face communication, event segmentation, etc. These are part of normal language use, but are often ignored in cognitive neuroscience because of the perceived difficulties in studying them. One message of the chapters in this book is that those difficulties can be overcome.

An issue that I have left aside so far is what counts as 'natural language' and what does not. This is done on purpose. What counts as natural language to the student of phonology will likely look like a much impoverished version of language to those with an anthropology background. What I think will not help the field forward is quarrelling about how natural the language under study has to be, or, even less productive, me providing a definition of 'natural language use'. It should be obvious from what I have written so far that natural language use refers to richer language stimuli than often used in lab experiments (less simplified), understood in richer contexts than is often the case. There will often be a tension between the degree of naturalness and the possibility of doing well-controlled and interpretable experiments. No hard rules can be given

on how to strike the balance between the two, and researchers will have to decide what they find the optimal balance for themselves.

Overview of chapters

The chapters in this volume can all be read as stand alone, and are grouped thematically only loosely. Each chapter starts with a brief abstract to orient the reader, and emphasis has been put on readability to a broad readership. All chapters are a combination of overview of recent experimental work, and opinion by the authors. Speculation was not discouraged, making the chapters well suited as starting points for discussion.

In Chapter 2, **Andric** and **Small** give an overview of methods and techniques used in fMRI studies that use more natural language stimuli than is usually done. Such experiments ask for unorthodox analysis methods, and the state of the art is reviewed in this chapter. In Chapter 3, **Ash** and **Grossman** provide a coherent overview of their substantial neuropsychological (patient) work on natural speech production. Using speech elicitation methods such as narrative retellings, their work illustrates the powerful combination of exact characterization of the deficits in natural speech production and measures of neural impairments in a variety of patient populations. In Chapter 4, **Kurby** and **Zacks** describe behavioural and neural studies that give insight into how readers construct and update situation models during narrative and discourse comprehension. The study of situation models is a nice example of a question that can by definition only be asked using stimuli which go beyond the word or sentence level, and the chapter illustrates the explanatory power of cognitive neuroscience in investigating this issue. In Chapter 5, **Knoeferle** shows how studying language comprehension can be enriched by taking non-linguistic (visual) context into account. Listeners quickly use visual information when available, extending the scope of traditional research beyond the study of isolated language stimuli. In Chapter 6, **Skipper** introduces his model of the Natural Organization of Language and the Brain, a model specifically designed to understand the neural instantiation of natural language comprehension. In Chapter 7, **Jacobs** sketches a neurocognitive model of an instance of natural language in optima forma, namely the reading of literature. Drawing upon insights from literary science, as well as on findings from experimental studies, he outlines what a model of literary reading looks like, providing an important theoretical basis for future work. In Chapter 8, **Kristensen** and **Wallentin** add novel insights into the debated role of Broca's area in language production and comprehension. They focus specifically on the influence of context, an often neglected but important factor in understanding how

this part of cortex adds to language, as is evidenced by their overview. In Chapter 9, **Kuhlen**, **Allefeld**, **Anders** and **Haynes** look into the growing literature on the cognitive neuroscience of dialogue. The chapter shows what has been established from studying brain signals from two people while they are in dialogue, and what remains to be done. In Chapter 10, **Stolk**, **Blokpoel**, **van Rooij** and **Toni** illustrate how studying brain activation while participants play a novel communication game, provides new insights into the neural basis of our communicative abilities. They moreover present a model for human referential communication, which – while based on studies that did not employ linguistic codes – has implications for our understanding of communication by linguistic means as well. Finally, in Chapter 11, **Hasson** and **Egidi** outline how processes at higher levels of language comprehension (e.g. discourse comprehension) can be understood in terms of more basic neural processes. They warn against misinterpretation of studies into natural language comprehension, and provide guidelines that should inform future studies of natural language comprehension.

Away from the sterile

In preparation of this chapter, I came across the following quotation, that I much sympathise with:

> that gradually, more and more scientific psychologists were willing to abandon the somewhat sterile environment of nonsense syllables and other laboratory-produced elements, for the study of objects less removed from real life. (De Groot, 1946, p. 56; transl. RW)[1]

This quote from renowned chess psychologist Adriaan de Groot's doctoral dissertation (*Thought and Choice in Chess*) voices a struggle that every student of language has to face. How close do you stay to real life when translating your topic of interest into a testable experiment? De Groot's solution was 'guided introspection', a methodology made popular by the Würzburg school in psychology (see Levelt, 2012), and one that most current psychologists would frown upon.

What I like about the original wording is the use of 'sterile': De Groot speaks about the 'sterile environment of nonsense syllables and other

[1] The original Dutch: 'dat gaandeweg steeds meer wetenschappelijke psychologen bereid bleken de ietwat steriel geworden sfeer van de zinlooze lettergrepen en andere alleen in het laboratorium te isoleeren 'elementen' te verlaten voor excursies naar minder levensvreemde gebieden' (De Groot, 1946, p. 56).
A digital copy of the English version of this lovely book (*Thought and Choice in Chess*) is freely available from the Amsterdam University Press website.

laboratory-produced elements'. Just like a surgeon who sterilizes his equipment before surgery, we have sterilized language by stripping off unwanted elements, until we end up with purely controlled language stimuli such as nonsense syllables.

The difference is that the surgeon has good reason to sterilize his equipment. After all it is not his purpose to understand what bacteria exist and how they grow, he needs to get rid of them, full stop. Students of language, on the contrary, lose something when they sterilize their material. They lose the capacity to understand language as it is in reality, dirty and complicated. It is my hope that the chapters in this book will help in re-establishing the cognitive neuroscience of language, away from the sterile.

References

Çukur, T., Nishimoto, S., Huth, A. G. & Gallant, J. L. (2013). Attention during natural vision warps semantic representation across the human brain. *Nature Neuroscience*, *16*(6), 763–770. doi:10.1038/nn.3381

De Groot, A. D. (1946). *Het denken van den schaker*. Amsterdam: Noord-Hollandse Uitgevers Maatschappij.

Donders, F. C. (1869 / 1969). On the speed of mental processes. *Acta Psychologica*, *30*, 412–431.

Levelt, W. J. M. (2012). *A History of Psycholinguistics: The Pre-Chomskyan Era*. Oxford University Press.

Olshausen, B. A. & Field, D. J. (2005). How close are we to understanding v1? *Neural Computation*, *17*(8), 1665–1699. doi:10.1162/0899766054026639

Peelen, M. V. & Kastner, S. (2014). Attention in the real world: toward understanding its neural basis. *Trends in Cognitive Sciences*. doi:10.1016/j.tics.2014.02.004

Wolfe, J. M., Võ, M. L.-H., Evans, K. K. & Greene, M. R. (2011). Visual search in scenes involves selective and nonselective pathways. *Trends in Cognitive Sciences*, *15*(2), 77–84. doi:10.1016/j.tics.2010.12.001

2 fMRI methods for studying the neurobiology of language under naturalistic conditions

Michael Andric & Steven L. Small

Abstract People ordinarily use language in complex, continuously occurring contexts. These contexts include rich and varied sources of information that can combine and exert shared influence. A common example is face-to-face conversation. When two people talk in person, there is not only spoken auditory information, but also visual information from facial and manual movements, as well as from the surrounding environment in which the conversation takes place. There may also be endogenous signals that a person experiences in context, such as memories relating prior conversations with a given speaker. In short, it is typical that a person mediates multiple multifaceted information sources when using language. By contrast, fMRI studies of the neurobiology of language often use conditions that present only features of language in isolation. In large part, this is because researchers need rigorous, reliable experimental protocols that minimize potential sources of variance ("noise") not directly relating a feature of interest. But such traditional protocols also often minimize, if not eliminate, the way people actually know and use language in their naturalistic, "everyday" experience. Thus, a fundamental challenge for researchers studying the neurobiology of language is to understand brain function as it might occur in typical experience. In this chapter, we highlight some available approaches that can address this challenge using fMRI. With specific examples, we discuss the importance of context and ways to incorporate it in understanding brain function when people process language under more naturalistic conditions.

Introduction

People ordinarily use language in rich and varied contexts, continuously mediating information from diverse sources. This includes information a person perceives in the external world, as well as from his or her own internal states and processes. Moreover, these sources typically combine in exerting shared influence. For example, when two people talk face-to-face, they might each focus on the other's vocalizations, facial and manual movements, and the meanings these convey. Understanding brain processes as they relate these features is an important area in neuroimaging

research, and numerous studies have aimed experimentally to characterize the neural bases of such context effects.

However, naturalistic experience – or "everyday life" – typically presents many information sources in continuous, complex combinations. In other words, it is rare, if ever the case, that people encounter a feature of language in isolation, as they do in experimental settings. Whether or not the entire amount of contextual information is directly relevant to a particular conversation's meanings or effects, it is present nonetheless. For example, imagine a face-to-face conversation on a busy street corner, where there is ongoing audiovisual information from passersby, automobile traffic, and other diverse sources. Also present are a person's internal ongoing signals – anything from their remembering that the person with whom they are talking is known for telling lies, to the indigestion from something they just ate – that can also influence the way a person interprets what their conversational partner says. Such complex rich natural settings are not ideal for isolating a controlled variable. Yet it is these kinds of settings in which people actually know and use language on a regular basis.

Certainly, rigorous experimental protocols require controlled variables. To determine an effect for a given feature of interest, researchers typically seek to minimize influence from potentially surrounding sources of non-interest ("noise" variance). But this necessity contrasts with naturally occurring contexts, which inherently comprise diverse, dynamic, and interacting information. This contrast among settings poses a great challenge for researchers: as a researcher, how is it possible to study language and these myriad sources of information in the way people actually use and experience them, yet at the same time employ experimental methods that are reliable and valid, and that ultimately yield informative results?

Increasingly, fMRI researchers acknowledge this challenge. Recently developed fMRI protocols better address natural experience than was possible previously. Some research now incorporates stimuli and methods that give and use contextual information, i.e., information that surrounds a feature of primary interest, rather than going to great lengths to minimize it. Importantly, results from these efforts are proving informative. This is not just because these results uniquely characterize ways that the brain functions. It is also because such results characterize ways that the brain functions under experimental conditions that people recognize – that better generalize to people's actual use and experiences – than has been typical in earlier fMRI studies.

In this chapter, we outline and discuss several approaches and methods for language fMRI research that move towards maintaining rigorous and reliable experimental methods in protocols that, at least partly, resemble

people's typical naturalistic language experiences. First, for perspective, we briefly highlight a number of mainstream findings in language fMRI, along with the traditional experimental designs used to acquire them. Next, we reiterate the importance of context in studying brain and language, as well as their accompanying features. We then highlight more recent studies. Our focus is on investigations that use methodological approaches that allow analysis of data acquired under stimulus conditions that resemble naturalistic, everyday situations. To convey the accessibility of these approaches, we discuss – over multiple sections of this chapter – specific examples of these methodologies. Finally, we conclude by summarizing and looking ahead to the utility of fMRI in future studies that seek to understand language function in naturalistic contexts.

Early findings and the methods they use

As a whole, fMRI characterizations of the functional anatomy of language encompass a large scope. In large part, this is because language affords many levels of processing, with many diverse questions that fMRI can help answer. For example, questions are asked about brain responses at multiple levels, from discourse, to words and sentences, to syllables and phonemes. Within just these linguistic levels, further specificity of analysis can be performed. For example, a topic of recent interest for fMRI research is whether processing action words involves responses that might be present when people perform those actions (see Pulvermuller, 2005). There is also much research focusing on brain function when people perceive actions that accompany and convey meaning with language, as with "co-speech gestures" (for review, see Andric & Small, 2012; Willems & Hagoort, 2007). Even such "higher-level" aspects as brain encodings for semantic concepts represent an active research area (see Binder *et al.*, 2009). Not only for receptive language, questions that focus on brain function when people produce language also branch into many categories and levels of analysis. Put simply, the dimensions, questions, and analyses of brain processes that can classify as "language" are broad, numerous, and multifaceted.

Interestingly, as diverse as the questions and angles from which people examine language are, to this point, the fMRI protocols and analyses researchers use are generally uniform, comprising relatively few techniques. In the following, we highlight some of the earlier questions and findings from language fMRI studies, and briefly discuss some of the common approaches researchers use to derive such findings.