Exploring social biases in language processing

ISBN: 978-94-92910-09-7

Cover artwork by Silvia Iacozza

Printed and bound by Ipskamp Drukkers b.v.

Exploring social biases in language processing

Proefschrift

ter verkrijging van de graad van doctor

aan de Radboud Universiteit Nijmegen

op gezag van de rector magnificus prof. dr. J.H.J.M. van Krieken,

volgens besluit van het college van decanen

in het openbaar te verdedigen op maandag 23 maart 2020

om 12.30 uur precies

door

Sara Iacozza

geboren op 8 maart 1988

te Rome (Italië)

Promotor:

Prof. dr. Antje S. Meyer

Copromotoren:

Dr. Monique E.P. Flecken (MPI)

Dr. Shiri L. Lev-Ari (Royal Holloway University of London, Verenigd Koninkrijk)

Dr. Alastair C. Smith (MPI)

Manuscriptcommissie:

Prof. dr. Asli Ozyurek

Prof. dr. Jos J.A. van Berkum (Universiteit Utrecht)

Dr. Atsuko Takashima

This research was supported by the Max Planck Society for the Advancement of Science, Munich, Germany.

To those who shared this journey with me

Table of Contents

1 | General introduction

> No man is an island entire of itself; every man is a piece of the continent, a part of the main [...]
>
> John Donne, 1624

Reading John Donne's *No man is an island* feels like reading a contemporary activist's message against isolationism, talking about the importance of showing solidarity with each other, regardless of differences in nationality and ethnicity. A message that strongly resonates in our contemporary communities. The European Union (EU) has in fact invested substantial economic and political resources in creating a common European identity that would be shared amongst the citizens of its 28 member states.

The creation of a shared identity is based upon the idea that citizens should feel like they belong to the same group, or to the same union of groups, at least. From a socio-psychological point of view, the emergence of a sense of belonging is a complex phenomenon, rooted in the tendency of human beings to discriminate others based of their affiliations and categorize them as either members of their own group (i.e., in-group members) or as members of other groups (i.e., out-group members) (see Turner & Tajfel, 1986, for Social Identity Theory).

While sharing a group membership may be beneficial in the evolution and survival of communities (e.g., supporting the fulfillment of both intra-community and intra-individual needs; see A. L. Johnson et al., 2006), it has also fostered the tendency of favoring in-group members over out-group members, a phenomenon known as **in-group bias** (see Hewstone, Rubin, & Willis, 2002; Dunham, 2018, for extensive reviews). In-group bias results in a series of beliefs that determine how groups interact, but also how single individuals behave with each other. Not only are out-group members associated with positive traits to a lesser extent (e.g., Dovidio, Kawakami, Johnson, Johnson, & Howard, 1997; Dovidio & Gaertner, 2000) and deemed less trustworthy than in-group members (e.g., Insko, Schopler, Hoyle, Dardis, & Graetz, 1990; Insko, Schopler, & Sedikides, 1998), they also tend to be considered as a threat to the survival of the in-group (see Stephan & Stephan, 2000).

This tendency could be reflected in the increasing nationalist sentiment found across many of the EU member states in reaction to the recent refugee movements and migration waves that reach the European shores every day. The growing number of migrants and asylum-seeking refugees in the EU, who are typically seen as out-group members, has been perceived as a threat by many European citizens and governments - strengthening the distinction between *us* and *them*. Together with economic and political reasons, the perceived threat led many to support policies that reinforce the control of the borders, instead of showing solidarity with asylum-seekers. Arguably, the migration issue is a far more complex phenomenon than what I describe here, and it must not be simplified. Nevertheless, here I make use of this analogy to illustrate how the discrimination of in-group vs. out-group members can have cascading effects not only on the single individuals but also on the society.

By drawing the analogy with the current socio-political scenario in the EU, I have introduced some of the cardinal aspects of the well-studied phenomenon known as in-group bias, which is central to this thesis, entitled **An exploration of social biases in language processing**. The research reported in this thesis attempts to bridge knowledge from two distinct fields, namely, social psychology and psycholinguistics. My intention is to provide a better understanding of how the basic cognitive mechanisms underpinning language learning and language processing may be influenced by social factors, such as the group membership status of speakers from whom new linguistic information is learned and learners' own in-group biases.

The need for such a research agenda comes from the fact that many models of language processing posit that linguistic input is stored in the form of abstract units (i.e., word forms or phonological representations) and information specific to speakers and contexts is left out (abstractionist accounts, e.g., Gaskell & Marslen-Wilson, 1997; McClelland & Elman, 1986; Norris, 1994). There are several reasons why such models appear to offer a rather incomplete account of how people process language.

First, most of the knowledge we acquire throughout our lives is in fact learned by observing and interacting with others in social contexts, a phenomenon known as *social learning* (Bandura & Walters, 1977) (see Heyes, 1994, 2011, 2012, for recent accounts). Arguably, much of this knowledge we learn from others eventually helps us deliver information and content to third parties. Since very early on in life, we sharpen our abilities to produce speech sounds and to refer to our surrounding world by copying others' labels and syntactic structures, and by receiving feedback from those others. Due to its natural unfolding in social contexts, language is, in fact, considered a clear-cut example of socially learned system (e.g., usage-based theory, e.g., Tomasello, 2000; Lev-Ari, 2016). Social interactions strongly contributed in shaping the human cognitive system to its current form in ways that supports language use (e.g., Levinson, 1995, 2006).

Second, it has been claimed that it is the tendency to store and continuously update information about partners in conversation (e.g., their perspectives, identity, knowledge and beliefs) that allows people to establish a common ground [1] (see Brown-Schmidt, Yoon, & Ryskin, 2015, for a theoretical account). Crucially, the emergence of such a common ground, based on partner-specific representations, would facilitate ongoing conversations, but also encounters with the same interlocutors in the future.

Finally, there is growing empirical evidence that indexical information, such as a speaker's social identity and voice characteristics, can influence language processing at many different levels (e.g., Hay, Nolan, & Drager, 2006; Hay, Warren, & Drager, 2006; K. Johnson, Strand, & D'Imperio, 1999; Lev-Ari & Keysar, 2012; Niedzielski, 1999; van Berkum, Van den Brink, Tesink, Kos, & Hagoort, 2008).

Considering these arguments, it is plausible that information about speakers and the way listeners socially perceive them is maintained in some way while processing the language produced by those speakers. What is still unclear is whether, and how, the perceived status and social relevance ascribed to speakers by listeners can affect how language is represented and used. Existing sociolinguistic evidence suggests that social factors, such as speakers' likeability or prestige, may mediate people's tendency to imitate others' speech (e.g., Babel, 2012; Heyselaar, Hagoort, & Segaert, 2017; Lev-Ari & Peperkamp, 2014; Lev-Ari, San Giacomo, & Peperkamp, 2014; Lev-Ari, 2016). These findings provide some support that social factors may actually influence how language is used. However, it is not clear yet how and via which cognitive mechanisms these social factors determine how language is *represented* in the mind.

Departing from these findings, I tested the prediction that learning language from speakers to whom a special social status is ascribed can influence how lexical representations are encoded and retrieved. Since in-group membership (vs

[1] With *common ground*, I mean here the mutual understanding and shared assumptions that are built among partners involved in a conversation (e.g., Clark & Carlson, 1981; Clark & Marshall, 1981)

out-group membership) has been shown to affect basic cognitive processes such as memory and learning (see below), group membership status offered a case study to test the hypothesis that language learning and language processing may also be influenced by the social relevance attributed to the speakers.

Previous research

Growing evidence has shown that the identity of the speaker can influence speech perception (e.g., Hay, Nolan, & Drager, 2006; Hay, Warren, & Drager, 2006; K. Johnson et al., 1999; Niedzielski, 1999) and lexical access (e.g., Kim, 2016; Walker & Hay, 2011). For instance, Niedzielski (1999) manipulated the information made available to the participants regarding the speakers' place of residence and asked them to identify diphthongs produced by those speakers. Despite the linguistic material being constant, participants reported hearing what they believed was more representative of the linguistic community of the supposed speakers. For example, when told the speaker was from Canada, they would report hearing raised-diphthong tokens which are stereotypically seen as a feature of the Canadian dialect. Likewise, Walker and Hay (2011) used an auditory lexical decision task and showed that, as compared to neutral-age words, listeners were faster to respond to words uttered by speakers that represented the stereotypical word users, based on the age inferred by their voice. For instance, when a young speaker uttered the word *lifestyle* – rated as typically produced by young people, as compared to the word *electricity* – rated as neutral, the former was faster responded than the latter.

In the same manner, the social characteristics of speakers have been shown to shape listeners' expectations about upcoming linguistic input (e.g., van Berkum et al., 2008; Boland & Shana'e, 2014; Martin, Garcia, Potter, Melinger, & Costa, 2016). For instance, in an EEG study by van Berkum and colleagues

(2008) participants listened to utterances produced by speakers from different social categories (e.g., female/male; upper/lower-class; child/adult). The word *sandcastles* produced by an adult speaker in a sentence like "On the beach I made *sandcastles* by the sea" elicited more negative N400 amplitudes than when produced by a child speaker, reflecting a relative cost in integrating this information in the context, during online sentence comprehension. The reported N400 effect and the classical N400 effect, typically seen for semantic violations, had similar topographical and temporal distributions (see Kutas & Federmeier, 2011, for a review). Such findings were interpreted as showing that the social identity of the speaker may be involved from early on in the same mechanisms responsible to comprehend the meaning of the sentence from the linguistic input (see also Boland & Shana'e, 2014; Martin et al., 2016).

Social identity of the speakers influences language processing

The evidence provided in the previous section challenges models assuming linguistic representations to be abstract units. Therefore, there is a growing consensus about the necessity to model language processing in a way that accounts for the encoding of both linguistic and contextual information (Cai et al., 2017; Hay, Nolan, & Drager, 2006; Kapnoula & Samuel, 2019; Münster & Knoeferle, 2018; Sumner, Kim, King, & McGowan, 2014) (see Drager & Kirtley, 2016, for a review).

Many proposals support the assumption that linguistic experiences may be stored in the mental lexicon as *exemplars*, that is, episodic memory traces that contain information related to both speaker/context characteristics and linguistic input (e.g., Goldinger, 2007; Sumner et al., 2014).

In a mechanistic model by Münster and Knoeferle (2018), it is claimed that on-line language processing may benefit from the concurrent encoding of speaker identity and linguistic information. According to this model, listeners calculate probabilistic contingencies regarding the language typically associated with the social group of the speaker. Following such speaker-driven probabilities, together with those guided by the linguistic input per se, a set of lexical candidates can be pre-activated. When those items are encountered, language comprehension is speeded up and eased, like in the case of the word *sandcastles* produced by a child voice (see Münster & Knoeferle, 2018, for details).

Variation in the strength of encoding: a case for in-group biases

While the findings mentioned so far show that linguistic information is represented along with information about the speaker's identity, an open issue remains: is the concurrent encoding of speaker- and language-related information equal across all speakers or does it vary as a function of the social relevance that the speakers are attributed?

A recent model proposed by Sumner and colleagues (2014) claims that the social context is not only stored together with the linguistic input but can also affect the strength of its encoding. To support this account, Sumner et al. refer to evidence that contrasts reaction times and memory accuracy in relation to phonetic variants that are either idealized, yet non-frequent, or non-idealized, but more frequent. Idealized variants are processed and remembered as fast and as accurately as non-idealized variants, despite the difference in frequency of occurrence suggesting otherwise. These results are interpreted as indicating that the strength of encoding of linguistic variants may differ, with idealized variants, which receive a higher social weight, being encoded more strongly than non-idealized variants.

A prediction from Sumner's model is that the social weight attributed by language users to contexts and speakers would determine the strength of the encoding of linguistic input. This also means that the social weight that is ascribed may vary among individual language users as a function of their own social perception of those speakers. In this thesis, I looked at differences in encoding due to speakers' group membership status to test this prediction. Specifically, one of my questions was whether the group membership status of the speakers from whom participants learned novel words, in interaction with participants' own in-group bias, would affect the encoding of the content in the lexical representations.

The indirect evidence in support of the prediction that in-group language may be encoded more strongly than out-group language comes mostly from studies in experimental social psychology that reports memory and processing superiority for in-group vs out-group related information (e.g., Frable & Bem, 1985; Hugenberg, Young, Bernstein, & Sacco, 2010; Van Bavel, Packer, & Cunningham, 2008; Wilder, 1990). For instance, in old/new recognition tasks, people tend to be better at recognizing old in-group faces than old out-group faces (e.g., Van Bavel et al., 2008; Hugenberg et al., 2010). Similarly, in source memory tasks, they remember the source of information (i.e., who said what) more accurately when it is an in-group vs out-group member (e.g., Frable & Bem, 1985; Greenstein, Franklin, & Klug, 2016; Wilder, 1990). These findings suggest that the input related to in-group membership receives benefits during processing that lead to advantages in memory.

Along the same lines, Moradi, Sui, Hewstone, and Humphreys (2015) used a perceptual matching task in which participants first learned arbitrary associations created by pairing neutral stimuli (i.e., geometrical shapes) with either the logo of their favorite football team (i.e., in-group team) or the logo of another football team (i.e., out-group team). Then, participants performed a match/mismatch task where they saw shape-logo pairs and had to judge whether they matched what had been learned. Results showed that they responded to in-group related associations in a faster and more accurate way, as compared to associations that involved the out-group team (see also Enock, Sui, Hewstone, & Humphreys, 2018; Moradi, Sui, Hewstone, & Humphreys, 2017). Such findings support the claim that in-group membership can enhance general input processing and learning.

One of the mechanisms proposed as the source of in-group biases is **attentional prioritization** (e.g., Van Bavel & Cunningham, 2012). According to such a hypothesis, the input related to in-group membership, which is considered as highly salient, captures attentional resources to a larger degree than out-group input does. As a consequence of such extra allocation of attention, in-group input is perceptually processed to a deeper extent than out-group input, resulting in memory superiority for the former.

Although the attention hypothesis sounds plausible, the evidence supporting it is mixed. Some studies demonstrated that attention is oriented to in-group members in a preferential manner (e.g., Chauhan, Visconti di Oleggio Castello, Soltani, & Gobbini, 2017; Park, Van Bavel, Hill, Williams, & Thayer, 2016; Van Bavel & Cunningham, 2012). On the other hand, there exists also some evidence that conflicts with such an interpretation, where faces of out-group members were stronger attentional cues than in-group faces (e.g., Brosch & Van Bavel, 2012; Park et al., 2016). Considering the mixed evidence, it is unclear whether and how the allocation of attentional resources is modulated by group membership status. One of the goals of the current thesis was also to shed further light on this issue to better understand the relationship between attention and in-group biases.

Independently of attention being the underlying source of in-group biases, the overall evidence supports that in-group membership boosts input processing and general learning. Understanding whether in-group biases affect also language learning and language processing is the aim of this thesis.

The current thesis

Working framework

Departing from the above-mentioned evidence, I developed a working framework and generated predictions about how speakers' group membership status and learners' individual biases may affect the representation and processing of novel words. The working model is based on three main working hypotheses:

1. Greater individuation of in-group information.
2. Facilitated lexical retrieval for in-group words.
3. Attentional prioritization by in-group membership.

Greater individuation of in-group information

Evidence for superior source memory and face-recognition for in-group vs out-group members (e.g., Frable & Bem, 1985; Van Bavel et al., 2008; Wilder, 1990) suggests that the former may be encoded in more individual-specific details than the latter (see Hugenberg et al., 2010, for a theoretical account). Such increased individuation of in-group members, together with the general tendency to find out-group members similar to each other (i.e., homogeneity effect, see Ostrom & Sedikides, 1992, for a review) led to the prediction that the representations of words learned from in-group speakers may be more speaker-specific than representations of out-group words. Out-group representations are in fact predicted to include mostly group-based details (see Figure 1.1, for a simplifying graphical illustration of the hypothesis).

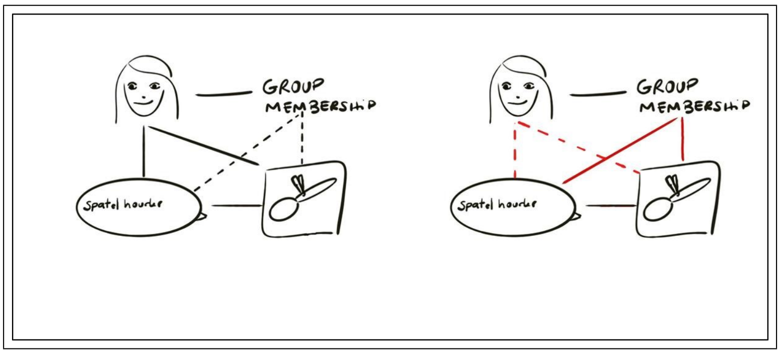

Figure 1.1: Illustration of the individuation hypothesis. A representation of a word learned from an in-group speaker (left), and a representation of a word learned from an out-group speaker (right). Both contain details about the linguistic input (i.e., the label) and its visual referent (i.e., the picture of the gadget). Additionally, the representations contain information about the speaker from whom the label was learned. In the in-group representation it is speaker-specific (personal and affiliative identity), whereas in the out-group representation this information is only group-based (affiliative identity)

Additionally, based on evidence showing that people differ in how strongly they prefer in-group over out-group members (e.g., Amodio, Harmon-Jones, & Devine, 2003; Hein, Silani, Preuschoff, Batson, & Singer, 2010; Platow, McClintock, & Liebrand, 1990; Van Bavel, Swencionis, O'Connor, & Cunningham, 2012), I predicted that the effect of speakers' group membership status on the individuation of lexical representations would interact with participants' individual in-group bias. Specifically, the stronger the individual bias, the larger the differences in individuation are expected to be.

Facilitated lexical retrieval for in-group words

Following the evidence showing that in-group membership boosts a) the learning of stimuli associations (e.g., Moradi et al., 2015, 2017) and b) the ability to recognize old stimuli (e.g., Van Bavel et al., 2008; Hugenberg et al., 2010), the prediction that follows is twofold. First, I predicted that words learned from in-group speakers may be more strongly encoded and associated with contextual information (e.g., source and pictorial referent), as compared to out-group words. Therefore, I predict that in a paradigm where representations are reactivated by contextual cues (e.g., the presentation of speakers' faces and gadgets), the associated lexical item will be retrieved in an easier and/or faster way. By looking at N400 amplitudes as a marker of the ease with which lexical items are processed, I expected to see differences between in-group and out-group words.

As in the individuation hypothesis, listeners' individual in-group bias is expected to modulate the difference in N400 effects elicited by in-group and out-group words, with the stronger the in-group bias, the larger the difference.

Attention prioritization by in-group membership

Given the mixed findings related to the hypothesis that attentional priorities to in-group membership may be the underlying source of in-group bias, one of the goals of the current thesis was to further test this possibility. The working hypothesis is that in-group membership would be perceived as highly salient, and therefore attentional resources would be automatically captured and devoted to processing in-group related input (vs out-group input).

Thesis outline and methods

I present here the results of four experimental chapters, which were designed to test the working hypotheses outlined in the previous section. **Chapters 2 and 3** examined the role of speakers' group membership status and learners' in-group bias on the encoding of speaker-specific information in the representations of novel words. In these chapters I used a **newly developed word learning paradigm combined with a source memory test**.

To design the stimuli, I previously carried out a multiple-step norming procedure (see Appendix 1) where gadgets were selected for eliciting at least a pair of competing labels. In the next step, I selected pairs of labels that had similar frequency of occurrence and goodness-of-fit ratings. These preliminary steps were critical because they allowed me to test for group membership effects on lexical representations while controlling for amount of prior exposure, goodness-of-fit of the labels and frequency of occurrence.

In the actual studies, participants learned pairs of competing labels for the selected gadgets. Crucially, for target items, one label was produced by one of the in-group speakers and the other label was produced by an out-group speaker. Participants' source memory for the novel words was then tested. In this test, they saw speaker-label pairs and had to decide whether the displayed speaker had produced the label in the previous learning phase. All speaker-label combinations were presented, so participants saw matches and mismatches. Among the mismatches, there were speaker-label pairs where the speaker belonged to the same affiliation of the correct speaker (within-affiliation mismatches), and pairs in which the speaker belonged to the other affiliation (between-affiliation mismatched).

When asked to decide if the pair matched or not with what they saw in the previous phase, participants were expected to produce errors that reflected the fact that they encoded information about the group membership status of the speaker (within-affiliation source memory confusion). This confusion, though, was expected to be less prominent for in-group information as a consequence of the stronger individuation of in-group labels. The source memory test functioned as a proxy to test this hypothesis and to investigate the level of detail of speaker-specific information encoded in the novel lexical representations. To test whether differences between in-group and out-group source memory were predicted by how strongly participants preferentially processed their in-group over the out-group membership, individual measures of in-group bias were also collected via an implicit in-group bias task.

By using a similar **word learning paradigm but in combination with an N400 task**, **Chapter 4** examined the role played by the group membership status of the speakers on the ease of integration of newly acquired words, preceded by in-group vs out-group cues. The main difference in the learning design was that participants only learned one label per gadget. Half of them were named by in-group speakers and the other half was named by out-group speakers. On the following day, participants came back for an EEG session where they saw a series of cueing stimuli (speakers' faces + gadgets) followed by a spoken label. The label was either the one they learned on the previous day (i.e., trained label) or a novel, equally fitting, one they did not learn (i.e., untrained label). I targeted the N400 effect (i.e., the difference in amplitude between untrained and trained labels) as indexing relative integration of trained vs untrained stimuli in the context. Differences due to group membership status of the labels were expected in the magnitude of the N400 effect. Additional behavioral measures during both the learning phase, the EEG task and afterwards were collected to shed further light on potential in-group biases.

Finally, in **Chapter 5** I tested the effect of faces' group membership on the orientation of spatial attention using a **Posner cueing task**, a non-linguistic task. In this task, faces of in-group and out-group members were used as central cues to subsequent peripheral targets. Crucially, the cues were, in equal measures, either valid, invalid or neutral in predicting the upcoming target location. The three conditions allowed to test for facilitatory and/or inhibitory effects of group membership on attention orientation.

In **Chapter 6**, I summarize the main findings, discuss theoretical implications, and highlight potential avenues for future research.

2 | How in-group bias influences source memory for words learned from in-group and out-group speakers[1]

Abstract

Individuals rapidly extract information about others' social identity, including whether or not they belong to their in-group. Group membership status has been shown to affect how attentively people encode information conveyed by those others. These findings are highly relevant for the field of psycholinguistics where there exists an open debate on how words are represented in the mental lexicon and how abstract or context-specific these representations are. Here, we used a novel word learning paradigm to test our proposal that the group membership status of speakers also affects how speaker-specific representations of novel words are. Participants learned new words from speakers who either attended their own university (in-group speakers) or did not (out-group speakers) and performed a task to measure their individual in-group bias. Then, their source memory of the new words was tested in a recognition test to probe the speaker-specific content of the novel lexical representations and assess how it related to individual in-group biases. We found that speaker group membership and participants' in-group bias affected participants' decision biases. The stronger the in-group bias, the more cautious participants were in their decisions. This was particularly applied to in-group related decisions. These findings indicate that social biases can influence recognition threshold. Taking a broader scope, defining how information is represented is a topic of great overlap between the fields of memory and psycholinguistics. Nevertheless, researchers from these fields tend to stay within the theoretical and methodological borders of their own field, missing the chance to deepen their understanding of phenomena that are of common interest. Here we show how methodologies developed in the memory field can be implemented in lan-

[1]Adapted from *Iacozza, Meyer, and Lev-Ari (2019). How in-group bias influences source memory for words learned from in-group and out-group speakers. Frontiers in Human Neuroscience, 13.*

guage research to shed light on an important theoretical issue that relates to the composition of lexical representations.

Introduction

Previous findings have shown that people utilize any cue they have available (e.g., gender, social class) to establish whether or not others are members of their own in-group (e.g., Bargh, Schwader, Hailey, Dyer, & Boothby, 2012). Group membership can affect how people process and remember information related to those others, with in-group information receiving more attention and being better remembered than out-group information (Hugenberg et al., 2010; Greenstein et al., 2016). While advantages for in-group members have been reported to affect a wide range of cognitive phenomena (see Xiao, Coppin, & Van Bavel, 2016; Molenberghs & Louis, 2018, for reviews), they have not been directly tested in the context of language processing and language learning, yet.

Such effects are relevant for models of language processing because they have consequences for an ongoing debate on how words are represented in the mental lexicon. One aspect of this broad issue is how well listeners maintain information that is not strictly linguistic but that relates to the context, such as the social identity of the speaker producing a word. In the memory literature, this type of information is referred to as source memory and it is a topic that has been extensively studied. By using memory tests developed to probe source memory, researchers in the field of psycholinguistics can gain a better understanding of how speaker-related information is encoded in the representations of words, and whether the encoding of such information is modulated by social factors, such as the group membership status of the speakers.

The aim of the current study is to investigate the proposal that in-group biases permeate language processing as well, and that they affect the level of detail of speaker-related information that is encoded when learning new words. We propose that representations of words learned from in-group members are more likely to contain highly specific speaker-related information, as compared to representations of words learned from out-group members, and that such differences are in turn influenced by how strongly each learner prefers their in-group members over out-group members.

Before turning to the current study, we review the relevant literature. We start by reporting evidence that shows that the social identity of the speaker affects how listeners process language. We then describe existing exemplar-based theories of language processing that provide a theoretical framework for understanding effects of speaker identity on language processing. We then point to a potential limitation of these models, namely, their tendency to assume that the speech of all speakers is treated equally. We propose that existing models should integrate parameters that allow different degrees of encoding specificity and assigning different weight to linguistic input depending on speaker group membership status. Specifically, we propose that linguistic information provided by in-group speakers is encoded in more detail than information from out-group speakers. We motivate our proposal with evidence from non-linguistic studies in social psychology that report group membership effects on memory and information processing.

Previous research indicates that when interacting with others, information about their social identity is rapidly extracted (see Bargh et al., 2012, for a review) and can influence people's attitudes and preferences towards those others (e.g., Greenwald & Banaji, 1995; Jones & Fazio, 2010; Kinzler, Corriveau, & Harris, 2011). There exists diverse evidence showing that others' social identity can influence how listeners process language. For instance, it has been shown that, when a speaker's social identity is made available via the speaker's voice, listeners take the identity into consideration and have particular expectations about what will likely be said. If these expectations are not met, such as when the desire of looking like Britney Spears is reported in a man's voice, language processing becomes harder (van Berkum et al., 2008; Martin et al., 2016; Walker & Hay, 2011). Similarly, speaker social identity can affect how listeners perceive speech sounds (e.g., Hay, Nolan, & Drager, 2006; Hay, Warren, & Drager, 2006; K. Johnson et al., 1999; Niedzielski, 1999). For example, changing listeners' expectations of a speaker's place of residence affected their responses in a diphthong identification task. Participants reported hearing what they believed to be more representative of the supposed speaker's linguistic community, independent of the actual linguistic input, which was identical across the two conditions (Niedzielski, 1999). This suggests that information about the speaker affects speech perception.

In short, this body of evidence shows that information related to the speaker's identity is extracted along with the linguistic input and can influence the processing of the latter. Existing exemplar-based models of speech processing argue that the reason that social information is used in language processing is because it is encoded along with linguistic input. These models state that linguistic experiences are encoded as rich episodic memories (i.e., exemplars) (e.g., Goldinger, 2007; Hay, Nolan, & Drager, 2006; Hay, Warren, & Drager, 2006) that contain information which is both language-specific (e.g., includes phonetic, lexical, and syntactic details) and context-specific (e.g., includes pragmatics, speakers' characteristics) (see Drager & Kirtley, 2016, for a detailed review).

Recently, in a new model by Münster and Knoeferle (2018) the contributions of encoding speakers and listeners' characteristics during on-line language processing were formally defined. Grounded on a large body of empirical evidence, the model posits that comprehending language in context, by, for instance, extracting both speaker-specific and language-specific input in tandem, may speed up and/or ease comprehension. For example, consider a scenario in which the utterance "Every evening I drink some wine before I go to sleep" is produced by an adult speaker. Based on the age of the speaker, listeners can build up probabilistic inferences about what is more likely to follow the verb *drink* (e.g., in the case of an adult, the word *wine* is more probable than the word *milk*). By pre-activating lexical items that are more probable, listeners can easily make sense of the new piece of information, i.e., the word *wine*, speeding up comprehension (see Münster & Knoeferle, 2018, for details).

Crucially, Sumner et al. (2014) proposed that the social context might not only be encoded with the linguistic input but might modulate the strength of its encoding. In support of this account, Sumner and colleagues showed that idealized phonetic variants are encoded with greater weight than common, therefore more frequent, phonetic variants. According to their model, phonetic variants with higher prestige (i.e., idealized ones) receive an advantage in representation and processing as compared to variants characterized by lower prestige. Extending their theory to more general linguistic processes, one could hypothesize that people would encode linguistic variations more strongly if they are associated with contexts and speakers that have a special status.

Here, we propose that learning new words from speakers that are ascribed a special status might lead to lexical representations that are richer in contextual information (e.g., speaker-related information), as compared to representations of words learned from speakers without a special status. An example of speakers that are ascribed a special status is the case of in-group members. Indeed, there is evidence suggesting that group membership influences input processing and learning. For instance, memory is usually better for in-group faces than out-group faces (e.g., Van Bavel et al., 2008; Hugenberg et al., 2010) and for information delivered by in-group than by out-group members (e.g., Frable & Bem, 1985; Wilder, 1990). Furthermore, people learn better and process more quickly new associations between previously neutral stimuli (e.g., geometrical shapes) and in-group membership (e.g., the logo of their favorite football club) than associations involving out-group membership (Enock et al., 2018; Moradi et al., 2015).

One way in which in-group biases may work is via the recruitment of additional cognitive resources (Meissner, Brigham, & Butz, 2005; Van Bavel & Cunningham, 2012). Such additional resources have been suggested to lead to in-group representations that are characterized by a higher level of detail than out-group representations. For example, when processing in-group related information, people were shown to encode the source of information in more detail than when the information was related to out-group members. This resulted in them being better in a source memory task when identifying in-group sources than out-group sources (e.g., Greenstein et al., 2016), suggesting that being exposed to in-group membership boosts the encoding of individual-specific information (see Hugenberg et al., 2010, for a similar account).

No study has tested whether lexical representations for the same words can depend on the identity of the speaker that tends to use them. If this is the case, this will have implications for language learning, language processing, and linguistic representations. It would extend current theories that examine the role of input and its distribution in language acquisition and representation by showing that the same distribution can have different effects depending on who are the speakers that provide different tokens in the input. As a first step, the current study was designed to investigate which social information learners encode when they learn new words from speakers who either belonged to the learners' social group (i.e., in-group members) or did not do so (i.e., out-group members).

The current study

We hypothesize that listeners encode the social identity of the speakers from whom they learn novel words and that the social identity influences how detailedly speaker-specific information is encoded. To test these predictions, we carried out the current study in which we examined participants' source memory for words learned from speakers from different social groups. In the Main Experiment, participants were exposed to a learning context in which they learned new words from speakers who supposedly shared their university affiliation (i.e., in-group speakers) and from speakers with a different affiliation (i.e., out-group speakers). In the Control Experiment, participants learned from two groups of speakers who supposedly attended two foreign universities. Since in the Control Experiment the group membership was not manipulated, because both universities were unrelated to the participants, we could check that the patterns hypothesized to be found in the Main experiment were indeed a reflection of the social saliency ascribed to speakers' group membership and not simply a consequence of the contrastive nature of our manipulation (i.e., teaching competing labels spoken by different groups of speakers).

During the word learning task, all participants in both experiments learned novel labels for uncommon gadgets. Crucially, target gadgets received two competing but equally fitting labels, one from a speaker of each affiliation (e.g., *citrus-peller* vs *citrus-schiller*, in English *lemon peeler* vs *lemon stripper*). Afterwards, source memory for these words was tested in a recognition memory test. Participants were shown one speaker and one label at a time and asked if the speaker had produced the label in the previous phase (i.e., forced choice: yes/no). Lastly, we collected participants' implicit in-group bias (see Materials and Methods section for details).

In the Main experiment, we predicted that participants would spontaneously monitor the speakers' group membership status. Consequently, when asked to recognize the source of the new words, we expected participants to remember speaker social group but to struggle remembering the exact speaker that produced each word. Therefore, they should be more likely to misattribute words to incorrect speakers within the same affiliation than between different affiliations, i.e., there should be source memory confusion. Following our hypothesis about different levels of detail depending on social salience and group membership of the speakers, we predicted that words learned from in-group speakers would contain a higher level of detail about who produced them, compared to words learned from out-group speakers. This would result in in-group linguistic representations that are more speaker-specific and less prone to source memory confusion than out-group representations. Crucially, this in-group advantage should be stronger for participants exhibiting stronger in-group bias. This pattern is expected to result in a significant interaction involving speaker group membership and individual in-group bias. In the Control Experiment, we expected no differences between the two speaker affiliations. This would show that differential processing of information learned from different groups is specific to cases where group membership is socially salient.

Materials and Methods

Participants

One-hundred-twenty-four native Dutch speakers (age range: 18-26 years) participated in the study after providing their informed consent, as approved by the Ethics committee of the Social Sciences department of Radboud University (project code: ECSW2014-1003-196). All participants were students or recent graduates of Radboud University Nijmegen. All participants were female, as were the speakers from whom they learned the labels. This was done to avoid that an additional social dimension (i.e., gender) of in-group status could interact with the one we manipulated (i.e., academic affiliation). Participants were randomly assigned to either the Main Experiment (n=62) or the Control Experiment (n=62).

Materials

Materials for the word learning task

Speakers: Eight fictitious speakers were created by pairing female faces selected from the Chicago Face Database (Ma, Correll, & Wittenbrink, 2015) with the voices of native Dutch female speakers recorded in our laboratory. Prior to the experiment, voices were matched for perceived typicality and attractiveness (paired t-tests, ps>.05) via a norming on-line survey in which twenty different participants participated. Each speaker was a unique combination of one face and one voice, consistent across participants. Speakers' academic affiliation was randomized across participants and indicated by the logo of the supposed affiliation displayed underneath the photo.

Affiliation logos: For the Main Experiment, original-color pictures of the logos of Radboud university (i.e., in-group affiliation) and ROC Nijmegen (i.e., out-group affiliation) were used. For the Control Experiment, original-color pictures of the logos of Pisa and Florence universities were used.

Gadgets and labels: Twenty-four images of unfamiliar gadgets (e.g., lemon peeler) and their corresponding labels were selected via a norming study (see Appendix 1, for details). Half of the gadgets, hereinafter referred to as target gadgets, were presented with two competing labels, which were equated for goodness-of-fit and frequency. The other twelve gadgets were presented with a single label and served as fillers. All labels were produced by each speaker and audio-recorded.

Materials for the individual in-group bias task

Affiliation logos: We used the logos used in the word learning task.

Geometrical shapes: We used black shapes for triangle, square and circle.

Procedure

Word learning task

The word learning task consisted of an exposure phase and a test phase. The exposure phase was presented as a communication task in which participants were instructed to pay attention to all the stimuli presented (i.e., faces, gadgets and labels) and to select gadgets based on what the speakers said. There was no explicit reference to the academic logos. Participants saw 24 gadgets, each named by speakers of both groups. Half were target gadgets, for which the two groups of speakers provided competing, but equally fitting, labels, whereas the other half were fillers, for which unique labels were provided. Fillers were included to minimize participants' awareness of the nature of the experimental manipulation (i.e., the contrastive nature of the labels). Note that not all speakers referred to all the gadgets. In fact, each gadget was only labeled by two of the eight speakers (one per group of speakers). Speaker group affiliation, speaker-label pairing, and label-group affiliation pairings were fully randomized per participant.

On each trial, a photo of a speaker, together with the corresponding affiliation logo, was displayed (800 ms). Then, while the photos of speaker and logo were still on screen, the audio-recording related to the gadget label was played. Simultaneously, the written form of the label was superimposed upon the speaker's mouth (1500 ms). Next, three gadgets appeared on the screen and participants selected the one that fit the audio and the written label (see Figure 2.1 for an example of the learning display [2]). If the response was wrong, the audio was repeated. Two exposure blocks were administered with half of the gadgets (i.e., six fillers and six targets) introduced in the first block, and the other half introduced in the second block. The gadgets were randomly allocated in the first or second exposure block per participant. Three exposure rounds were administered per block so that each display was repeated three times, once per round, in a randomized trial order.

Figure 2.1: Example of the learning display in the word learning task. Participants had to select the gadget that was mentioned. In this case, they had to select the first image. Stimuli are not drawn to scale.

After each exposure block, participants performed a surprise source-memory recognition test on the gadgets introduced in the preceding exposure block only. In each trial, they saw a photo of a speaker with their affiliation logo and a written label (see Figure 2.2). Participants indicated whether the speaker had produced the label in the previous exposure phase via key press (i.e., forced

[2] Due to copyright issues, none of the pictures of the gadgets in the example corresponds to actual stimuli, but they provide a good approximation of the type of stimuli we used.

choice: yes/no). Decisions were self-paced. Across the two memory test blocks, there were 288 trials in which all possible speaker-label pairings were shown. Of those 288 trials, 96 were filler-related trials (subsequently excluded from the analyses) and 192 were trials in which target gadgets were shown. Of the 192 target-gadget trials, 24 were matching trials (i.e., the speaker had indeed produced the label) and 168 mismatching trials (i.e., the label had not been used by the speaker). Of those mismatching trials, 72 were within-affiliation mismatching trials (showing a label along with a wrong speaker from the same affiliation as the correct one), 72 were between-affiliation mismatching trials (showing a label a speaker from the wrong affiliation). The remaining 24 trials showed a speaker with a label that competed with the one she used (e.g., the speaker that had used "citrus-schiller" was displayed with "citrus-peller"). They were only included to make all possible speaker-label combinations available, but they were not analyzed. Note that in all mismatching trials, the correct answer was that the pairing was incorrect because the speaker depicted in the photo had not used the displayed label in the exposure task.

Figure 2.2: Example of a trial in the source memory task. Participants indicated if the speaker had produced the label in the exposure task. Stimuli are not drawn to scale.

Implicit in-group bias task

Participants' individual in-group bias was measured in a perceptual matching task (Moradi et al., 2015), which has been shown to provide results that are reliable within individuals and across different test sessions (Stolte, Humphreys, Yankouskaya, & Sui, 2017). Three geometric shapes (circle, square, triangle) were randomly paired with logos of three academic affiliations. For the Main experiment, the logos depicted the in-group university –Radboud University, and two out-group affiliations –ROC Nijmegen and Tilburg University. To keep the two experiments comparable, participants in the Control experiment performed the task with logos of the Italian universities that appeared in the word learning task (Pisa and Florence) and a third Italian university, Bologna. Each association was initially presented ten times. Then, participants performed a practice block of 24 trials, followed by two blocks of 120 experimental trials each. In both practice and test trials, a fixation cross (500 ms) preceded a blank screen (between 1000 and 2000 ms) and the simultaneous presentation of logo and shape (600 ms), following the timings utilized in Moradi et al. (2015). Participants had 1500ms to judge the accuracy of the pairing. Feedback was given only during practice. In-group bias in this task is usually indexed by faster and more accurate responses for stimuli that are newly associated with in-group membership compared to stimuli associated with out-group membership (e.g., Moradi et al., 2015).

Results

All analyses were performed with mixed-effects modelling as implemented in the lme4 package (Bates, Maechler, Bolker, & Walker, 2014) in R (R Core Team, 2016) and the models' random structures were determined following the procedure suggested by Bates, Kliegl, Vasishth, and Baayen (2015).

Before turning to the main analyses from the source memory test, we performed a sanity check to confirm that, at the group level, participants in the Main experiment showed the expected in-group bias in the perceptual matching task used to extract individual in-group bias measures.

Group-level in-group bias

Analyses over RTs

Prior to analyses, trials with incorrect responses or with RTs faster than 200ms or slower than 2100ms were excluded. For these sanity-check analyses, we selected only matching trials (i.e., in which the logo of the university was displayed with the associated geometrical shape) which referred to the in-group university and the out-group university used in the study (i.e., ROC Nijmegen). We then performed an outlier removal procedure by removing trials with RTs 2.5 SDs or higher from the mean per condition, per participant. The resulting dataset was analyzed using linear mixed-effect model in which log(10)-transformed RTs were predicted by the fixed effect for Group Membership (In-group vs Out-group, reference level: In-group). We added per-participant random intercept and by-participant random slope for Group Membership. Results confirmed the usual patterns for this task: participants were faster at recognizing in-group-related associations than out-group-related associations (in-group: mean=709 ms, SD=212 vs out-group: mean=754 ms, SD=199; beta=-0.01, SE=0.003, t=-5.03, $p < .0001$).

Analyses over Accuracy

As with the RT analysis, the analysis included only matching trials (i.e., trials in which the logo of the university was displayed with the associated geometrical shape) which referred to the in-group university and the out-group university used in the study (i.e., ROC Nijmegen). Accuracy was analyzed using a logistic mixed-effect model with a fixed effect for Group Membership (In-group vs Out-group, reference level: In-group). We added per-participant random intercept and by-participant random slope for Group Membership. Results confirmed the usual patterns for this task: participants were better at recognizing in-group-related associations than out-group-related associations (in-group: mean=94.70% SD=22.4 vs out-group: mean=92.88% SD=25.72; beta=0.4, SE=0.1, t=3.17, $p < .01$).

The analyses confirmed that in the Main experiment, at the group-level, participants showed a strong in-group bias for their own university. Successively, we extracted individual measures of in-group bias by calculating a per-participant measure of effect size, namely Cohen's d, from both accuracy and RTs over in-group versus out-group matching trials. The measure calculated over RTs was not a significant predictor in any of the models we ran; thus, we will focus on the measure derived from accuracy.

Next, the results from the Main and Control experiments are presented separately because the in-group vs out-group contrast only applies to the former experiment. The data from each experiment was analyzed following the outlined steps: (1) planned analyses on matching and mismatching trials, separately; (2) post-hoc analyses over d-prime and response bias values.

Main Experiment

After each exposure round in the word learning task, participants were tested with a recognition memory test. In this test, they were presented with matching or mismatching speaker-label pairings and had to decide via key press if the label had or had not been produced by the speaker. We carried out analyses over matching and mismatching trials separately. We predicted that participants would show more accurate source memory of in-group labels, as compared to out-group labels, and that such advantage would be modulated by participants' own in-group biases.

Matching trials

To test whether source memory was better for in-group than for out-group words, we ran a logistic mixed effects model with accuracy as the dependent measure and fixed effects for Group Membership (In-group vs Out-group, reference level: In-group), In-group Bias (centered continuous predictor), and their interaction. Block (Block1 vs Block2, reference level: Block1) was included as covariate to control for potential confounds[3]. We added per-participant and per-items random intercepts and a by-participant slope for Group Membership.

Overall, participants' accuracy in the matching trials was 63.08% (SD=48.28) and above chance level, as confirmed by a one-sample t-test (i.e., 50%) (t=10.41, p<.001). Results showed that neither Group Membership (beta=0.10, SE=0.13, z=0.75, p=0.45) nor its interaction with In-group Bias significantly predicted accuracy (beta= 3.13, SE=3.23, z=0.97, p=0.33). Participants' accuracy did not differ between Block1 and Block2 (beta=0.02, SE=0.11, z=0.19, p=0.34). However, participants' In-group Bias significantly predicted accuracy, but only at the reference level, i.e., in-group membership (beta=-6.90, SE=3.17, z=-2.18, p<.05). By re-leveling Group Membership with Out-group as the reference level, we saw that accuracy for out-group speaker-label pairs was not modulated by the individual measure of In-group Bias (beta=-3.76, SE=2.93, z=-1.29, p=.20) (see Figure 2.3). This means that the more in-group

[3]To ensure that the patterns of results were comparable across both testing blocks, we also ran a mixed-effect model where response accuracy was modelled by Group Membership (In-group vs Out-group, reference level: In-group), In-group Bias (centered continuous predictor), Block (Block1 vs Block2, reference level: Block1) and their interactions. We added per-participant and per-items random intercepts and a by-participant slope for Group Membership. Results from this analysis showed that neither the main effect of Block (p=.37) nor its interactions with the other variables (ps>.16) significantly predicted response accuracy.

biased participants were, the less accurate they were at recognizing speaker-label pairs, in particular when the speaker-label pairs were of their in-group.

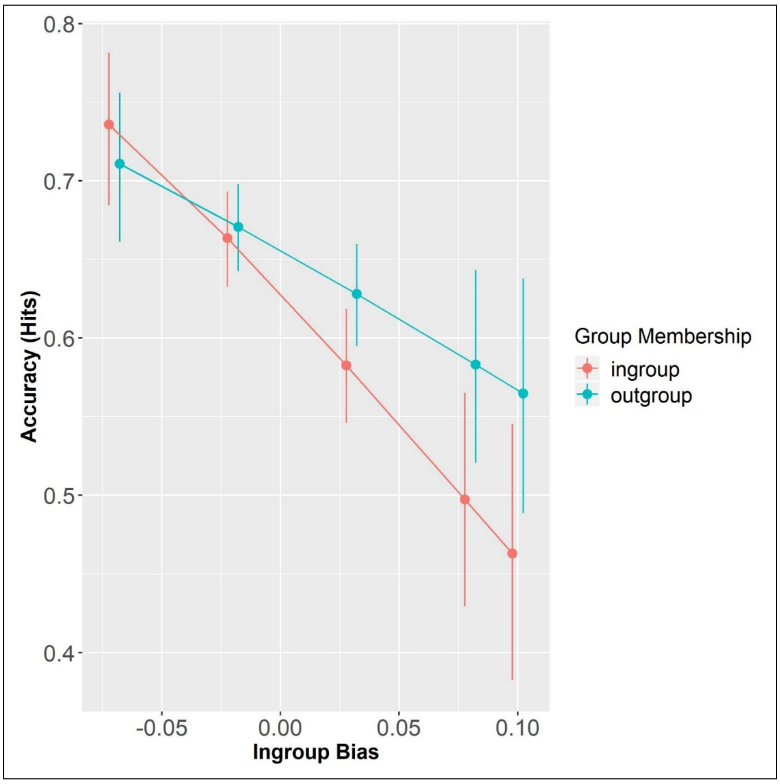

Figure 2.3: Accuracy on matching trials in the Main Experiment as a function of Group Membership and In-group Bias (centered). Error bars represents standard errors.

Mismatching trials

To test whether speaker group membership influenced the level of detail for speaker-specific information encoded with the new words, we analyzed accuracy on mismatching trials. By looking at participants' performance on within-affiliation mismatching trials, where labels were paired with incorrect speakers but belonging to the same affiliation as the correct source, we were able to test whether the source-related information for novel words was speaker-specific (participants should have rejected the wrong source) or group-specific (participants would have incorrectly accepted the wrong source). We hypothesized that people would encode more speaker-specific information with in-group labels than with out-group labels. We therefore predicted greater confusion among out-group speakers than among in-group speakers in the within-affiliation mismatching trials. We also predicted that this difference in accuracy would depend on individual In-group Bias, such that the greater In-group Bias participants exhibited, the greater difference they should show between in-group vs out-group trials. Conversely, in between-affiliation mismatches (i.e., where an in-group label was shown with out-group members, and vice versa) no differences were expected.

To test these hypotheses, we ran a logistic mixed model analysis with fixed effects for Mismatch Type (Within- vs Between-affiliation, reference level: Within), Group Membership (In-group vs Out-group, reference level: In-group), In-group Bias (centered continuous measure), and their interaction terms. We added Block as covariate, per-participant and per-item random intercepts and by-participant slopes for Group Membership and Mismatch Type.

Overall, participants' accuracy on mismatching trials was 65.79% (SD=47.45) and above chance level (i.e., 50%), as confirmed by a one-sample t-test ($t=31.31$, $p<.001$). As expected, participants were more accurate for between–affiliation mismatches than for within–affiliation mismatches (beta=0.53, SE=0.14, z=3.10, $p<.0001$; mean=70.35 %, SD=45.68 and mean=61.22%, SD= 48.73, respectively). This shows that participants encoded speakers' affiliations. Due to a practice effect, they were also more accurate in Block2 than in Block1 (beta=0.79, SE=0.05, z=15.95, $p<.0001$; mean=73.61 %, SD=44.08 and mean=58.09 %, SD=49.35, respectively). Participants' performance was also significantly predicted by In-group Bias at the reference levels (beta=7.98, SE=3.03, z=2.64, $p<.01$) and by a marginally significant interaction of In-group Bias with Group Membership (beta=-3.52, SE=1.96, z=-1.80, p=0.07), which suggests that participants with different strengths of In-group Bias were differently affected by speaker Group Membership. Specifically, simple effect analyses revealed that the larger the In-group Bias, the better participants were at correctly rejecting pairings involving the in-group membership (beta=7.98, SE=3.03, z=2.64, p<.01). On the other hand, participants' In-group Bias did not predict their performance with pairings involving the out-group membership (beta=4.46, SE=2.78, z=1.6, p=.11) (see Figure 2.4).

Furthermore, neither the two-way interaction between In-group bias and Mismatch Type (beta=-5.21, SE=3.48, z=-1.5, p=.13), nor the three-way interaction between Mismatch Type, Group Membership, and In-group Bias reached significance (beta=3.10, SE=2.54, z=1.22, p=.22). Therefore, participants' performance in both between- and within-affiliation mismatches was comparably affected by the Group Membership x In-group bias interaction.

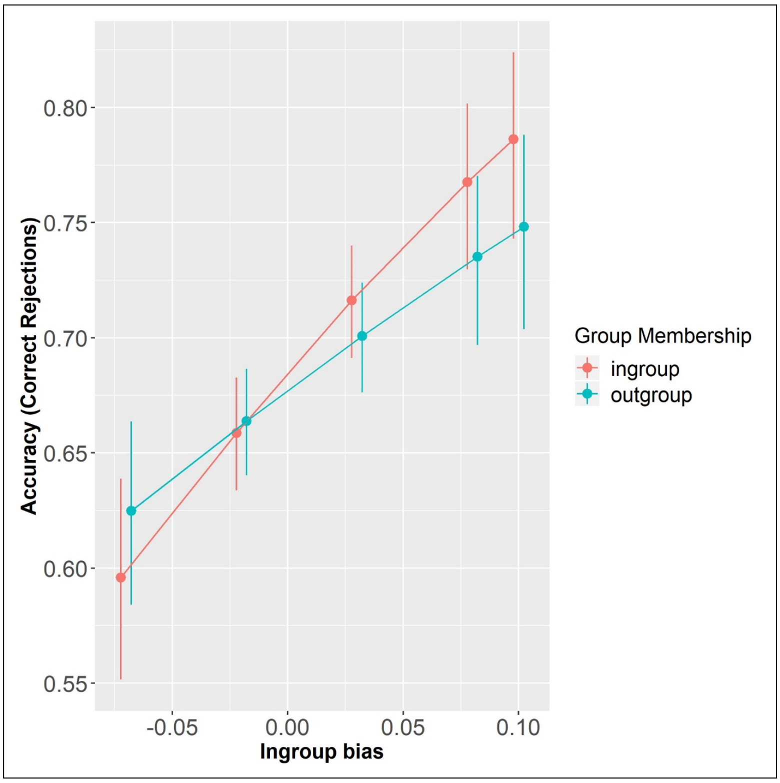

Figure 2.4: Accuracy on mismatching trials in the Main Experiment as a function of Group Membership and In-group Bias (centered). Error bars represent standard errors.

In short, results from the matching trials revealed a negative relationship between In-group Bias and response accuracy, especially for in-group pairings. This pattern suggests that participants with stronger in-group bias were more likely to produce misses with in-group speaker-label pairs. On the other hand, results from the mismatching trials revealed a *positive* relationship between In-group Bias and accuracy, meaning that those strongly biased participants also produced fewer false alarms when in-group pairings were involved. These seemingly contradictory results can be reconciled by stepping away from simple accuracy analyses and by relying on signal detection theory measurements which capture detection sensitivity (namely, d-prime) and response bias (namely, C).

D-prime and C values

Analyses over d-prime and C measures allow us to test whether participants' sensitivity and response bias during decision making processes differed for in-group vs out-group related decisions. We calculated two d-prime values and two C values per participant for in-group and out-group trials separately. In order to generate values that reflected participants' decisions to purely in-group or out-group trials, d-prime and C values were calculated from participants' performance in matching trials (i.e., hit rates) and within-affiliation mismatching trials (i.e., false-alarm rates)[4]. Between-affiliation mismatches were not considered for these analyses because they were created by having an element (either label or speaker) from each group and were therefore not purely in-group or out-group related. We ran two linear mixed-effect models with either d-prime or C values as the dependent variable and Group Membership (In-group vs Out-group, reference level: In-group), In-group Bias and

[4] To calculate C and d-prime values, we firstly followed Macmillan and Creelman (1991) and converted 0 values in False Alarms to 1/2N and 1 values in Hit rates to 1-1/2N. Next, we subtracted the zscored False Alarms rate from the zscored Hit rate. C values, were calculated using the following formula: $(-0.5) * (zscored(HitRate) + zscored(FalseAlarmRate))$.

their interaction as fixed effects. The models included per-participant random intercepts.

The model that explored the relationship between individual d-prime and the independent variables showed no significant main effects or interactions (ps>.57), suggesting that participants' sensitivity was not modulated by speaker Group Membership or their own In-group Bias, nor the interaction between them (see Figure 2.5, left panel).

On the other hand, the model exploring C values showed a significant main effect of In-group Bias (beta=8.60, SE=2.55, t=3.38, p<.001) so that the more in-group biased, the more conservative participants were in their decision (i.e., having a bias for "no" responses). Importantly, there was a significant interaction between In-group bias and Group Membership (beta=-4.26, SE=1.97, t=-2.16, p<.05), showing that participants with different In-group Bias strength were differently affected by speaker Group Membership. Simple effect analyses revealed that while In-group Bias strongly modulated participants' response bias with in-group labels (beta=8.60, SE=2.55, t=3.38, p<.001), this was only marginally so with out-group labels (beta=4.34, SE=2.55, t=1.71, p=.09). These findings show that participants differed in their response bias as a function of Group Membership and In-group Bias, so the more in-group biased they were, the more conservative they were in their in-group related decisions, as compared to out-group related decision (see Figure 2.6). In other words, they were more careful in attributing in-group words to any in-group speaker.

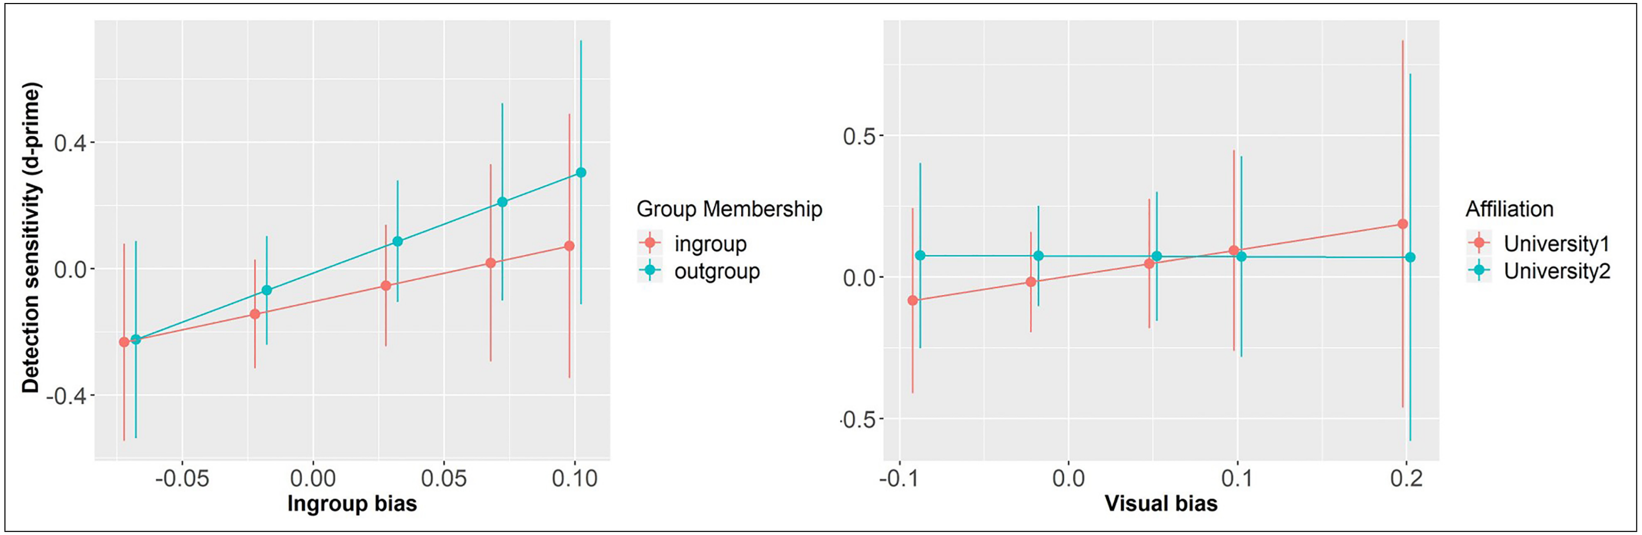

Figure 2.5: Detection sensitivity in the Main Experiment as a function of Group Membership and In-group Bias (centered), left panel, and as a function of Affiliation and Visual Bias (centered), right panel. Error bars represents standard errors.

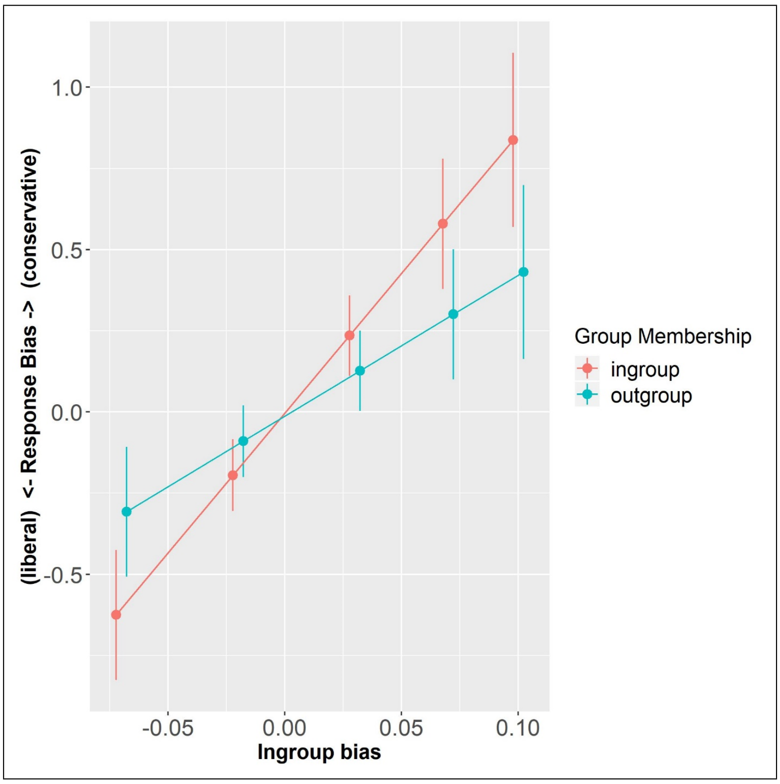

Figure 2.6: Response bias as a function of Group Membership and In-group Bias (centered). Error bars represents standard errors.

Control Experiment

We hypothesized that the tendency to monitor speaker social identity was dependent on whether the affiliations were perceived as socially salient, or relevant. To test this, we ran a control experiment in which participants learned new words from Dutch native students attending two Italian universities, as part of an exchange program. In this experiment, group membership was not manipulated. Participants still learned from two groups of speakers, like in the Main Experiment, but here the speakers' affiliations were supposed to be socially neutral because the speakers belonged to two foreign universities. Therefore, no differences were expected between the two groups. To control for potential visual dissimilarities between the logos used, participants performed the same perceptual matching task as in the Main experiment, responding to pairings involving the logos of the Italian universities. Similar to what we did in the Main experiment, we calculated an individual measure that in this case can be seen as an index of visual bias. This individual measure was entered in the statistical analyses.

Matching trials

We ran a logistic mixed effects model with accuracy as the dependent measure and fixed effects for Affiliation (University1 vs University2, reference level: University1), Visual Bias (centered), and their interaction. Block was included as covariate to control for potential confounds. We added per-participant and per-items random intercepts and by-participant slope for Affiliation.

Overall, participants' accuracy in the matching trials was 57.52% (SD=49.48) and above chance level, as confirmed by a one-sample t-test (i.e., 50% t=5.84, $p<0.0001$). Neither Affiliation, nor Visual Bias or their interaction significantly predicted accuracy (ps>.27). Participants' accuracy was better in Block2 than in Block1 (beta=0.28, SE=0.11, z=2.51, $p<.05$).

Mismatching trials

We ran a logistic mixed model analysis with fixed effects for Mismatch Type (Within- vs Between-affiliation, reference level: Within), Affiliation (University1 vs University2, reference level: University1), Visual Bias (centered continuous measure), and their interaction terms. We added Block as covariate, per-participant and per-item random intercepts and by-participant slopes for Affiliation and Mismatch Type.

Overall, participants' accuracy on mismatching trials was 69.37% (SD=46.10) and above chance level (i.e., 50%), as confirmed by a one-sample t-test (t=39.53, $p<.0001$). Generally, participants were more accurate in the between–affiliation mismatches than in the within–affiliation mismatches (beta=

0.20, SE=0.07, z=2.64, $p<.01$; mean=70.55% SD=45.59 and mean=68.18% SD= 46.58, respectively), indicating that even the irrelevant social affiliations were encoded to some degree. Participants were also more accurate in Block2 than in Block1 (beta=0.92, SE=0.05, z=18.00, $p<.0001$; mean=78.14% SD= 41.33 and mean=60.73, SD=48.84, respectively). None of the other main effects or interactions resulted significant (ps>.16), showing that, unlike the modulating effect of in-group bias in the Main Experiment, participants' memory for speaker-label pairings was not modulated by visual bias.

D-prime and C values

To be consistent, we also performed analyses over d-prime and C values, as we did in the Main experiment. Crucially, we did not expect any differences between the two academic affiliations. We calculated two d-prime values and two C values per participant for the two affiliations separately. We ran two linear mixed-effect models with either d-prime or C values as the dependent variable and Affiliation University1 vs University2, reference level: University1), Visual Bias (centered continuous measure), and their interaction terms. The models included per-participant random intercepts.

The model that explored the relationship between individual d-prime and the independent variables showed no significant main effects or interactions (ps>.67), suggesting that participants' sensitivity was not modulated by speaker Affiliation or their Visual Bias, nor the interaction between them.

Similarly, the model exploring C values showed no significant main effect of Visual Bias or interaction (ps>.24). There was a marginal effect of Affiliation (beta=0.11, SE=0.06, t=1.90, p=.06) with decisions made about University2 being numerically more conservative than decisions involving University1 (see Figure 2.7).

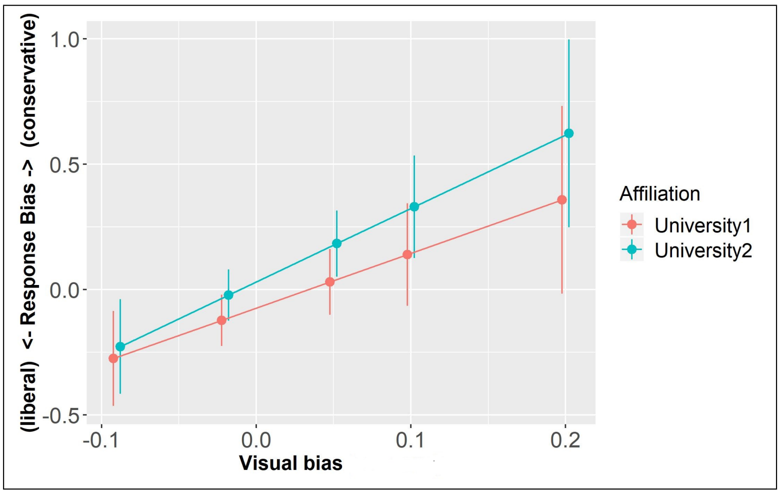

Figure 2.7: Response bias in the Control Experiment as a function of Affiliation and Visual Bias (centered). Error bars represents standard errors.

Discussion

We used a novel word learning paradigm to test whether learners of new words monitored speakers' social identity, such as their group and individual identity. Furthermore, we asked whether group membership status of the speakers and individual in-group biases of the learners affected the level of detail of speaker-specific information encoded in the novel lexical representations. We additionally performed a control experiment and ensured that the patterns found in the Main experiment were indeed a reflection of the social salience ascribed to speakers' group membership and not simply a consequence of the contrasting nature of our manipulation (i.e., teaching competing labels spoken by different groups of speakers).

In the test phase of the word learning task, participants' source memory for the new words was tested in an alternative forced-choice task (i.e., yes/no) where they decided whether displayed speaker-label pairs matched or mismatched what they learned in the exposure phase. This task offered a proxy for investigating the level of detail of speaker-specific information in the novel representations. Results confirmed our prediction regarding the general tendency to encode in parallel both linguistic content and speakers' social identity (i.e., speakers' affiliation). This tendency was reflected in the fact that participants made more within-affiliation errors than between-affiliation errors, i.e., source memory confusion. This finding provides further support for models of word learning where linguistic units are encoded together with speaker-related information (exemplar models e.g., Goldinger, 2007; Hay, Nolan, & Drager, 2006; Nielsen, 2011, see Drager & Kirtley, 2016, for a review).

Concerning our hypotheses about the effects of Group Membership and In-group Bias, the results revealed a more complex pattern than we had predicted. We had predicted that participants would encode in-group labels with a higher level of detail of speaker-specific information, as compared to out-group labels. This phenomenon was expected to be reflected in a) a higher proportion of hit rates for matching in-group speaker-label pairs and b) a higher proportion of correct rejections for within-affiliation in-group speaker-label pairs. Both effects were predicted to be positively modulated by the individual In-group Bias, so that the stronger the bias, the stronger the effects. We found that indeed participants with stronger in-group bias were better at correctly rejecting wrong in-group pairings (i.e., in the mismatching trials). However, when looking at the matches, the results revealed that those participants with stronger in-group bias were also more likely to miss matching in-group speaker-label pairs.

These seemingly contradictory results are hard to reconcile when relying only on accuracy (i.e., correct/incorrect). For this reason, we relied on signal detection theory measurements, such as d-prime and C values, to gain a deeper understanding of the phenomenon. These measures capture both hit rates and false-alarm rates for conceptually similar items and allow us to test whether participants' detection ability and/or response bias differed for in-group vs out-group speaker-label pairs. Results showed that participants' detection sensitivity was not modulated by our social manipulations such that they were equally sensitive to in-group and out-group speaker-label pairings. On the other hand, the model exploring C values showed that the more in-group biased, the more conservative participants were in their decision (i.e., having a bias for "no" responses), and this was particularly applied to in-group related decisions. That is, participants' in-group bias and speakers' group membership influenced *how liberally* decisions were made, so that participants with stronger in-group bias were more careful in attributing in-group labels to any speaker. This pattern explains why participants' in-group bias negatively predicted hit rates and positively predicted correct rejection rates: the stronger the in-group bias, the more likely participants responded "no" to in-group speaker-label pairs.

How do our findings reconcile with the initial predictions and with previous literature? While previous studies showed that source memory was more accurate for information related to in-group membership, compared to information related to out-group membership (e.g., Hugenberg et al., 2010; Greenstein et al., 2016), in the current study we showed that the scenario can be more complex. Participants with a stronger bias were more accurate at correctly rejecting mismatches involving in-group labels, but they were also more likely to miss in-group matches. Looking closely at these patterns, we could deduce, and confirm with our analyses, that it was participants' response bias that was mainly affected by our social manipulation of group membership, and by participants' in-group bias. Participants with stronger in-group bias were in fact more cautious when attributing in-group labels to any speakers.

Our results resemble previous findings by Castano, Yzerbyt, Bourguignon, and Seron (2002), who investigated if high vs low in-group identifiers differed in their decision preferences when they had to categorize ambiguous faces as either in-group (i.e., Northern Italians) or out-group (i.e., Southern Italians) members. They found that participants that strongly identified with their in-group membership were less likely to classify a target face as in-group member, as compared to participants with a lower in-group identification score (see Blascovich, Wyer, Swart, & Kibler, 1997; Yzerbyt, Leyens, & Bellour, 1995, for similar results). The authors claimed that such a pattern was supportive of the *In-group overexclusion hypothesis* (Leyens & Yzerbyt, 1992), which states that when people are in doubt about classifying targets as either in-group or out-group, they tend to exclude them from their in-group. Such a hypothesis seems to apply to our dataset as well where participants with stronger in-group bias were more conservative when attributing in-group labels to speakers.

We consider why it is that learners' in-group bias and speakers' group membership status might lead to differences in response preferences, but not in detection sensitivity, as we had predicted. In other words, what might it mean that an individual with strong in-group bias is selectively more conservative when making a decision that involves her in-group membership? Originally, we had predicted group membership and in-group biases to play a role during the *encoding* of novel words, leading to in-group representations with more highly detailed speaker-specific information, as compared to out-group representations. The lack of modulation on the detection sensitivity measure by these social variables suggests that in-group and out-group labels did not differ in how they were encoded. Instead, we found a significant Group Membership x In-group bias effect on response bias, so that the stronger the in-group bias, the more conservative participants' responses were in relation to in-group labels, but not in relation to out-group labels.

We believe that these differences in decision bias might reflect asymmetries during retrieval processes for in-group related episodic events, as compared to out-group related events. Previous research has shown that response bias acts during memory retrieval processes (Windmann, Urbach, & Kutas, 2002) and depends on criterion setting functions of the prefrontal cortex (Schacter, Norman, & Koutstaal, 1998; Swick & Knight, 1999; Miller, Handy, Cutler, Inati, & Wolford, 2001). During recognition decision-making processes, this brain region is considered to be involved in initiating, monitoring and controlling item-retrieval from memory to maintain a description of the information being sought and actively inhibit memory traces that do not match this description (Buckner, 1996; Fletcher, Shallice, Frith, Frackowiak, & Dolan, 1998; Wagner, Desmond, Glover, & Gabrieli, 1998; Henson, Shallice, & Dolan, 1999; Tomita, Ohbayashi, Nakahara, Hasegawa, & Miyashita, 1999). Therefore, Windmann and colleagues (2002) suggest that differences in response bias, especially when independent of the accuracy of the memory, can be explained by the fact that decision makers differ in what they prioritize in the task (i.e., the detection of matches or mismatches).

In light of this evidence, our findings might reflect differences in recognition threshold for in-group vs out-group memory traces. During the decision processes, the inhibitory system of those participants who were more in-group biased was activated to a larger extent to avoid creating false positives and attributing in-group information to any source. Attributing in-group labels to incorrect speakers might have been perceived as more hurtful than missing the detection of correct in-group speaker-label pairs, as the in-group overexclusion hypothesis states. If this was indeed the case, these findings would validate the claim that in-group membership information recruits the control system to a larger degree than out-group membership does, as has been previously suggested (Meissner et al., 2005; Van Bavel & Cunningham, 2012). Furthermore, such a response bias could contribute to the effect known as out-group homogeneity in face recognition and categorization tasks (Castano et al., 2002), where new out-group faces produce more false alarms than new in-group faces do, supporting the claim that out-group members are perceived as more homogeneous.

Of course, it is important to replicate the present novel findings using different groups of speakers and different tasks, as to ensure that these effects and biases do not reflect poor recognition and/or high cognitive load in general. While the analyses revealed that participants' accuracy was above chance level, it was still relatively low. Note that participants learned about the affiliations of the speakers during the word learning task, by seeing the faces of the speakers together with the logos of the supposed academic affiliations. This means that during the source memory test, they were potentially retrieving from their memory multiple pieces of information (e.g., speaker's affiliation, label's source). On that point, it is worth mentioning that even though in-group trials included a logo that might be more familiar than the out-group logo, as it is participants' own university logo, participants did not exhibit superior memory for in-group items. Future studies should test whether our finding replicates when the source memory task is simplified, for instance, by participants learning the group membership status of speakers in an earlier experimental session, and in a more natural way (e.g., by listening to speakers referring to their university lives).

Similarly, to gain a deeper understanding of how speakers' group membership and individual in-group biases influence language learning, it would be important to test whether source memory (i.e., the speaker) and item memory (i.e., the word) are equally affected by these social factors. While in this study we investigated the encoding of context-related information in the representations of novel words, and tested if its specificity was modulated by group membership and individual in-group biases, further research should test whether these factors influence the linguistic component of the representations, too.

According to our general hypotheses, labels learned from in-group speakers would be easier to remember than words learned from out-group speakers.

If these patterns are substantiated, they will have far-reaching implications for theories of language learning and processing, as well as theories concerning prejudice and stereotyping. For instance, the results suggest that interlocutors' group membership status and listeners' individual biases may influence how likely newly acquired information is to be generalized to other interlocutors. In particular, for in-group speakers, listeners with a strong in-group bias appear to be more cautious when attributing in-group related information to other speakers, preventing over-generalization, whereas speakers with low in-group bias may be more liberal in their generalizations. One may wonder whether this greater caution relates to social stereotypes as well. It is well known that people tend to homogenize out-group members whereas they are aware of the heterogeneity of their own in-group. It would be interesting to examine to what degree such findings relate to the findings from this study about individuals' greater cautiousness in attributing information to in-group compared with out-group members.

Further research should explore more how social characteristics that are ascribed to both speakers and contexts during language processing, and information processing more generally, influence encoding and storage, and how these, in turn, affect decision processes during memory retrieval. Such research would shed further light on the intersection between memory and processing, including language processing, and, importantly, how this intersection is influenced by the social properties of the input.

3 | How in-group biases influence the level of detail of speaker-specific information encoded in novel lexical representations[1]

Abstract

An important issue in theories of word learning is how abstract or context-specific representations of novel words are. One aspect of this broad issue is how well learners maintain information about the source of novel words. We investigated whether listeners' source memory was better for words learned from members of their in-group (students of their own university) than for words learned from members of an out-group (students from another institution). In the first session, participants saw six faces and learned which of the depicted students attended their own or a different university. In the second session, they learned competing labels (e.g. *citrus-peller* and *citrus-schiller*, in English *lemon peeler* and *lemon stripper*) for novel gadgets, produced by the in-group and out-group speakers. Participants were then tested for source memory of these labels, and for the strength of their in-group bias, i.e. for how much they preferentially process in-group over out-group related information. Analyses of the accuracy of source memory demonstrated an interaction of speaker group membership status and participants' in-group bias: the stronger the bias, the less accurate was participants' source memory of out-group labels as compared to in-group labels. These results add to the growing body of evidence on the importance of social variables for adult word learning.

[1]Adapted from *Iacozza, Meyer, & Lev-Ari (2019). How in-group biases influence the level of detail of speaker-specific information encoded in novel lexical representations. Journal of experimental psychology. Learning, memory, and cognition.*

Introduction

Previous research has established that people can quickly make use of any available cue (e.g., Babel, 2012). Such group membership status has been shown to influence how attentively people process information about and from others, resulting in more accurate memory for information related to in-group compared to out-group members (e.g., Hugenberg et al., 2010; Judd & Park, 1988; Van Bavel & Cunningham, 2012). for instance, after seeing a series of unfamiliar faces, people showed superior memory for those faces that they believed to be in-group members, as compared to out-group faces (e.g., Van Bavel & Cunningham, 2012). Similarly, people were better at recalling in-group biographic information than out-group related information (e.g., Judd & Park, 1988).

Here, we test our proposal that in-group biases permeate language learning as well. Specifically, we ask whether the group membership of those from whom new words are learned modulates the level of detail with which information about speaker identity is encoded. We suggest that representations of novel words are more speaker-specific when learned from in-group members than when learned from out-group members. Furthermore, we predict that the difference in level of detail between in-group and out-group representations are modulated by the strength of learners' own in-group biases, namely, by how much people preferentially process in-group over out-group related information.

To test our predictions, we designed a word learning study in which participants learned novel words from in-group and out-group speakers. Successively, they performed a task that measured their in-group bias. We then tested their source memory for the words learned from in-group vs out-group members and assessed how source memory was related to individual in-group biases.

Before turning to the current study, we review relevant literature. Firstly, we focus on evidence that shows that speaker social identity plays a role in language processing. Secondly, we describe existing exemplar-based models of language processing that account for the effects of speaker social identity on language processing. At this stage, we also propose modifications for existing models. We argue that existing models should not assume that all speakers are treated equally but instead assume that information provided by in-group speakers is encoded more fully than information from out-group speakers. We end by providing support for our propose modifications from non-linguistic studies in social psychology that report group membership effects on memory and information processing.

Speaker social identity can influence language processing

Previous research indicates that people rapidly extract information about others' social identity (see Babel, 2012, for a review). Such information has been shown to activate expectations and attitudes (e.g., Greenwald & Banaji, 1995; Jones & Fazio, 2010). For instance, children who were given silent demonstrations of unfamiliar objects' different functions selectively endorsed the functions provided by native-accented rather than foreign-accented speakers (Kinzler et al., 2011), showing a native-speaker preference even when the language was not involved in the task. Using others' social identity to guide processing extends to language processing in a variety of ways. For instance, a speaker's social identity can shape listeners' expectations. Upon hearing a speaker's voice, people activate expectations about what will likely be said. If these expectations are not met, such as when a habit of drinking wine is reported in a child's voice, language processing becomes harder (van Berkum et al., 2008; Martin et al., 2016; Walker & Hay, 2011). Similarly, speaker social identity can affect listeners' word processing (e.g., Hay, Nolan, & Drager, 2006;

Hay, Warren, & Drager, 2006; K. Johnson et al., 1999; Niedzielski, 1999) . For example, the gender of a face displayed when listening to ambiguous vowels affected participants' responses in a vowel identification task, suggesting that information about the speaker affects speech perception (K. Johnson et al., 1999). Finally, listeners detected fewer word changes in a story after listening to non-native as compared to native speakers (Lev-Ari & Keysar, 2012). Importantly, these differences disappeared when participants were instructed to pay attention to the exact wording of the stories, rather than simply comprehending them (Lev-Ari & Keysar, 2012). These results confirmed that people were in principle able to process the language of native and non-native speakers but did so less when listening for comprehension to the non-natives (unless given explicit instructions).

This body of evidence shows that people encode speakers' social identity and this information can influence speech processing. Based on this evidence, one might expect speaker's social identity to influence language learning as well. When learning new words, the novel linguistic information might be encoded and stored in representations that contain details about both its linguistic content (e.g., its word form, its meaning) and its source (i.e., who said it). Indeed, existing exemplar-based models of speech processing assume such encoding of social information.

Exemplar-based models: concurrent encoding of social and linguistic input

Exemplar-based models of speech processing posit that linguistic experiences are encoded as rich episodic memories (i.e., exemplars) (e.g., Goldinger, 2007; Hay, Nolan, & Drager, 2006; Nielsen, 2011). Exemplars contain information that is both language-specific (e.g., includes phonetic, lexical, and syntactic details) and related to the speaker's social identity, spanning from more intrinsic indexical characteristics such as age and sex to broader social categories, such as professions (e.g., a doctor) and membership in a specific community (e.g., health care professionals) (see Drager & Kirtley, 2016, for a review). Associating speaker-related information to linguistic input can be beneficial for language users. It allows listeners to pre-activate those associations in the appropriate context and with the relevant interlocutors. This may exploit the socially conditioned variation in speech to speed up and/or ease comprehension (Münster & Knoeferle, 2018, for a recent account of socially situated language processing).

In a recent extension of exemplar models, Sumner et al. (2014) have proposed the *socially weighted encoding of spoken words*. They propose that social characteristics, such as perceived prestige, might affect linguistic encoding, such that prestigious forms would receive greater weight. In support of a socially weighted encoding account, Sumner and colleagues referred to empirical evidence from work that contrasts standard to non-standard phonetic variations. Extending their theory, one could hypothesize that people may encode linguistic variations more strongly if they are associated with speakers to whom a special status is ascribed.

Here, we propose that learning new words from speakers that are ascribed a special status, such as in-group members, might lead to lexical representations that are richer in contextual information (e.g., speaker-related) compared to representations of words learned from speakers without a special status. In the current study, we tested the hypothesis that representations of novel words learned by in-group speakers contain individual-specific detail, whereas new words learned from out-group speakers contain broader group-specific information, and that such a pattern is further modulated by learners' in-group bias. This finding would contribute to the open debate in psycholinguistics on how words are represented in the mental lexicon and call for an extension of existing exemplar-based theories that assume an encoding of contextual information that does not vary in relation to the identity of the speakers. Though the role of group membership in language processing has not yet been explored, there is evidence to suggest that it influences processing and learning.

In-group biases

Well established self-biases, such as improved memory and faster processing of stimuli related to the self vs stimuli related to others (see Symons & Johnson, 1997, for a meta-analysis) have been shown to extend to in-group members as well (e.g., Aron, Aron, Tudor, & Nelson, 1991; Sui & Humphreys, 2015; Symons & Johnson, 1997; Cadinu & Rothbart, 1996). Thus, memory is better for in-group faces than out-group faces (e.g., Van Bavel et al., 2008; Hugenberg et al., 2010) and for information delivered by in-group than by out-group members (e.g., Frable & Bem, 1985; Wilder, 1990). People also learn and process new associations between previously neutral stimuli (e.g., geometrical shapes) and in-group membership (e.g., the logo of their favorite football club) more quickly than associations involving out-group membership (Enock et al., 2018; Moradi et al., 2015). Like self-biases, in-group biases may rely on the recruitment of additional attentional resources (Meissner et al., 2005; Van Bavel & Cunningham, 2012). These additional resources have been suggested to lead to representations with a higher level of detail for in-group than out-group representations. For example, in-group faces are easier to recognize than out-group faces because more individual-specific information is encoded during the exposure phase (e.g., Hugenberg et al., 2010).

The better processing of in-group-related information extends to the individuals providing the information. Greenstein and colleagues (2016) showed that source memory was superior when participants believed the sources to be in-group members than when they believed them to be out-group members.

Such findings are highly relevant for language learning and language processing because they suggest that when listeners comprehend speech, the amount of both speaker-specific detail and linguistic information that is encoded might depend on whether the speaker is an in-group or an out-group member of the listener.

The current study

We propose that listeners encode the group membership status (i.e., in-group vs out-group) of the speakers from whom they learn new words, and that group membership status and individual in-group bias influence the level of detail with which speaker-specific information is encoded. To test this idea, we probed participants' source memory for novel words. We predicted that words learned from in-group speakers would contain a higher level of detail about who produced them, compared to words learned from out-group speakers. This would result in in-group linguistic representations being more speaker-specific than out-group representations. Crucially, this in-group advantage should be stronger for participants exhibiting greater in-group bias.

We conducted a two-day experiment. On day 1, speaker group membership was established via a speaker familiarization task in which participants listened to six fictitious speakers refer to facts and habits about their lives that implied either attendance of the participants' university (n=3; in-group speakers) or of a different university (n=3; out-group speakers). On day 2, participants performed the word learning task, which was a communication task. In the exposure phase of this task, they learned novel labels for 24 unfamiliar gadgets by listening to the speakers from day 1 refer to these gadgets. Crucially, target gadgets received one label from an in-group speaker and an alternative competing, but equally fitting, label from an out-group speaker (e.g., *citrus-peller* vs *citrus-schiller*, in English *lemon peeler* vs *lemon stripper*). Afterwards, source memory for these words was tested. Participants were shown a photo of one speaker and one label at a time and asked if the speaker had produced the label in the previous phase (eliciting a forced choice yes/no). This was done with all speakers and labels from the exposure phase. Lastly, participants' implicit in-group bias was measured by their performance in a perceptual matching task.

Following our hypothesis about different levels of detail depending on speaker group membership, we predicted that representations of words learned from in-group members would include more detailed information about the speaker than representations of words learned from out-group members. This would be evidenced by greater accuracy for in-group than for out-group label-speaker pairs. Additionally, we expected participants to spontaneously encode the speakers' group membership status (even though they were not explicitly instructed to apply this strategy). Consequently, when asked to recognize the source of the new words, we expected participants to struggle remembering the exact source for each word, and instead to rely on speaker group membership to inform their decisions. Therefore, they should be more likely to misattribute words to incorrect speakers within the same university than between different universities, i.e., there should be source memory confusion. Since in-group representations should be more speaker-specific, the source memory confusion for in-group words should be of a smaller magnitude as compared to confusion for out-group words. Lastly, we predicted that these effects would be modulated by participants' own in-group bias, such that the stronger the in-group bias, the larger the difference between speaker-specific details in in-group vs out-group representations.

Methods

Participants

Sixty native Dutch speakers (age range: 18-27years, mean=21.98, SD=1.89) participated in the study after providing their informed consent, as approved by the Ethics committee of the Social Sciences department of Radboud University (project code; ECSW2014-1003-196). All participants were students or recent graduates of Radboud University Nijmegen[2]. Eight additional participants were tested but their data were excluded from the study because they reported during the debriefing, at the end of the experiment, to have strong connections with the out-group university (e.g., their sister or best friend studying there). All participants were female, as were the speakers from whom they learned. This was done to avoid having an additional social dimension (gender) of in-group status interact with the one we manipulated (i.e., same/different university).

Materials

Materials for the speaker familiarization task

Speakers: Six fictitious speakers were created by pairing female faces selected from the Chicago Face Database (Ma et al., 2015) with the voices of native Dutch female speakers recorded in our laboratory. Prior to the experiment, voices were matched for perceived typicality and attractiveness (paired t-tests, ps>.05) via an on-line norming in which twenty different participants participated. Each speaker was a unique combination of one face and one voice, consistent across participants.

[2]One participant's accuracy was more than 3SD away from the population mean in the perceptual matching task. We excluded her data from the word-learning analyses.

Facts: After a norming procedure (see Appendix 1, for details), we selected twenty-four sentences about student life and recognizable landmarks or actions. Twelve sentences referred to Radboud University and twelve to Groningen University. This was achieved by creating twelve unique sentential frames in which key words explicitly related to either Radboud or Groningen University were embedded. All statements were audio-recorded by each of the female voices described above.

Materials for the word learning task

Speakers: The same speakers from the speaker familiarization task were used.

Gadgets and labels: Twenty-four images of unfamiliar gadgets (e.g., lemon peeler) and their corresponding labels were selected via a norming study (see Appendix 1, for details). Half of the gadgets, hereinafter referred to as target gadgets, were presented with two competing labels, which were equated for goodness-of-fit and frequency. The other twelve gadgets were presented with a single label and served as fillers. All labels were audio-recorded by the speakers.

Materials for the individual in-group bias task

University logos: Black and white pictures of the logos of Radboud and Groningen universities were used.

Geometric shapes: Black shapes for triangle, square and circle were used.

Materials for the gadget familiarization task and prior familiarity task

The same 24 images of gadgets used in the word learning task were employed in these tasks. However, for the purpose of the gadget familiarization task, an additional exemplar for each gadget was selected online.

Procedure

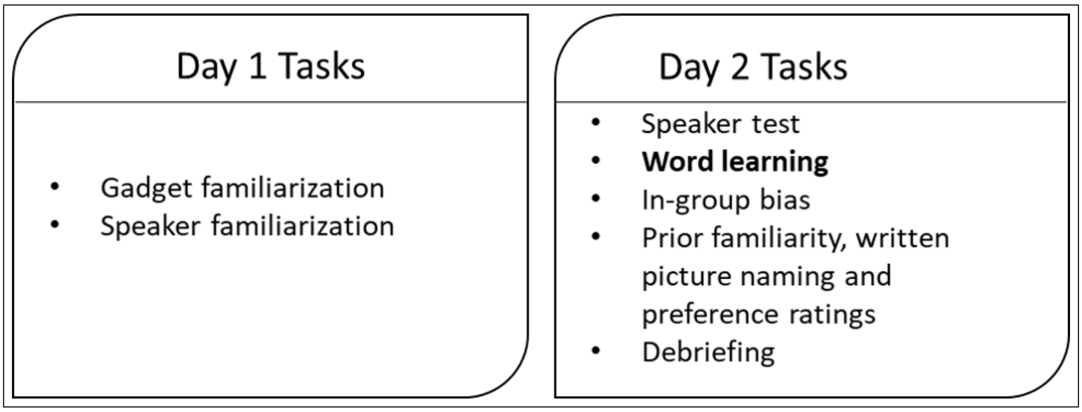

Figure 3.1: Overview of the procedure

Day 1

Gadget familiarization task

The goal of the gadget familiarization task was to familiarize participants with the gadgets to reduce demands in the subsequent word learning task. This task was a picture-matching task. On each trial, participants saw two images and had to indicate whether they depicted the same gadget. In half of the trials, different items were shown, whereas in the remaining trials, two instances of the same gadget (only one of which was later used in the word learning task) were displayed. The gadgets were shown in two exposure rounds with a randomized trial order. This led to 96 trials in total (48 in each round). Participants' responses were self-paced. Performance on the task was at ceiling, as expected, and will not be discussed further.

Speaker familiarization task

The goal of the speaker familiarization task was to familiarize the participants with the speakers and ensure that they knew their group membership status before they did the word learning task on day 2. Participants were told this task was part of a joint project of Radboud University, their in-group university, and Groningen University, an out-group university, in which they would see students of those universities describing their student life. The information provided was supposed to guide inferences about the speakers' academic affiliation: all habits implied attendance at either the in-group or out-group university. In order to make sure these affiliations had been correctly inferred, participants performed a surprise memory test. In each trial of this memory test, they saw a photo of a speaker and six alternative sentences written down. They selected which sentence among the given six was the one that the speaker had previously uttered. The memory test was performed both on Day 1, immediately after the exposure phase, and on Day 2, at the beginning of the experimental session. Performance on these indirect tests allowed us to ensure that participants indeed learned speakers' academic affiliations.

On each trial, participants saw a fixation cross (500 ms) followed by a photo of a speaker (500 ms). Then the sentence was played while the photo was still displayed. Two exposure rounds of 24 trials each were administered with a randomized trial order. In each round, participants heard all six speakers, three per university, produce four statements each. Speaker affiliation was randomized across participants.

After exposure, a surprise source memory test was administered. On each test trial, participants saw a photo of a speaker and six written sentences. These sentences were: the correct response (i.e., the one sentence that the speaker had previously uttered), two incorrect sentences produced by other speakers from the speaker's university, and three incorrect sentences produced by speakers from the other university. Participants were instructed to indicate which sentence had been previously uttered by the speaker by pressing the corresponding key on a keyboard. Responses were self-paced and immediate visual feedback about the accuracy of their choice was provided along with the correct option.

Responses were coded as either a) correct, b) within-university error (selection of a sentence produced by a speaker from the speaker's university), or c) between-university error (selection of a sentence produced by a speaker from the other university). If participants made a between-university error, the trial was repeated at the end. This was done until no between-university errors were made to ensure participants correctly learned the speakers' affiliation before moving to the word learning task. We did not repeat trials that led to within-university errors, since the goal of the training was only to teach participants which speakers were members of their in-group and which belonged to the out-group. This was critical for the assessment of our hypothesis concerning the effect of in-group bias on novel word learning.

Day 2

Speaker test

The test of the speaker familiarization task from day 1 was repeated (without repeating the exposure phase) to make sure participants still remembered the speakers' academic affiliations before the word learning task. Speakers and sentences were the same as on Day 1, and the procedure of repeating trials that yielded between-university errors that was employed on day 1 was also employed on day 2.

Word learning task

The word learning task was presented as a communication task. Thus, participants were told that they would select pictures according to information provided by the speakers from the previous task. In the exposure phase, participants saw 24 gadgets, all named by both in-group and out-group speakers. Half were target gadgets, for which in-group and out-group speakers provided competing, but equally fitting, labels, whereas the other half were fillers, for which unique labels were provided. Fillers were included to minimize participants' awareness of the nature of the manipulation (i.e., in-group vs out-group) in the experiment. Note that not all speakers referred to all the gadgets. In fact, each gadget was only labeled by two of the six speakers (one in-group and one out-group speaker). Speaker group membership, speaker-label pairing, and label-group membership status pairings were randomized per participant.

On each trial, a fixation cross (500 ms) preceded a photo of a speaker (800 ms). Then, the recording of the gadget label was played while the photo was still displayed. Simultaneously, the written form of the label appeared superimposed upon the speaker's mouth (1500ms). Next, three gadgets appeared on the screen and participants selected the one that fit the audio and the written label (see Figure 3.2 for an example of the learning display[3]). If the response was wrong, the audio was repeated. Three exposure rounds were administered so that each display was repeated three times, once per round, in a randomized trial order.

Figure 3.2: Example of the learning display in the word learning task. Participants had to select the gadget that was mentioned. In this case, they had to select the first image. Stimuli are not drawn to scale.

[3]Due to copyright issues, none of the pictures of the gadgets in the example correspond to actual stimuli, but they provide a good approximation of the type of stimuli we used.

Next, participants performed a surprise source memory recognition task on the target gadget labels only. In each trial, they saw a photo of a speaker and a written label (see Figure 3.3). Participants indicated whether the speaker had produced the label in the previous phase via key press (yes/no). Trials were self-paced. There were 148 test trials in total in which all possible speaker-label pairings were shown. This means that there were 24 matching trials (i.e., the speaker had indeed produced the label) and 120 mismatching trials (i.e., the label had not been used by the speaker). Of those mismatching trials, 48 were within-university mismatching trials (showing a label along with a wrong speaker from the same university as the correct one), 48 where between-university mismatching trials (showing a label with a wrong speaker from the other university). The remaining 24 trials showed a speaker with a label that competed with the one she used (e.g., the speaker that had used "citrus-schiller", instead of "citrus-peller") and were included only so that all possible combinations of speaker-label pairings were available; these were not analyzed. Note that in all mismatching trials, the correct answer was that the pairing was incorrect as the speaker depicted in the photo had not used the displayed label in the word learning task.

Figure 3.3: Example of a trial in the source memory task. Participants indicated if the speaker had produced the label in the exposure task. Stimuli are not drawn to scale.

In-group bias task

The participants' in-group bias was measured in a perceptual matching task (Moradi et al., 2015), which has been shown to provide results that are reliable within individuals and across different test sessions (Stolte et al., 2017). Three geometric shapes (circle, square, triangle) were randomly paired with logos of three universities: the in-group university –Radboud University, and two other universities -Groningen University and Tilburg University. Each association was initially presented ten times. Then, participant performed a practice block of 24 trials, followed by two blocks of 120 experimental trials each. In both practice and test trials, a fixation cross (500 ms) preceded a blank screen (between 1000-2000 ms) and the simultaneous presentation of logo and shape (600 ms), following the timings utilized in Moradi et al. (2015). Participants had 1500 ms to judge the accuracy of the pairing. Feedback was given only during practice.

In-group bias in this task is typically indexed by faster and more accurate responses for stimuli that are newly associated with in-group membership compared to stimuli associated with out-group membership (e.g., Moradi et al., 2015). We replicated these general patterns (see Appendix 2) and extracted individual measures of in-group bias by calculating a per-participant measure of effect size, Cohen's d, from both accuracy and RTs over in-group versus out-group matching trials. The measure calculated over RTs was not a significant predictor in any of the models we ran; thus, we will focus on the measure derived from accuracy (see Appendix 2, for the distribution of the individual measures).

Ratings of prior familiarity, written picture naming and preference

After the perceptual matching task described above, participants carried out three further tasks. First, in order to exclude gadgets that participants already had labels for before the experiment, we asked them to indicate whether they already knew any gadget and their names before the experiment. Based on these ratings, gadgets familiar to a participant were excluded from their data set (11.81% of data across all participants). Second, participants saw all the pictures of the gadgets again and were asked to write down the labels they had learned, as well as any additional labels they might know or come up with. Finally, they were asked to indicate which label among those they had just written down (regardless of whether or not they had seen them in the experiment) was the most appropriate for each gadget. This task was meant to assess the participants' preference for labels acquired from in-group versus out-group speakers. However, probably due to unclear instructions, many participants were very creative and came up with their own alternatives, or rarely wrote down two labels for one object. Therefore, this analysis could not be conducted for scarcity of data.

Results

All analyses were performed with mixed-effects modelling as implemented in the lme4 package (version 1.1-15; Bates et al., 2014) in R (R Core Team, 2016). The models' random structures were determined following Bates et al. (2015).

Speaker familiarization task

We tested whether participants correctly learned speaker group membership status. Overall, participants answered 78.65% of trials correctly (SD=41) (81.23% on day 1 and 76.05% on day 2). Their performance was above chance level (16.67% since there was one correct response out of six presented options), as confirmed by a one-sample t-test (t=82.47, p<.001). The overall accuracy rate was calculated across all trials, including those in which participants made within-university errors. This type of errors indicate that participants could not remember which exact statement corresponded to the speaker, but that they did correctly infer speakers' affiliation. When re-coding within-university errors as correct responses (i.e., correctly identifying the university the speaker was from), the average of correct responses increased to 96.51% (SD=18.35) for day 1 and to 97.1% (SD=16.79) on day 2. These analyses confirmed participants successfully learned which speakers belonged to which academic affiliation.

Word learning task

The main task in the current study was the word learning task. After an exposure phase, participants were tested with a memory recognition test. In this test, participants were presented with trials in which the speaker-label pairings either matched or mismatched with what they had seen in the learning phase. That is, the label had or had not been produced by the displayed speaker. We carried out two main analyses over the memory data. We looked first at the matching trials and, secondly, at the mismatching trials.

Matching trials

To test whether source memory was better for words learned from in-group than from out-group speakers, we ran a logistic mixed effects model with accuracy per trial as the dependent measure and fixed effects for Group Membership (In-group vs Out-group, reference level: In-group), In-group Bias (z-scored continuous measure), and their interaction. We added per-participant and per-items random intercepts and by-participant slope for Group Membership.

Overall, participants' accuracy in the matching trials was 68.11% (SD=46.62) and above chance level, as confirmed by a one-sample t-test (i.e., 50%, t=14.35, p<.0001). Neither Group Membership (out-group: mean=70.38, SD=45.69; in-group: mean=65.84, SD=47.46; beta=0.19, SE=0.13, z=1.45, p=.15) nor In-group Bias (beta=-0.02, SE=0.33, z=-0.07, p=.95) significantly predicted accuracy at the reference level. The Group Membership x In-group Bias interaction, however, was marginally significant (beta= -0.72, SE=0.39, z=-1.84, p=.07)[4]. Follow-up simple effects analyses confirmed that In-group Bias affected the accuracy of participant responses differently for the in-group and the out-group matching trials. While the simple effect of In-group Bias did not predict response accuracy for in-group matches (beta=-0.02, SE=0.33, z=-0.07, p=.95), it was a significant negative predictor of participants' performance for the out-group matches (beta=-0.75, SE=0.27, z=-2.73, p=.006). In agreement with our predictions, the larger participants' In-group Bias, the less accurate they were at identifying out-group speakers (see Figure 3.4).

[4] Visual inspection of the figure might imply that participants with low in-group bias (and potentially out-group bias) exhibit higher accuracy with out-group members. An analysis that focuses only on participants in the bottom quartile of in-group bias (all showing no bias or reversed bias on the In-group Bias measure) does not show any significant effect though, and an analysis that excludes the bottom quartile of participants is qualitatively similar to the reported analysis.

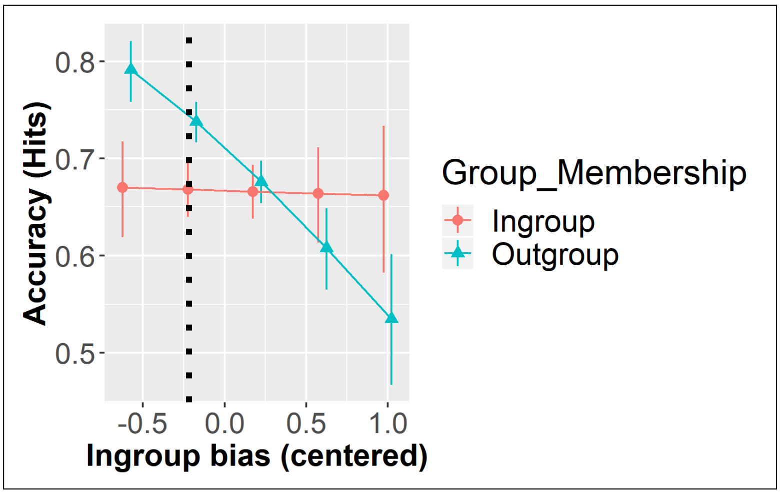

Figure 3.4: Accuracy on matching trials as a function of Group Membership and In-group Bias (centered; the black, dotted line represents in-group bias values equal to zero). Error bars represents standard errors.

Mismatching trials

To test whether speaker group membership influenced the level of detail in the speaker-specific information encoded with the new words, we analyzed accuracy on mismatching trials. By looking at participants' performance on within-university mismatching trials, where labels were paired with incorrect speakers who belonged to the same university as the original speaker, we could test whether the source-related information for novel words was speaker-specific (participants should have rejected the wrong source) or group-specific (participants would have incorrectly accepted the wrong source). Our hypothesis was that people encode more speaker-specific information from in-group members. We therefore predicted greater confusion among out-group speakers than in-group speakers in the within-university mismatching trials. We also predicted that this difference in accuracy would depend on In-group Bias, such that greater In-group Bias participants exhibited, the greater difference they should show between their accuracy on in-group vs out-group trials. Conversely, in between-university mismatches (i.e., where an in-group label was shown with out-group members, and vice versa) no differences due to our manipulation were expected.

To test these hypotheses, we ran a logistic mixed model analysis with fixed effects for Mismatch Type (Within- vs Between-university, reference level: Within), Group Membership (In-group vs Out-group, reference level: In-group), In-group Bias (z-scored continuous measure), and their interaction terms. We added per-participant and per-item random intercepts and by-participant slopes for Group Membership and Mismatch Type.

Overall, participants' accuracy on mismatching trials was 57.97% (SD=49.36) and above chance level (i.e., 50%, t=12.88, p<.0001).

Generally, participants were less accurate in the between–university mismatches than in the within–university mismatches (beta=-0.19, SE=0.08, z=-2.27, p=.02; mean=56.39%, SD=49.6 and mean=60.08%, SD=48.98, respectively). Participants' performance was also predicted by a significant interaction of Group Membership with In-group Bias (beta=-0.82, SE=0.28, z=-2.96, p=.003), suggesting that participants with different strength of In-group Bias were differently affected by speaker Group Membership. Crucially, both effects were further modulated by the expected 3-way interaction between Mismatch Type, Group Membership, and In-group Bias (beta=0.95, SE=0.37, z=2.61, p=.009). The remaining main effects and interactions were not significant (ps>.57).

Before further evaluating the three-way interaction, we tested its reliability. We ran a series of split-half analyses by randomly selecting subsets of participants composed of half the number of the original participants (N=30) over which we ran the same mixed-effect model. The interaction pattern was replicated in 15 out of 15 subsets, with the effect being significant in 13 subsets and marginally significant in two subsets. This procedure suggests that the significance of the interaction was well supported by the data and that the analysis was sufficiently powered such that even random subsets of half the sample led to comparable results.

To unpack the interaction, we ran separate analyses for each Mismatch Type condition (see Fig. 4). The analysis of the between-university trials included fixed effects for Group Membership, In-group Bias, and their interaction terms, as well as per-item and per-participant random slopes and by-participant slopes for Group Membership. Results did not show any significant predictor (ps>.42) (see Figure 3.5, left panel).

The analysis of the within-university mismatches included Group Membership, In-group Bias and their interaction as fixed effects, and per-participant and per-item random intercepts and by-participant slopes for Group Membership. In agreement with our predictions, the results showed a significant Group Membership x In-group Bias interaction (beta=-0.83, SE=0.28, z=-3.00, p=0.003). Post-hoc analyses revealed a cross-over interaction, but non-significant simple effects. The stronger the in-group bias, the more participants showed a difference between their ability to distinguish in-group members compared with out-group members (see Figure 3.5, right panel). These patterns show greater receptiveness to individual in-group speakers over out-group speakers and are in line with our predictions and with previous research showing that in-group bias stems from greater individuation of in-group members while viewing out-group members as less distinct (e.g., Hugenberg et al., 2010).

Next, we used signal detection theory to capture detection sensitivity (namely, d-prime) and decision criterion (namely, C). These analyses allow us to test whether participants' sensitivity and criterion during decision making differed for in-group vs out-group related decisions. We calculated two d-prime values and two C values per participant for in-group and out-group trials separately. In order to generate values that reflected participants' decisions to purely in-

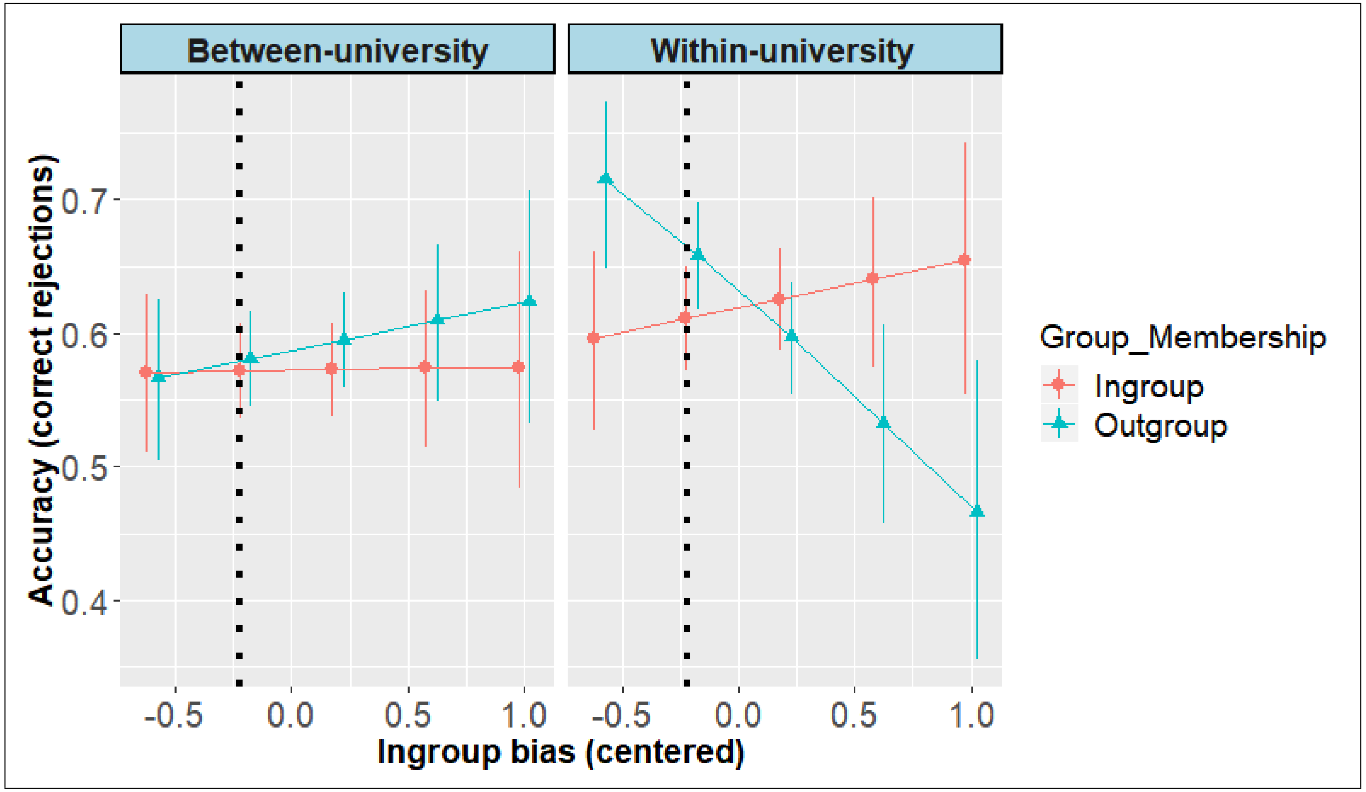

Figure 3.5: Accuracy for between-university mismatches (left panel) and within-university mismatches (right panel) as a function of In-group Bias (centered; the black, dotted line represents in-group bias values equal to zero) and the Group Membership of the speaker from whom the word was originally. Error bars represent standard error.

group or out-group trials, d-prime and C values were calculated from participants' performance in matching trials (i.e., hit rates) and within-university mismatching trials (i.e., false alarm rates). Between-university mismatches were not considered for these analyses because they contained an element (either label or speaker) from each group and were therefore not purely in-group or out-group related. We ran two linear mixed-effect models with either d-prime or C values as the dependent variable and Group Membership (In-group vs Out-group, reference level: In-group), In-group Bias and their interaction as fixed effects. The models included per-participant random intercepts.

The model which explored the relationship between individual d-prime and the independent variables showed no significant main effects for Group Membership (beta=0.13, SE=0.1, t=1.42, p=.16) or In-group Bias (beta=-0.01, SE=0.30, t=-0.04, p=.97). However, there was a significant interaction between these two factors (beta=-0.63, SE=0.29, t=-2.13, p=.04) (see Figure 3.6). Simple effects analyses revealed that the larger the participants' In-group Bias, the less sensitive they were to out-group speakers (beta=-0.64, SE=0.30, t=-2.12, p=.04), whereas sensitivity to in-group speakers did not change as a function of In-group Bias (beta=-0.01, SE=0.30, t<1, p=.97).

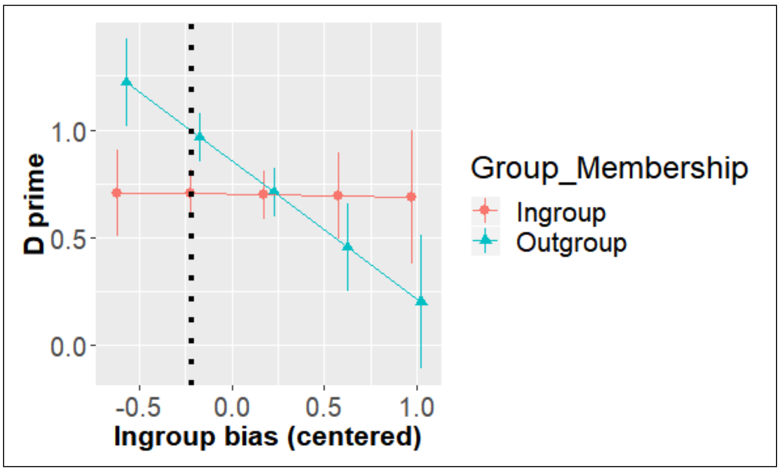

Figure 3.6: D-prime as a function of Group Membership and In-group Bias (centered; the black, dotted line represents in-group bias values equal to zero). Error bars represents standard errors.

The model exploring C values did not lead to any significant results ($p>.23$) suggesting that participants did not differ in decision criterion as a function of Group Membership or In-group Bias. This means that participants' tendencies to produce positive or negative responses were not affected by their social biases. That is, they were not more likely to say "yes" when seeing in-group speakers than when seeing out-group speakers, or vice versa.

Discussion

In a novel word learning study, we tested whether listeners encoded speakers' social identity, such as their group membership status and their personal identity, and whether the level of detail of such information depended on the group membership status of the speakers and the individual in-group bias of the listeners.

In the word learning task, participants had to identify matching speaker-label pairs in a forced-choice task (i.e., yes/no). Our results show that the level of detail with which participants encoded the identity of the speakers varied as a function of speakers' group membership status and participants' own in-group bias. Specifically, the analyses indicated that more in-group biased individuals exhibited greater confusion for out-group over in-group members. Our results are consistent with previous findings that showed greater individual-specific encoding for in-group than for out-group members (e.g., Frable & Bem, 1985; Greenstein et al., 2016; Wilder, 1990) and extend them to the case of word learning.

To better understand participants' behavior, we also analyzed their decision criterion and detection sensitivity for in-group and out-group speakers. Again, these analyses showed results that were comparable to those obtained in the previous analyses: the larger the in-group bias, the greater participants' detection sensitivity to in-group as compared to out-group speakers. That is, across all analyses (for matching trials, for mismatching trials and of d-prime), we find consistent results that demonstrate that speaker group membership and learners' individual in-group biases modulated the level of detail of speaker-related information contained in the representations of novel words.

One interesting aspect of our findings was that while we expected individual in-group biases to simultaneously increase source detection for the in-group and decrease it for the out-group, as Greenstein et al.'s results suggested, our data indicate mostly the latter phenomenon. The more in-group biased participants were, the less sensitive they were to individual out-group speakers, while their sensitivity to individual in-group speakers did not change as a function of in-group bias. There can be two potential explanations for these results. The most straightforward one relates to the possibility that in-group individuation, that is, the encoding of speaker-specific detail for in-group members, might be the default mechanism, regardless of individual in-group bias, and it is whether people also encode specific information about out-group speakers that depends on learners' own in-group bias. In the study by Greenstein and colleagues

(2016), participants were better at identifying in-group sources as compared to out-group sources. However, they did not include an individual measure of in-group bias in their analysis as a modulator of the effects of interest and, on the contrary, excluded participants who identified themselves with the in-group and out-group membership to a similar extent. Therefore, a direct comparison between our results and the results of their study is difficult to make.

An alternative explanation of this unexpected pattern is that the implicit measure we used to quantify in-group bias was better at capturing individual differences related to how well people attended to out-group members than differences related to the in-group. The distribution of accuracy rates in the perceptual matching task (see Appendix 2) shows that participants' performance for the in-group condition was almost at ceiling. Therefore, the variation in the individual difference measure we used was mostly driven by differences in the out-group condition. Future research should use a more taxing task that limits ceiling effects to determine whether individual measures of in-group bias would then be seen to also modulate participants' performance in in-group within-university trials. Although the in-group bias did not directly modulate the level of detail encoded in representations of in-group labels, the patterns of results confirmed our predictions.

Taking a broader scope, our results address a primary question in the field of psycholinguistics, namely, how words are represented in the mental lexicon. A growing body of literature suggests that social information, such as speaker identity, is stored together with the lexical information (e.g., Goldinger, 2007; Hay, Nolan, & Drager, 2006; Nielsen, 2011) (see Drager & Kirtley, 2016, for a review). Importantly, all the existing models assume that the same information will be stored to the same degree regardless of who the speaker is. In contrast, we show that the type of social information that is encoded (group level or individual level) and its robustness depend on speaker's identity (e.g., in-group vs out-group member) and the listeners' social biases. These findings, thus, call for the modification of existing models of language processing.

The nature of lexical representations is of great interest partly because it has implications for all aspects of language use – language processing, language production, and language learning. Our results therefore can have implications across the field of language use. We will next illustrate that with a couple of examples. There is great interest in how learners learn from input and what should be its distributional characteristics to optimize learning. Studies and models on the topic often contrast intra-individual and inter-individual variability. Our results, however, suggest that learners would be able to isolate intra-individual and inter-individual variability for input provided by in-group speakers, but less so for input provided by out-group speakers. These differences could lead learners to draw different inferences from the same input, depending on who provides it.

Another hot topic in psycholinguistic is the role of prediction in language learning and processing. Listeners are assumed to constantly predict utterance continuations, and use the prediction error (the difference between the input and the input they predicted) to update their representations, and thus learn (e.g., Chang & Bock, 2006; Garrod & Pickering, 2014; Kuperberg & Jaeger, 2016). Current research attempts to understand what information these predictions are based on and to which cases learned information is generalized (Kleinschmidt & Jaeger, 2015). Our results suggest that the identity of the speaker could influence the granularity of the predictions that listeners make, as well as how widely they would generalize the patterns they encounter. For example, listeners might be able to use speaker-specific information when predicting the speech of an in-group member, but only be able to use group-level information when predicting the speech of out-group members. Furthermore, poor individuation of out-group sources in people with high in-group bias may lead them to generalize information learned from a specific out-group member to all the out-group members. Imagine, for instance, what might happen if a politician from a disliked party labels an event in a particular way (say as "fake news" or "national triumph"). A member of an opposing party may overgeneralize the usage of this label ("fake news") to all members of that disliked party, thereby, perhaps, feeding unfounded beliefs about the views of the members of the disliked party. Further research should explore the consequences of different speaker identities and listeners' biases on language processing and learning as well as how in-group bias modulates the role of social factors in learning, processing, and transmitting language.

4 | Investigating on-line language processing of novel words learned from in-group and out-group speakers – an ERPs study

Abstract

Previous research showed that the information about the speaker's social identity is processed along with the linguistic content of the message. Furthermore, we recently showed that speaker group membership status (i.e., being an in-group or out-group speaker), in tandem with learners' individual in-group bias, determines the accuracy of the speaker-specific information encoded in the representations of novel words. Here, we investigated whether the on-line processing of novel words was influenced by the group membership of the speakers from whom they were learnt, after learning. Specifically, we tested whether words learned from in-group speakers were more easily processed than words learned from out-group speakers, in comparison to non-trained, unexpected labels. On day 1, participants learned labels for pictorially presented, unfamiliar gadgets. Half of the gadgets were labelled by in-group speakers and the other half by out-group speakers. Extensive training was provided to minimize group-membership-related differences in memory during learning. On day 2, we recorded EEG signal while participants saw speakers and gadgets. Crucially, they listened to labels in relation to those referents that were either trained labels, or non-trained, unexpected labels. ERPs were time-locked to the presentation of the label. Results revealed that non-trained vs trained labels elicited a predicted N400 effect, showing that the trained labels were represented in and retrieved from the mental lexicon. However, speaker group membership did not modulate this effect. These results suggest that when memory-related differences are controlled for during learning, the on-line processing of novel words might not be influenced by the group membership status of the speakers from whom those words were previously learned. Overall, these findings specify the role of social information during word learning and language processing and contribute to the understanding of socially situated language processing phenomena.

Introduction

Many models of word recognition posit that lexical representations are stored in the mental lexicon as abstract units (i.e., word forms or phonological representations) where information specific to speakers and contexts is left out (abstractionist accounts, e.g. Gaskell & Marslen-Wilson, 1997; McClelland & Elman, 1986; Norris, 1994). However, there is growing evidence that indexical information, such as speaker's social identity and voice characteristics, influences language processing at many different levels (e.g., Hay, Nolan, & Drager, 2006; Hay, Warren, & Drager, 2006; K. Johnson et al., 1999; Lev-Ari & Keysar, 2012; Niedzielski, 1999; van Berkum et al., 2008). In previous work, we have shown that not only is the speaker's social identity (in our case, the group membership status) encoded in the representations of novel words, but it can also influence, in interaction with learners' individual social biases, how accurately the source of the linguistic information is remembered (Iacozza, Meyer, & Lev-Ari, 2019b).

Departing from these findings, the current study tested whether learning novel words from speakers that belonged to the learners' in-group (i.e., in-group speakers) would facilitate the processing of these words in later encounters, as compared to words learned from out-group speakers. We tested this hypothesis by measuring electrophysiological responses (i.e., event-related brain potentials) while participants saw speakers and gadgets and listened to labels in relation to those gadgets. Crucially, the label was either *trained*, because it matched with what they learned in a previous learning session, or *non-trained*, because it had not been encountered before. The effect of group membership on N400 amplitudes in relation to non-trained vs trained labels allowed us to test the prediction that in-group words would receive a processing benefit as compared to out-group words.

Before turning to the current experiment, we review the relevant literature. First, we provide evidence that shows that information related to the speaker's social identity can influence language processing. Secondly, we review models of how language is processed and stored in the mental lexicon that assume concurrent encoding of both linguistic and indexical information. Finally, we discuss the proposal that the strength of such concurrent encoding may vary depending on the social relevance ascribed to the speakers by the listeners.

Previous evidence shows that speaker social identity can influence how listeners perceive speech sounds (e.g., Hay, Nolan, & Drager, 2006; Hay, Warren, & Drager, 2006; K. Johnson et al., 1999; Niedzielski, 1999). For example, responses in a diphthong identification task were influenced by what listeners believed was the speaker's place of residence. Despite the actual linguistic input being constant, participants reported hearing what they knew to be more representative of the supposed speaker's linguistic community, suggesting that information about the speaker can alter what is perceived (Niedzielski, 1999). Similarly, speaker social identity has been shown to affect how easily words are recognized (e.g., Kim, 2016; Walker & Hay, 2011). By using an auditory lexical decision task, Walker and Hay (2011) showed that the age of the speaker, inferred from the voice, speeded up listeners' responses to words rated as typical of speakers of that age, as compared to neutral-age words.

Along the same lines, listeners' expectations about upcoming linguistic input can be shaped by the social characteristics of the speaker they are listening to (e.g., van Berkum et al., 2008; Boland & Shana'e, 2014; Martin et al., 2016). Such social characteristics can in fact activate expectations about what will likely be uttered, given the specific speaker. Therefore, when these expectations are disconfirmed, language processing becomes more demanding (e.g., van Berkum et al., 2008; Boland & Shana'e, 2014; Martin et al., 2016). In an ERP study, van Berkum and colleagues (2008) recorded EEG signal while participants listen to utterances produced by speakers belonging to different social categories (e.g., female/male; upper/lower-class; child/adult). They found that when participants heard an adult speaker producing a sentence like "On the beach I made *sandcastles* by the sea", the word *sandcastles* elicited more negative N400 amplitudes than it did when produced by a child speaker. This N400 effect was found to have a similar topographical and temporal distribution to the classical N400 effect for semantic violations (see Kutas & Federmeier, 2011, for a review). Therefore, the authors interpreted the findings as evidence for an early involvement of information about the speaker' identity in the same language comprehension mechanisms that are used to extract sentence meaning from pure linguistic content.

To sum up, the evidence provided contradicts abstractionist models that assume linguistic representations to be stored as purely abstract units. There is a growing consensus about the need to have a theoretical account that models the encoding of both types of information (indexical and linguistic) (e.g., Cai et al., 2017; Hay, Nolan, & Drager, 2006; Kapnoula & Samuel, 2019; Nielsen, 2011; Münster & Knoeferle, 2018; Sumner et al., 2014) (see Drager & Kirtley, 2016, for a review). Most of the proposed models assume that linguistic experiences are stored in the mental lexicon as "exemplars", that is, episodic memory traces containing both speaker/context characteristics and linguistic information (e.g., Goldinger, 2007; Sumner et al., 2014).

Recently, Münster and Knoeferle (2018) put forward a formal model that defines how the concurrent encoding of both speaker identity and linguistic information would support on-line language processing. This model posits that, by extracting speaker-related information, listeners are able to calculate probabilistic contingencies related to the language typically produced by the speaker's linguistic community. Such a probabilistic approach would result in the pre-activation of a set of most likely lexical candidates. Therefore, when those items are indeed encountered, comprehension is speeded up and eased. As an example, take the sentence we referred to above "On the beach I made *sandcastles* by the sea" produced by a child speaker. Once the listeners extract the social features carried by the speaker's voice, i.e., the age, they can use such information to build up probabilistic inferences about the speaker's linguistic community in relation to the unfolding message. In this case, after hearing the speaker's reference to beach activities, the word *sandcastles* is more likely to be pre-activated by a child voice than by an adult voice, and sentence comprehension benefits from that lexical pre-activation (see Münster & Knoeferle, 2018, for details).

A critical open question is whether such encoding is the same across all speakers or varies as a function of the social relevance ascribed to the speakers. Sumner and colleagues (2014) proposed an account where the social context is not only encoded with the linguistic input but can also influence the strength of its encoding. The evidence in support of this account comes from work that compares reaction times and memory performance in relation to standard vs non-standard phonetic variants. Idealized, yet non-frequent, variants can be processed and remembered as fast and as accurately as non-idealized, but more frequent, variants. According to the authors, such evidence indicates that the social weight ascribed to contexts and speakers by language users can define the strength of the encoding of linguistic information, with socially weighted variants being more strongly encoded than non-socially weighted variants.

One group of speakers to whom a special status is ascribed are those that belong to the listeners' own social group (i.e., in-group speakers). Previous studies have showed that in-group-related information benefits from enhanced memory and processing, as compared to out-group information (e.g., Frable & Bem, 1985; Hugenberg et al., 2010; Van Bavel et al., 2008; Wilder, 1990). For instance, in-group faces are recognized more accurately than out-group faces (e.g., Van Bavel et al., 2008; Hugenberg et al., 2010) and the source of information (i.e., who said what) tends to be remembered better when the information is delivered by in-group vs out-group members (e.g., Frable & Bem, 1985; Greenstein et al., 2016; Wilder, 1990). Furthermore, people are faster and more accurate when learning arbitrary new associations between previously neutral stimuli (e.g., geometrical shapes) and their in-group memberships (e.g., the logo of their favorite football club) than when they learn associations involving an out-group membership (Enock et al., 2018; Moradi et al., 2015). This evidence suggests that in-group membership can boost general input processing and learning.

In previous studies, we expanded these findings to the case of language learning. Specifically, we tested whether there were differences in the level of detail of speaker-specific information encoded in the representations of novel words learned from in-group vs out-group speakers. The employment of novel words allowed us to control for prior exposure to the stimuli and to test for effects of group membership on the actual encoding of the novel lexical representations. Unlike previous studies (e.g., van Berkum et al., 2008), we focused on the mechanisms that integrate the social identity of the speakers during the *learning* of novel words. To achieve this purpose, participants learned competing labels for uncommon gadgets. Crucially, one label was produced by an in-group speaker (e.g., *citrus-peller*, in English *lemon peeler*) and another one by an out-group speaker (e.g., *citrus-schiller*, in English *lemon stripper*. Then participants' source memory of the new words was tested in a recognition test in which they saw speaker-label pairs and had to decide if the label had been previously produced by the displayed speaker. Additionally, we collected participants' own in-group bias, i.e., how strongly they preferentially encoded in-group vs out-group information.

Results showed that people with strong in-group bias a) had more accurate source memory of words learned from in-group vs out-group speakers (Iacozza et al., 2019b), and b) were more cautious in their decisions when in-group labels were involved (Iacozza, Meyer, & Lev-Ari, 2019a). These findings show that, firstly, the social identity of the speakers, such as their group membership status, can be encoded along with the novel linguistic input. These results provide further support for an exemplar-based account of language processing. Secondly, they show that speaker group membership status interacts with learners' own in-group bias in determining a) with what level of detail the speaker-specific information is encoded in the lexical representations, and b) how liberally decisions are taken when retrieving in-group vs out-group words.

A research question that directly sprouts from our previous findings relates to whether speaker group membership status and learners' individual biases would also affect the lexical processing of the novel words in later encounters, after they have been learned. That is, are words learned from in-group speakers more easily processed than out-group words? And if so, do individual in-group biases of learners modulate this tendency? These questions were addressed in the current study.

The current study

The main aim of the current study was to test the hypothesis that novel words learned from in-group members would be more easily processed than those learned from out-group members. We designed a two-day experiment in which we taught participants novel labels for uncommon gadgets, presented pictorially. Crucially, half of the gadgets were labelled by in-group speakers and the other half by out-group speakers. Since we were interested in testing group membership effects on the processing of the words *after* learning took place, we minimized memory-related differences between in-group and out-group words by training participants at ceiling via an extensive word learning phase (see Methods section for more details).

On day 2, we recorded participants' EEG signal while they saw the same speakers and pictorial referents presented along with spoken labels that either matched what had been learned on the previous day (i.e., *trained* labels) or did not (i.e., *non-trained* labels, they had never heard before).

From previous literature, we know that the processing cost experienced when a target word does not match the expectations driven by the preceding context leads to an increase in N400 amplitude, as compared to a word that fits the context (see Kutas & Federmeier, 2011, for a review). Here, we analyzed the N400 effect for non-trained, and therefore unexpected labels, compared to the trained labels (expected in the context of specific speakers and pictorial referents of objects). A significant N400 effect would show ease of processing and integration with the context of the newly learned labels, and reflect that they were indeed represented in and retrieved from the mental lexicon of our participants after extensive training.

Importantly, we predicted that when trained speaker-object pairs were presented to participants, in-group representations would be more strongly activated than out-group representations, because of the higher social relevance ascribed to in-group speakers. As a consequence, we predicted in-group labels to be pre-activated in a stronger and quicker fashion than out-group labels, resulting in a boosted lexical processing for in-group vs out-group labels. To test this prediction, we targeted potential modulation of the N400 effect by group membership of the speakers from whom the labels were learned (in-group vs out-group speakers) and participants' own in-group bias. We expected a larger N400 effect for in-group trials, as compared to out-group trials, and a modulation by in-group bias, such that the stronger the participant's in-group bias, the greater the difference between in-group and out-group N400 effect.

Additionally, we performed a series of analyses to gain a better understanding of participants' behavior during a) the learning session (day 1) and b) the memory tests that followed the EEG task (day 2). Specifically, we tested whether participants' learning trajectory and their source memory of new words, as well as the stability of the novel lexical representations, were influenced by speaker group membership status and participants' own in-group bias.

Finally, we calculated per-participant measures from their performance in the source memory test and the label memory test and investigated the relationship between these measures and the N400 effects for in-group and out-group words (see Methods section for more details).

Methods

Participants

Forty-four female native speakers of Dutch took part in the experiment (mean age=21.92 years old, SD=2.34). They were current or recent graduates from Radboud University Nijmegen. They gave informed consent as approved by the Ethics committee of the Social Sciences department of Radboud University (project code: ECSW2014-1003-196) and received monetary compensation for their participation. Two participants were excluded due to technical errors, another one was excluded because she reported afterwards being dyslexic, and an additional two participants were excluded because the number of trials that were rejected in the EEG preprocessing exceeded 20% of the trials in at least one experimental condition. The data from the remaining 39 participants were analyzed.

Materials

Six different tasks were implemented throughout the current study (see Procedure section for more details, and Figure 4.1 for an overview of the tasks). Below, we introduce the stimulus materials employed in each of these tasks, separately.

Materials for the speaker familiarization task and for the speaker memory test

Speaker faces: Four female faces were selected from the Chicago Face Database (Ma et al., 2015)for looking plausibly Dutch and of university age.

University logos: Black and white pictures of the logos of Radboud and Groningen universities were used.

Materials for the word learning tasks and for the source memory test

Speakers: The four female faces from the speaker familiarization task were paired with the voices of native Dutch female speakers recorded in our laboratory. Prior to the experiment, voices were matched for perceived typicality and attractiveness (paired t-tests, ps>.05) via a norming on-line survey in which twenty different participants participated. Each speaker was a unique combination of one face and one voice, consistent across participants.

Gadgets and labels: Sixty-three images of unfamiliar gadgets (e.g., lemon peeler) and their corresponding labels were selected via a norming study (see Appendix 1, for details). Forty-two were target gadgets and the remaining 21 were fillers. In the word learning tasks, all gadgets were presented with one label each. However, for the target gadgets, there were two possible labels (e.g., *citrus-peller*, in English *lemon peeler*, and *citrus-schiller*, in English *lemon stripper*), equated for frequency and goodness-of-fit. Prior to the experiment, for each participant one label was randomly selected among the two available labels. All labels uttered by the two speakers were audio-recorded.

Materials for the EEG task

Speakers: The same speakers from the previous tasks were used.

Gadgets and labels: The same gadgets from the word learning tasks were used. In this task, however, target gadgets were presented with both available labels: the one initially selected for the word learning task (hereinafter, referred to as trained label), and the one that had not been previously used (hereinafter, referred to as non-trained label).

A trained native speaker of Dutch identified the onset of the acoustic disambiguation point for each target label recording, which was operationally defined as the point in the acoustic signal where the two versions of the labels per speaker began to diverge in terms of their respective phonemes (see van Berkum et al., 2008, for a similar procedure applied to the acoustic inflection of critical adjectives). For the *citrus-peller–citrus-schiller* example pair, we estimated the acoustic disambiguation point to be at the offset of the sibilant -s- in citrus. However, for the *lepelklem–spatelhouder* example pair, the acoustic disambiguation point corresponded with the acoustic onset of the words, therefore was set at 0 ms. Across all target labels, and relative to their acoustic word onset, the mean disambiguation point was at 259 ms (range 0–787 ms). Due to this great variability that characterized our stimuli, the use of the disambiguation point to time-lock the ERP analyses was a much more fine-grained approach than the use of word onset.

Materials for the gadget-label recognition test

Gadgets and labels from the EEG task were used.

Procedure

The experiment was administered in two sessions of about 2.5 hours each that took place over two consecutive days (see Figure 4.1, for an overview of the procedure).

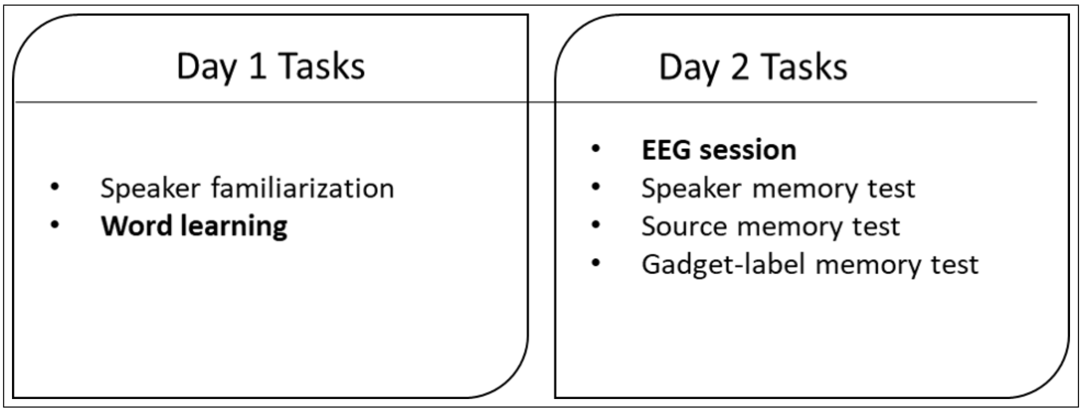

Figure 4.1: Overview of the procedure of the experimental sessions, across two consecutive days.

Day 1

Speaker familiarization

The goal of the speaker familiarization task was to familiarize participants with the speakers and their affiliations. In this task, participants learned to associate the four speaker faces with the logos of their supposed academic affiliations. Two faces were introduced as students from the participants' own university (i.e., Radboud University) and the two other faces were students from another university in the north of the Netherlands (i.e., Groningen University). Speaker-affiliation pairings were randomized.

During the exposure phase, participants saw face-logo pairs ten times and were asked to memorize which student belonged to which university. Afterwards, they performed a match/mismatch task in which they saw again face-logo pairs, but these either matched or mismatched the initially learned associations. In the task, a fixation cross (500 ms) preceded a blank screen (between 1000 and 2000 ms) and the presentation of logo-face pairs (600 ms). Participants' task was to indicate whether the pairs were correct or not via key press (y/n) within a time limit of 2100ms. One-hundred-sixty test trials, presented in two blocks of 80 trials each, were administered following 32 practice trials. Visual feedback was provided only during practice. Participants' performance was assessed after the speaker familiarization task, before moving to the word learning task. Participants performed on average at 93.7% accuracy (SD=24.23).

Word learning tasks

In the word learning tasks, participants saw the students from the speaker familiarization task referring to 64 unfamiliar gadgets, pictorially presented, with novel labels. Forty-two were target gadgets and 22 were fillers. Prior to the experiment, for each participant the gadgets were randomly divided into two halves: one half was referred to by in-group speakers exclusively, whereas the other half was referred to by out-group speakers, exclusively. Each participant heard only one label per target gadget, which was randomly selected from the two available labels. Due to the large number of gadgets, we randomly divided them into four groups of 16 gadgets each[1]. For each group of gadgets, participants performed three different tasks in a row. After the first learning round, which comprised all four groups of items, participants took a break of 5-10 mins and underwent a whole new learning round (see Figure 4.2 for an overview of the word learning procedure).

The three tasks are explained in details below.

[1]The first three groups consisted of 10 target gadgets and 6 fillers, whereas the last group of items consisted of 12 target gadgets and 4 fillers.

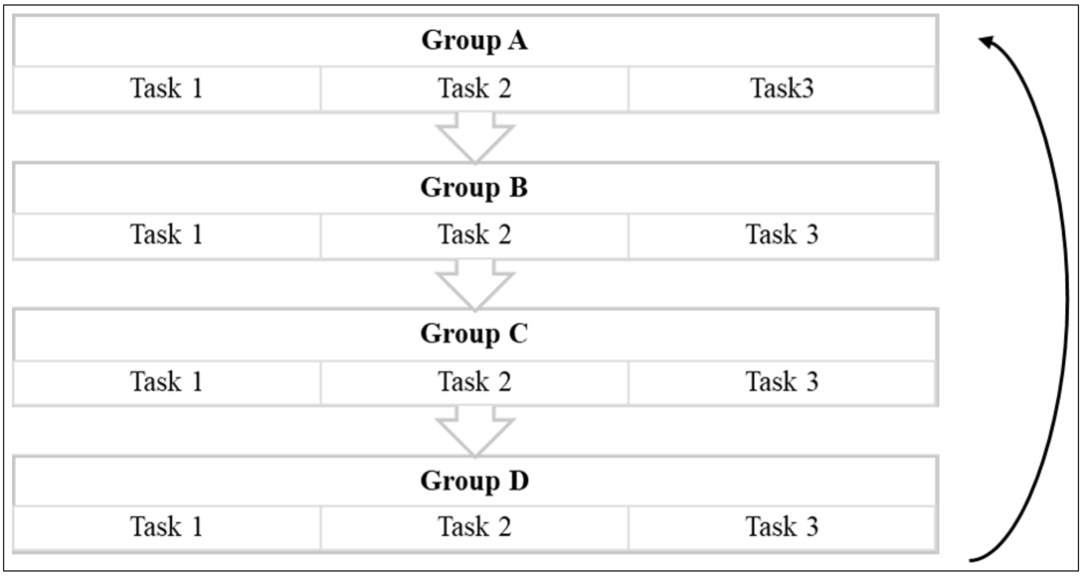

Figure 4.2: Overview of the procedure in the word learning tasks. Each group of gadgets was trained in three different tasks across two learning rounds.

Task 1: passive listening

In the first task, participants passively listened to the speakers referring to the gadgets. As Figure 4.3 shows, each trial began with a fixation cross (1000 ms), followed by a photo of a speaker. After 500 ms, a picture of a gadget was displayed as well (500 ms). Then, while speaker-gadget pairs were still present on the screen, the recording of the corresponding label was played. The next trial began 1500 ms after the presentation of the spoken label.

There was a total of 64 learning trials divided across two repetitions of 32 trials each. In each repetition, the gadgets were presented in blocks according to the speaker affiliations. This means that the first 16 trials of one repetition were related exclusively to one set of speakers (e.g., in-group speakers), whereas the last 16 trials were related to the other set of speakers (e.g., out-group speakers). The order of blocks per affiliation was counterbalanced across groups of gadgets, repetitions and participants. Each gadget was labelled by both speakers belonging to the specific affiliation (i.e., eight gadgets by two speakers equals 16 trials per group of items). The presentation of the trials and the order of the speakers was randomized.

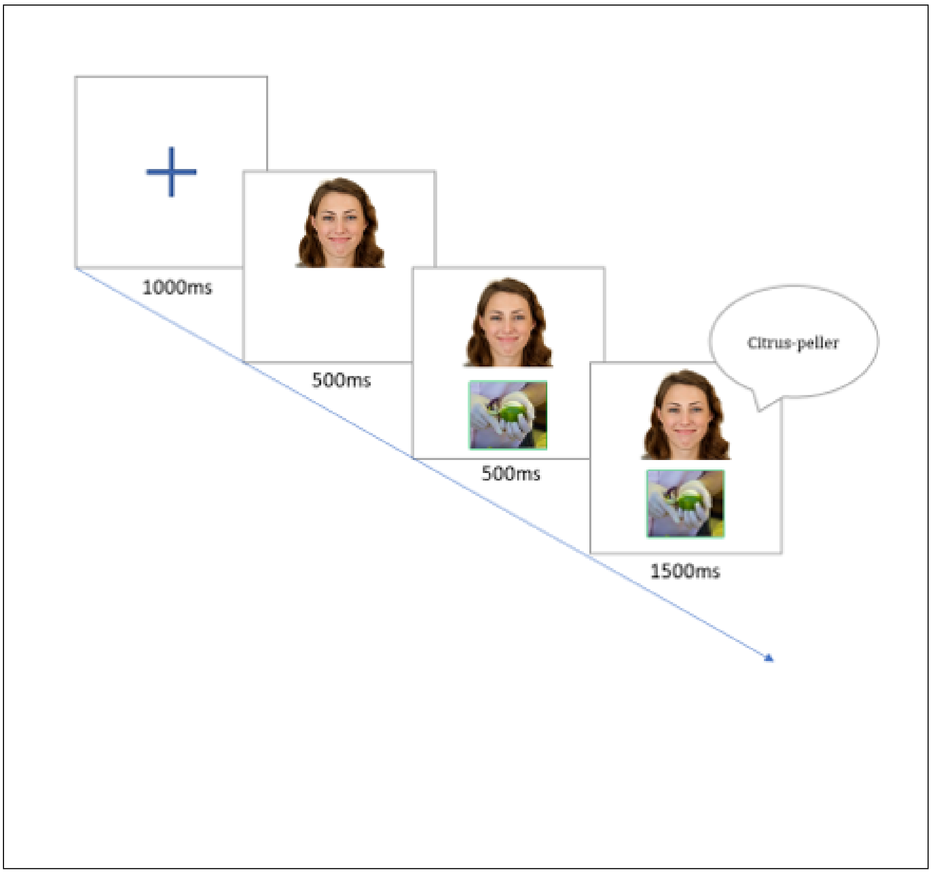

Figure 4.3: Trial sequence of the first learning task.

Task 2: speaker-voice association task

The aim of the second task was twofold: a) to create stable associations between the speakers and their voices via an associative task, b) to strengthen associations between the speakers and their corresponding academic affiliation. As Figure 4.4 shows, each trial began with a fixation cross (500 ms). A picture of a gadget was then displayed for 500 ms, followed by the photos of the four speakers and the auditory presentation of the corresponding label. Participants were instructed to identify the speaker that had produced the label within a time limit of 3500 ms after which auditory and visual feedback was provided for both correct and incorrect (or late) responses. Feedback was given by playing a sound that indicated a correct (i.e., ding sound) or a wrong answer (i.e., buzzer sound), and showing the participants what they correctly selected/should have selected.

There was a minimum of 32 trials, and if participants incorrectly selected a speaker belonging to the opposite affiliation of the correct one, the whole trial automatically repeated at the end. This procedure repeated until participants selected either speaker belonging to the correct affiliation. This was done to train participants at ceiling in relation to speaker-affiliation-gadget associations. On average, participants needed 34.45 total trials (SD=3.62) per group of items to reach ceiling.

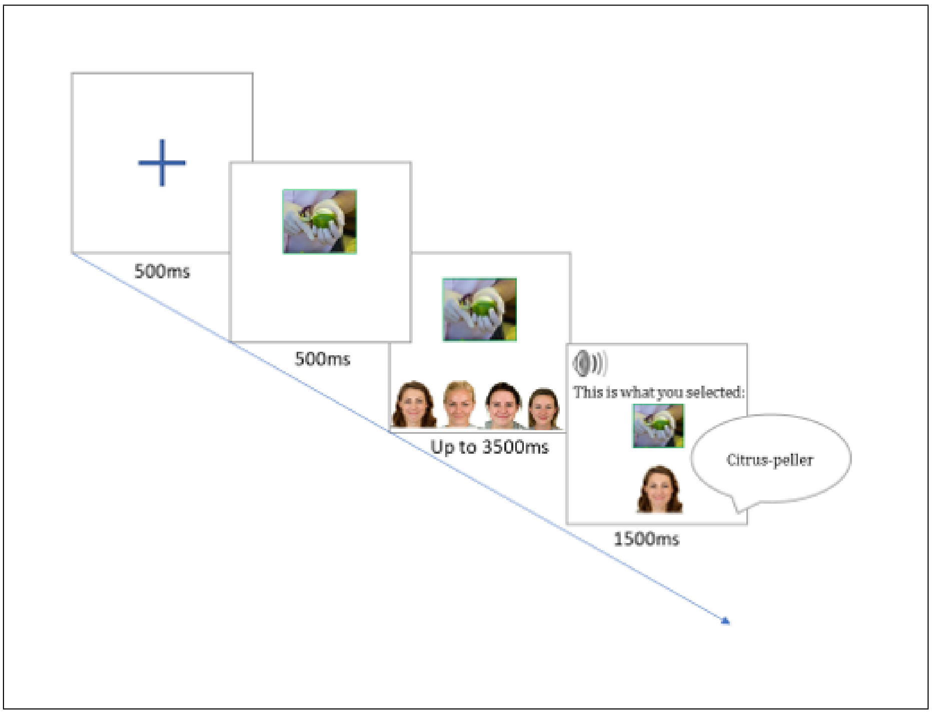

Figure 4.4: Trial sequence of the second learning task. The feedback was given after each response.

Task 3: Speaker-gadget associative task

The aim of the third task was to create stable associations between speakers and gadgets and make them aware that certain gadgets were exclusively labelled by a set of speakers, and others by the other set of speakers. As

Figure 4.5 shows, each trial began with a fixation cross (500 ms), followed by the presentation of speaker-gadget pairs along with three written options that they were instructed to choose from with a time limit of 3500 ms. The options were: a) the correct label of the gadget, b) a label of another gadget, randomly chosen, and c) a written sentence in Dutch stating that the speaker did not refer to the displayed gadget (see Fig. 4 for the trial sequence). The presentation order of the written options was randomized. Auditory and visual feedback was always given after both correct and incorrect responses by playing a sound associated to correct or wrong answer, and showing the participants both speakers that labelled the displayed gadget (one by one and playing the recording of the label).

If participants were always correct, 64 trials were administered, half of each were mismatches (i.e., where participants should have indicated that the speaker had not labelled the gadget) and the remaining half were matches (i.e., where indeed the speaker had labelled the gadget). If participants produced a false alarm (by incorrectly indicating that the speaker labelled the item), the whole trial was automatically repeated at the end. This procedure repeated until participants selected the correct response and was done to strengthen the associations between the speakers and gadgets and train participants at ceiling. On average, participants needed 68.42 total trials (SD= 5.44) per group of items to reach ceiling.

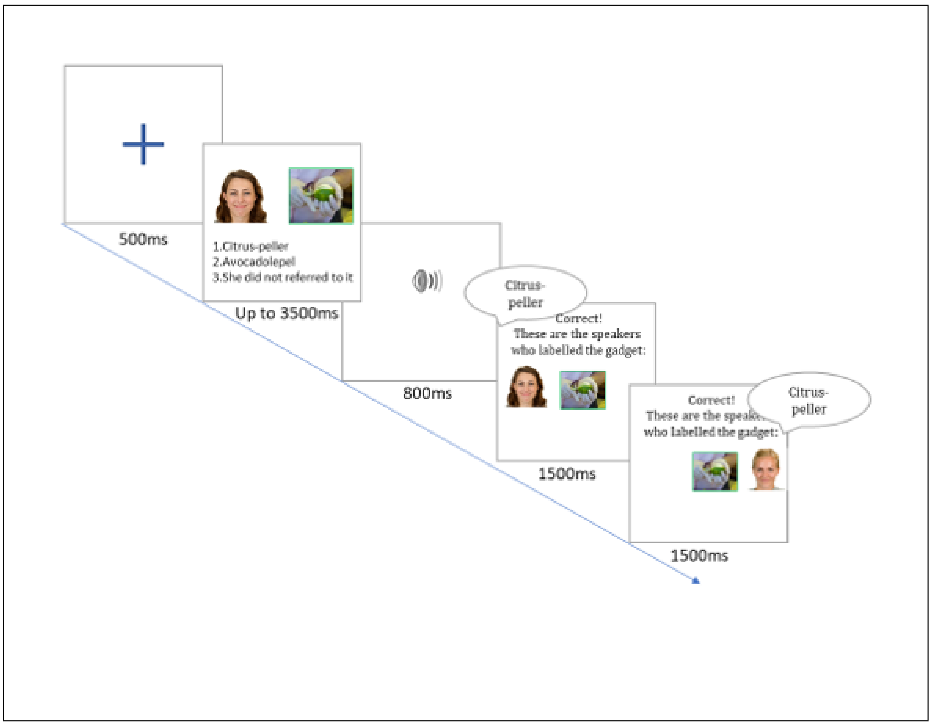

Figure 4.5: Trial sequence of the third learning task. The feedback was given after each response.

Day 2

EEG task

On Day 2, participants came back and were prepared for the EEG session. The aim of the EEG task was to test whether participants had learned the labels from the previous day so that they would create specific expectations about the labels when presented with the pictures of the gadgets. This would result in processing ease/difficulty when confronted with a trained or non-trained label, reflected in N400 amplitudes upon presentation of the label. Furthermore, we tested whether the N400 effect varied as a function of the speakers being in-group or out-group members. Expectations were activated by the presentation of a series of stimuli consisting of a photo of a speaker (800 ms), followed by the picture of a gadget (800 ms) and the recording of the label, which was delayed of 330 ms (see Figure 4.6 for the trial sequence).

Crucially, participants heard either the trained label (i.e., the one they learned on the previous day in relation to the displayed speakers and gadgets) or a non-trained label (i.e., a competing, yet equally fitting, label that they never heard before).

There were two blocks of presentation with a total of 256 trials: 88 trials involved filler items, which were shown always with the trained label, 84 were trials where target gadgets received a trained label (42 trials involving in-group speakers and 42 involving out-group speakers), and the remaining 84 trials contained target gadgets with a non-trained label (42 from in-group speakers and 42 from out-group speakers). Therefore, on two thirds of the trials trained labels were presented, and only in one third of the trials non-trained labels were presented. For target gadgets, participants saw each speaker-gadget pair in both trained and non-trained conditions, but across two presentation blocks. For instance, if speaker A from the in-group affiliation used the trained label for item 1 in the first block, then speaker B with the same affiliation referred to item 1 with the non-trained label, in that block. In the second block, speaker-gadget pairs were presented in the opposite conditions (i.e., for item 1, speaker A used the non-trained label and speaker B used its trained counterpart). The condition of speaker-label assignment was pseudo-randomized, so that speakers were presented with an equal number of trained/non-trained labels in each block.

Importantly, participants were not instructed to predict the labels according to the sequence of the stimuli presented to them. On the contrary, after each speaker-gadget-label sequence in the trial, they performed a 1-back cover task. They had to decide which stimulus had just been presented in the previous sequence, out of two options which were displayed side by side on the screen. Participants had to press a left/right key for identifying the correct stimulus, the position of which was randomized. For one third of the trials, the decisions concerned the face of the speaker (i.e., two faces were shown and participants selected the one that had just appeared), for another third of the trials, they concerned the gadget (i.e., pictures of two gadgets appeared) and for the remaining third of the trials, they concerned the label (i.e., two written labels appeared). This task was trivial for the participants (accuracy = 98.23% of the trials, SD=13.20) but it served to keep them alert.

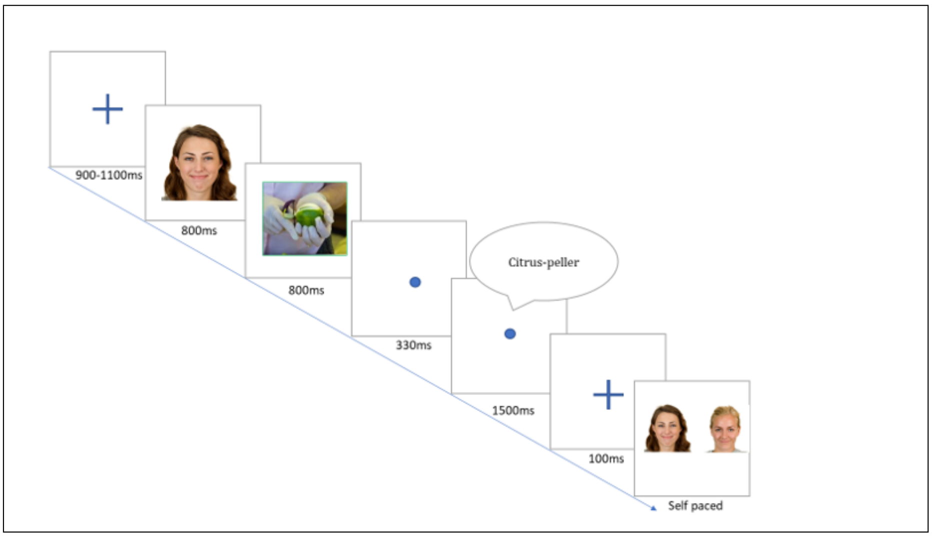

Figure 4.6: Trial sequence in the EEG task.

Speaker memory test

The procedure of the speaker memory test was identical to the test phase in the speaker familiarization task on Day 1, therefore without exposure and practice phases. The aim of the speaker memory test was to calculate individual in-group bias measures based on participants' performance in the task. Specifically, we obtained the individual measure by calculating per-participant Cohen's D values from accuracy rates for in-group versus out-group matching trials.

Source memory test

The source memory test was a sanity check to confirm that participants remembered the source of the labels they had learned on the previous day. In this task, they were presented with all possible combinations of pairing between the speakers and the trained labels. The trained labels were presented auditorily in the voice of the displayed speaker. Their task was to decide if the speaker had produced the label during the learning tasks. Responses were self-paced.

Gadget-label memory test

The gadget-label memory tested whether participants had indeed learned the labels on day 1. In this task, they saw each target gadget along with both labels (the trained and the non-trained one) and had to indicate which was the one they learned on the previous day by pressing the corresponding key. Responses were self-paced.

EEG recording and preprocessing

During the EEG task, EEG data were acquired from 59 active scalp electrodes, mounted in an elastic cap (ActiCAP), placed according to the 10-20 convention. Additionally, two electrodes were placed on the mastoids and four EOG-dedicated electrodes were used to detect eye-movements and blinks –to the left of the left eye, above and below the left eye and to the right of the right eye. The signal was referenced online to the left mastoid, amplified using BrainAmp amplifiers with a band-pass filter of 0.016-150Hz and recorded using Brain Vision Recorder (Brain products GmbH). All electrode impedances were kept below 5kΩ. The recording was done at a sampling rate of 500 Hz, with a time constant of 10 s.

In MATLAB (Version R2016a; the MathWorks Inc.), we used the FieldTrip toolbox (Oostenveld, Fries, Maris, & Schoffelen, 2011) to perform the off-line preprocessing procedure in the following order. We re-referenced the EEG data to the average of the two mastoids and band-pass filtered the data at 0.1-100 Hz. Channels exhibiting either excessive noise or strong drifts were manually detected and removed. We segmented the data into trial epochs using the time of the initial fixation cross display as starting point, and the end of the spoken word presentation as the ending point, which led to epochs of ca. 5.55 s[2]. Eye-movement artifacts were detected and removed using independent component analysis (ICA) with the runica algorithm implemented in Fieldtrip. Afterwards, the bad channels, initially excluded, were interpolated based on the activity of their neighboring channels, using the weighted triangulation method. All data was then low-pass filtered at 30 Hz. Subsequently, we again segmented the data into epochs that were time-locked to the disambiguation point of the target labels, starting at -100 ms and ending at 1000 ms after it. Segments containing artifacts, hence values ± 100 μV or outside a 150 μV range (min-max difference) were invalidated and thereby excluded from further analyses. We further excluded trials related to filler items and those with incorrect responses from all analyses.

The remaining trials of all participants were combined in a 3D matrix (channels, time points, trials by participants), which formed the basis for all further ERP analyses. In addition, single-trial mean baseline values were calculated by averaging over the EEG data corresponding to the 100 ms prior to the actual word onset. This time interval was comparable across trials, since it contained the activity elicited by the fixation dot display preceding the presentation of the spoken labels.

[2]This first segmentation into long epochs was performed to increase the effectiveness of the ICA procedure. By submitting a relatively large number of datapoints containing eye-movements (participants were instructed to blink on the fixation cross), we maximized ocular artifact detection.

Results

EEG

Analyses

Similar to Frömer, Maier, and Abdel Rahman, we followed an analysis pipeline in which we first generated a 3D matrix (channels, time points, trials by participants), then we used aggregated data by condition and by participant to perform cluster-based permutation tests, as implemented in Fieldtrip, for data screening and ROI selection. To limit circularity in the statistical analysis interpretations, permutation tests were run to define time windows and regions of interest only for the main effect of Label Identity (trained—non-trained), and not for its interaction with Group Membership (in-group—out-group). Secondly, we submitted the single-trial data to linear mixed model analyses in R-Studio (R Core Team, 2014). An overview of the analysis pipeline is shown in Figure 4.7.

In order to perform cluster-based permutation tests as implemented in Fieldtrip, data from the 3D matrix was aggregated by condition and by participants, and baseline corrected. Relevant time windows of the single trial EEG data, as determined using the permutation tests, were then exported for single-trial linear mixed model analyses (see Results for specific time windows and regions of interest). Linear mixed model analyses of ERPs were conducted using the lme4 R package (Bates et al., 2014) and p-values were calculated with the lmerTest package (Kuznetsova, Brockhoff, & Christensen, 2017), using Satterthwaite approximation for degrees of freedom.

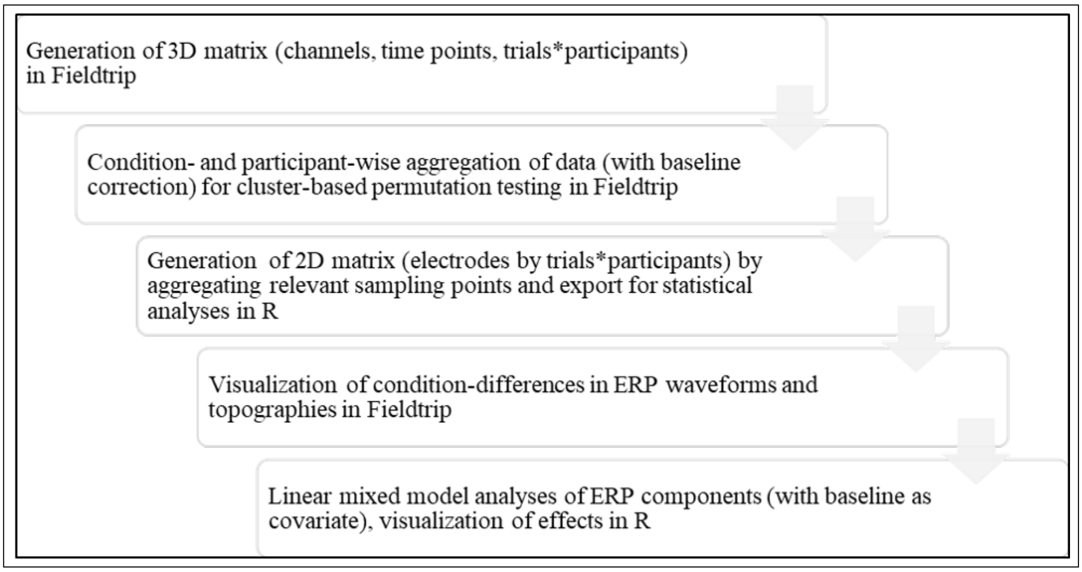

Figure 4.7: Flow chart of the different parts of the EEG analysis pipeline.

Defining time and region of interest for ERP analyses

Cluster-based permutation tests comparing mean amplitudes over the epoch 0-800ms revealed that the trained and non-trained label conditions differed significantly. As shown in Figure 4.8, differences started around 52 ms after disambiguation point onset and remained until 696 ms after disambiguation point. Three significant clusters were identified. Between 52 ms and 206 ms, the non-trained vs trained label condition evoked more positive amplitudes at central and centro-parietal electrode sites (p=0.002). Furthermore, two clusters with more negative amplitudes in the non-trained compared to the trained condition were observed, between 336 and 538 ms at centro-parietal and occipital electrodes (p¡.001) and between 582 and 696 ms at centro-parietal electrodes (p=.02). The N400 typically peaks between 300 and 500 ms and is thought to reflect integration processes. Therefore, we limited the follow-up analyses to a centro-parietal N400 ROI between 350 and 550 ms (electrodes POz, Pz, CPz, CP1, CP2, P1, P2, PO4, P03), which corresponded with the second identified cluster. Grand averaged ERP waveforms on label conditions across Group Membership conditions are shown in Figure 4.9.

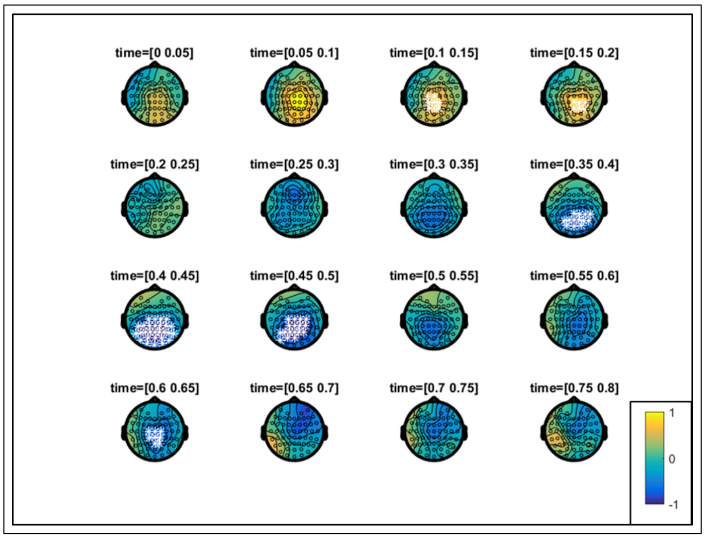

Figure 4.8: Results of the CBPT on the main effect of Label Identity (non-trained-trained). Electrodes that are part of clusters with p-values<.05 are highlighted in white in the corresponding time window. Color bar shows the difference in μV between the two Label Identity conditions.

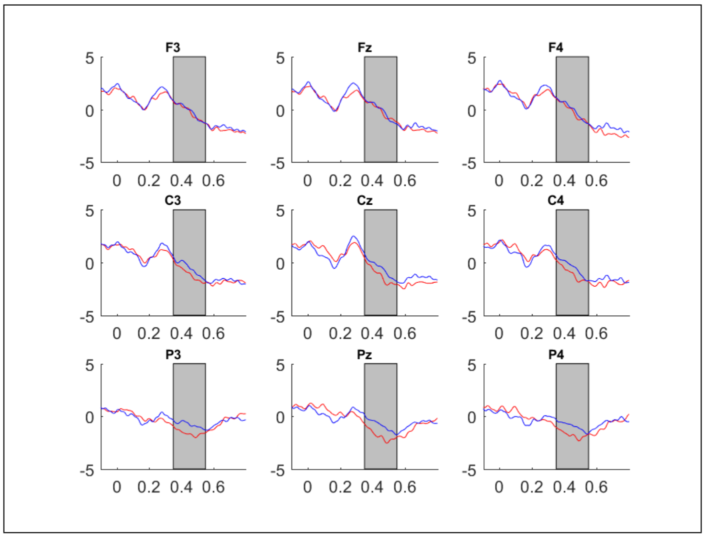

Figure 4.9: ERPs elicited by target labels in the trained (blue line) and non-trained (red line) conditions on individual channels. Shaded area represents the analyzed time window (350 - 550 ms). NB: negative polarity is plotted downwards.

N400 amplitude

To test the effects of our experimental manipulations on N400 amplitude, we regressed Label Identity (trained vs non-trained), Group Membership (in-group vs out-group) and their interaction on N400 amplitudes per trial. Following Alday (2019), we included the single-trial mean baseline values (scaled and centered) as covariate. As random effects, we modelled per-participant and per-gadget random intercepts and a by-participant random slope for Label Identity. This was the maximal model that converged. For all predictors, we applied sliding difference contrasts (i.e., -0.5, 0.5), thus the resulting estimates can be interpreted as the difference between conditions per predictor (trained minus non-trained; out-group minus in-group). The model estimates are summarized in the table in Figure 4.10 and can be read out directly as mean differences in μV for main effects.

	N400 amplitude				
Predictors	*Estimates*	*Std. Error*	*CI95%*	*t-value*	*p-value*
Intercept	-1.34	0.18	-1.70 – -0.98	-7.32	**<0.001**
Scaled Baseline	0.07	0.07	-0.07 – 0.20	0.94	0.345
Group Membership	-0.21	0.14	-0.48 – 0.06	-1.52	0.129
Label Identity	-0.85	0.20	-1.24 – -0.46	-4.31	**<0.001**
Group Membership x Label Identity	-0.01	0.27	-0.55 – 0.53	-0.04	0.971
Random Effects					
σ^2	29.92				
τ_{00} ItemID	0.24				
τ_{00} Participant	0.90				
τ_{11} Participant.LabelIdentity2-1	0.79				
ρ_{01} Participant	0.24				
ICC ItemID	0.01				
ICC Participant	0.03				
Observations	6361				
Marginal R^2 / Conditional R^2	0.006 / 0.049				

Figure 4.10: Effects of Group Membership and Label Identity on N400 amplitude.

We observed a significant main effect of Label Identity, with a reduced N400 for non-trained compared to trained labels. N400 amplitude was not predicted by either the main effect of Group Membership, nor by its interaction with Label Identity, suggesting an N400 effect of equal magnitude for in-group and out-group labels (see Figure 4.11). Figure 4.12 shows the scalp topographies for the main effect of Label Identity separately for in-group and out-group labels.

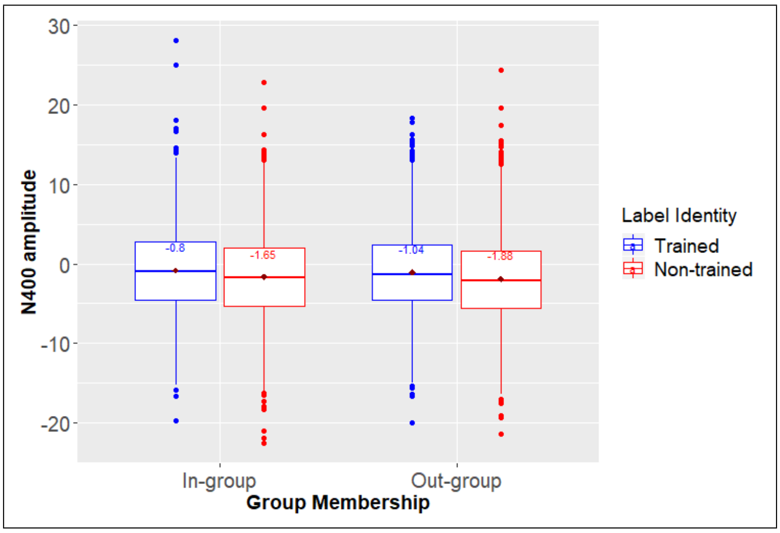

Figure 4.11: Distribution of N400 amplitude per condition. Box-plots provide information about median, and 1st and 3rd quartiles; red dots and digits correspond to mean values.

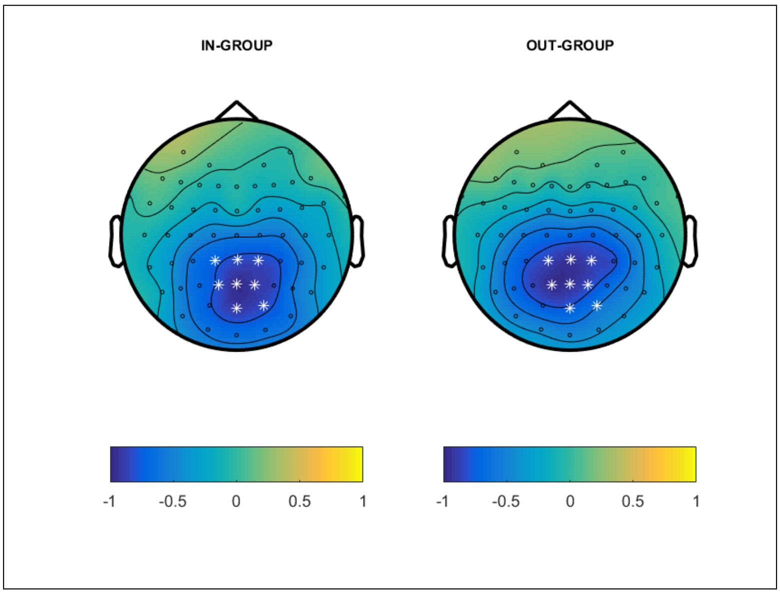

Figure 4.12: Topographical distribution of the difference in ERP amplitudes between the trained and non-trained conditions for target labels in in-group and out-group conditions separately, in the time-window 350–500 ms. The figure shows the ERP voltage of the non-trained minus the trained label. ROI channels are highlighted.

Behavioral data

Learning speaker - gadget associations

We focused on Task 3 of the word learning session (day 1) in which participants learned the associations between speakers and gadgets that were later measured in the EEG session. We tested whether participants' behavior a) reflected learning across the task (i.e., training trials decreased from round 1 to round 2) and b) was affected by speakers' group membership status and participants' in-group bias.

Note that on each trial in this task, participants saw a speaker and a gadget along with three written options. Among the options, there were the label they learned for depicted gadget (independent of whether the speaker had produced it), an incorrect label, and the option to indicate that the speaker had not referred to that gadget before. Whenever an error was made, the trial display was saved and repeated at the end of the task until the correct response was given. Therefore, the more errors a participant made (Total Trials hereafter), the more the trials per item they received. We operationalized learning behavior by counting per-item total trials. This was calculated for each round of exposure.

We regressed Round (round 1 vs round 2), Group Membership (in-group vs out-group), In-group Bias (centered) and their interaction on Total Trials per item within each round and for each participant. As random effects, we modelled per-participant and per-item random intercepts. This was the maximal model that converged. For all predictors, we applied sliding difference contrasts (i.e., -0.5, 0.5). The model estimates are summarized in the table in Figure 4.13.

	Tot Trials per item				
Predictors	*Estimates*	*Std. Error*	*CI95%*	*t-value*	*p-value*
Intercept	4.28	0.04	4.21 – 4.36	117.38	**<0.001**
Group Membership	-0.02	0.02	-0.06 – 0.02	-1.16	0.248
Round	-0.20	0.02	-0.24 – -0.16	-10.21	**<0.001**
In-group Bias	-0.05	0.79	-1.60 – 1.49	-0.07	0.947
Group Membership x Round	0.04	0.04	-0.04 – 0.12	1.03	0.305
Group Membership x In-group Bias	0.95	0.43	0.11 – 1.79	2.23	**0.026**
Round x In-group Bias	1.83	0.43	0.99 – 2.66	4.28	**<0.001**
Group Membership x Round x In-group Bias	-0.94	0.85	-2.61 – 0.74	-1.10	0.272
Random Effects					
σ^2	0.30				
$\tau_{00\ Item}$	0.00				
$\tau_{00\ PP}$	0.05				
ICC_{Item}	0.00				
ICC_{PP}	0.13				
Observations	3192				
Marginal R^2 / Conditional R^2	0.034 / 0.164				

Figure 4.13: Effects of Group Membership, Round and In-group Bias on Total Trials per item in the word learning tasks.

We observed a significant main effect of Round. On Round 2, they made fewer errors than on Round 1 (Round 1: mean=4.38, SD=0.71; Round 2: mean=4.18, SD=0.44). There was a significant Round x In-group Bias interaction, which was due to a cross-over interaction with no significant simple effects (ps>.3), as follow-up analyses confirmed. The total number of trials was not predicted by the main effect of Group Membership, or by its interaction with Round, suggesting that participants needed an equal number of trials for in-group and out-group associations across the two rounds.

Importantly, there was a significant Group Membership x In-group Bias interaction, suggesting that participants with different magnitude of in-group bias were differently affected by the Group Membership of the speakers from whom they learned (see Figure 4.14). Simple effect analyses revealed a cross-over interaction where simple effects were not significant (in-group: beta=-0.53, SE= 0.82, t=-0.65; out-group: beta=0.42, SE=0.82, t=0.52). The 3-way interaction failed at reaching significance.

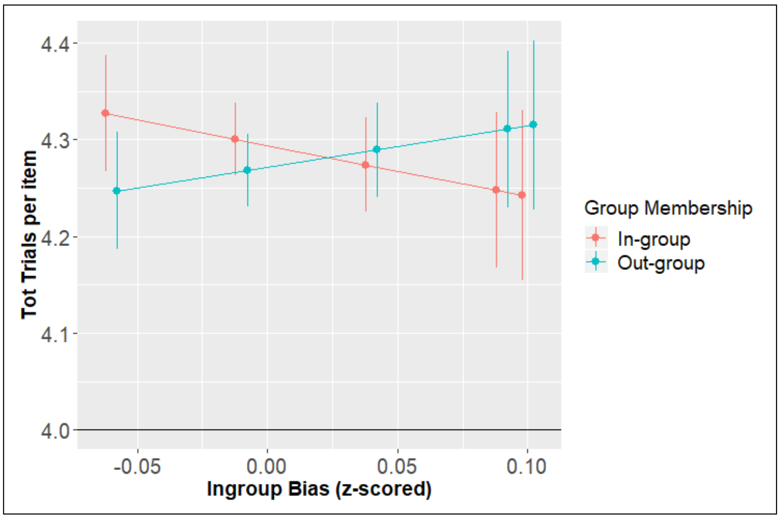

Figure 4.14: Effects of Group Membership and In-group Bias (z-scored) on the total number of trials per item across repetitions in the word learning tasks. Error bars represent Standard Errors. The horizontal black line at the bottom is set to the minimum number of trials when no errors were made.

Despite the 3-way interaction not being significant, we performed exploratory analyses on data subsets by Round levels to shed light on the patterns of results. We ran a separate model for each Round level with fixed effects for Group Membership, In-group Bias and their interaction on Total Trials per item. Per-participant and per-item random intercepts were included.

Results from the model referring to Round 1 showed a significant Group Membership x In-group Bias interaction (beta=1.40, SE=0.69, t=2.03, $p<.05$) (see Figure 4.15 left panel). Simple effect analyses showed non-significant trends: the stronger the in-group bias, the stronger was the tendency to need fewer training trials for both in-group and out-group items. Interestingly, the size of the in-group related simple effect was more than six times bigger than the size of the out-group related effect (in-group: beta= -1.67, SE=1.26, t=-1.32, p=.19; out-group: beta= -0.27, SE=1.26, t=-0.21, p=0.83), explaining the significant interaction. This suggests that during the first round of exposure participants with stronger in-group bias were tendentially more likely to commit fewer errors for in-group than for out-group associations.

In Round 2, the Group Membership x In-group Bias interaction was not significant anymore (beta= 0.5, SE=0.5, t=1.08, p=.28) (see Figure 4.15 right panel).

Source memory

First, we checked whether participants learned the speaker-label associations well. Overall, their accuracy was 93.74% (SD=24.23) and above chance level, as confirmed by a one-sample t-test (i.e., 50%) (t=340.07, $p<.0001$).

Second, we followed the procedure adopted in our previous studies (Iacozza et al., 2019b, 2019a) and analyzed participants' performance in the source memory test (on day 2) by looking at signal detection theory measurements. Specifically, we tested whether participants' sensitivity (indexed by d-prime

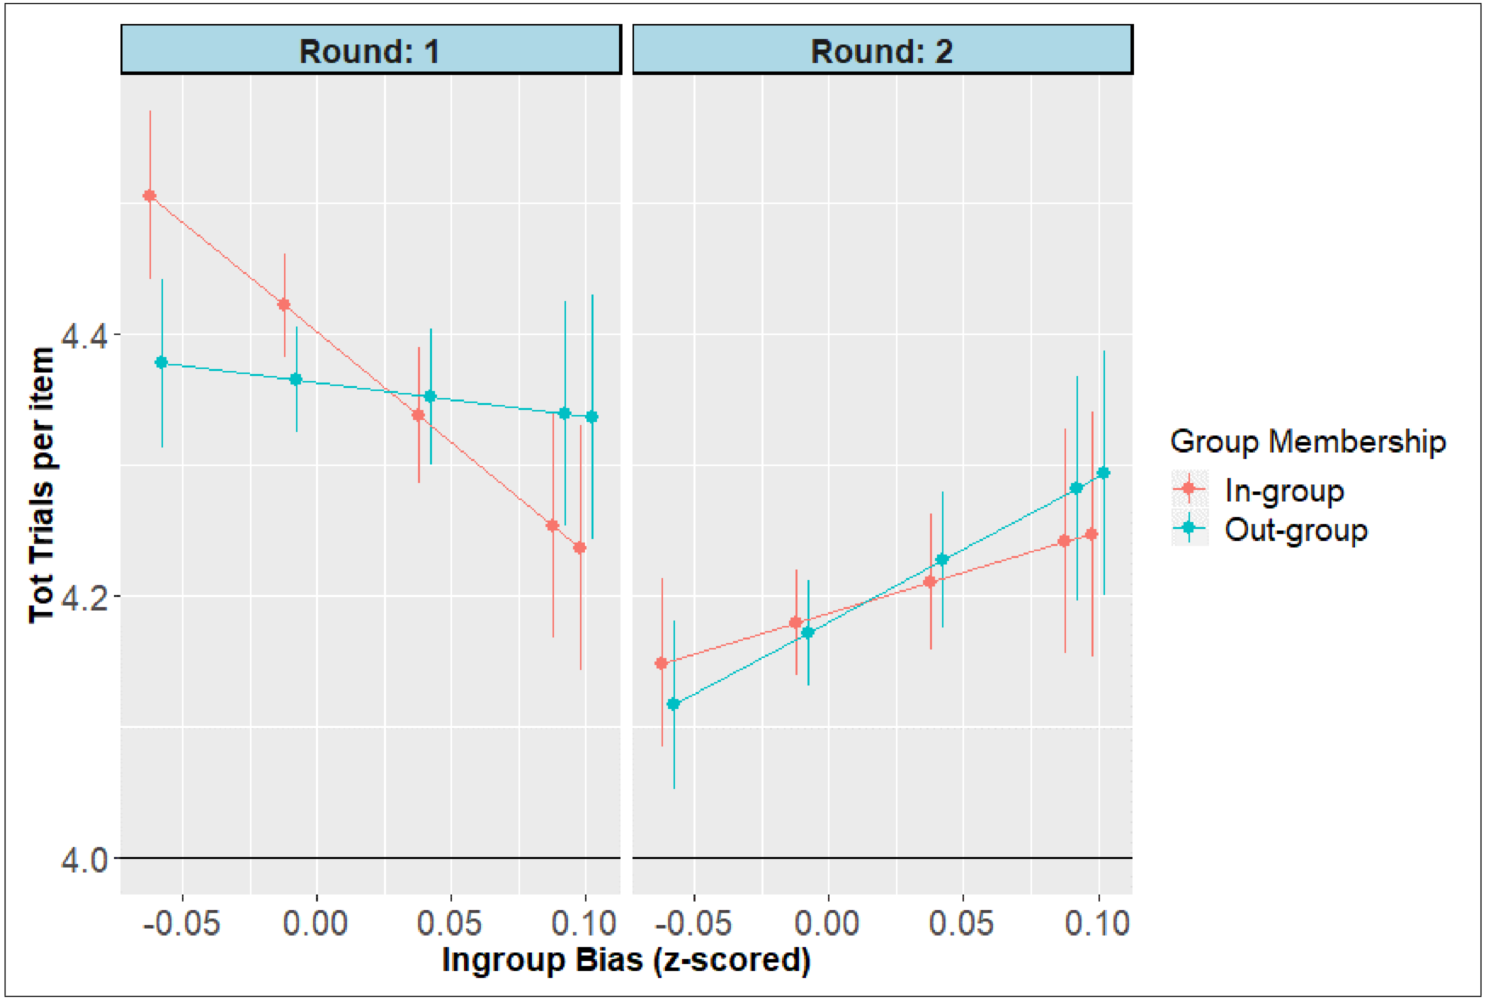

Figure 4.15: Effects of Group Membership and In-group Bias (z-scored) on total number of training trials in the two rounds of exposure of the word learning tasks. Error bars represent standard error.

values) and response bias (indexed by C values) during decision making differed for in-group vs out-group related decision[3].

We calculated two d-prime values and two C values per participant, separately for in-group and out-group trials, by considering participants' performance in matching trials (i.e., hit rates) and mismatching trials (i.e., false alarm rates). We ran two linear mixed-effect models with d-prime and C values as the dependent variable respectively and Group Membership (in-group vs out-group), In-group Bias (centered) and their interaction as fixed effects. The models included per-participant random intercepts.

The model that explored the relationship between individual d-prime and the independent variables showed a significant main effect of In-group Bias (beta=5.75, SE= 2.6, t=2.21, $p<.05$), suggesting that the stronger the in-group bias, the higher the sensitivity (see Figure 4.16). Neither Group Membership nor its interaction with In-group Bias were significant predictors ($ps>.32$), suggesting that sensitivity was equal for in-group and out-group labels.

The model exploring C values did not lead to any significant results ($ps>.11$) suggesting that participants did not differ in decision criterion as a function of Group Membership or In-group Bias. This means that participants' tendencies to produce positive or negative responses were not affected by their social biases.

[3]We also performed accuracy analyses separately for matching and mismatching trials, which led to similar results.

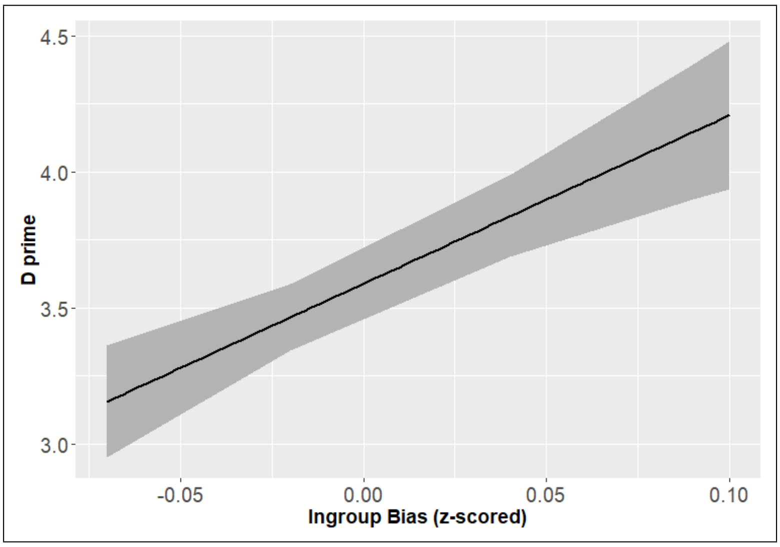

Figure 4.16: Effects of In-group bias (z-scored) on D prime values based on the performance in the source memory task. Shading represents standard error.

Gadget-label memory

First, we checked whether participants learned gadget-label associations well. Overall, participants' accuracy was 89.56% (SD=30.06) and above chance level, as confirmed by a one-sample t-test (i.e., 50%) (t=118.51, p<.0001).

Second, to test whether participants memory for in-group and out-group labels differed as a function of their own in-group bias, we modeled response accuracy in the gadget-label memory test as predicted by Group Membership (in-group vs out-group), In-group Bias and their interaction as fixed effects. The models included per-participant and per-item random intercepts and by-participant random slopes for Group Membership. Results show that none of our predictors significantly affected performance (ps>.31).

Relationships between N400 effect and behavioral measures

So far, we established that presenting non-trained vs trained labels affects N400 amplitudes at the group-level and that this effect was not modulated by speaker group membership. Secondly, we showed that speaker group membership and learners' individual in-group bias mildly affected learning behavior (only in Round 1) and their source memory of novel words. In the next sections, we report analyses in which we related the magnitude of the N400 to a) participants' in-group bias, b) participants' performance in the source memory task and c) performance in the label memory task.

Relationship between N400 effect and In-group bias

We assessed whether individual differences in in-group bias affected the magnitude of the N400 effect for in-group and out-group labels. We expected the in-group bias to positively relates to the magnitude of the N400 effect for in-group labels, whereas we expected it to negatively relates to the effect for out-group labels. To test this, we first calculated per-participant Cohen's D values related to N400 effects (non-trained vs trained) for in-group and out-group labels, separately. Then, in a linear mixed-effect model, we entered scaled N400 effects as the dependent variable, predicted by In-group Bias (centered), Group Membership (in-group vs out-group) and their interaction. A per-participant random intercept was added as well. Results showed no significant effects ($ts<1.36$, $ps>.05$; see Figure 4.17), suggesting that participants' in-group bias and speakers' group membership did not affect the magnitude of the N400 effect.

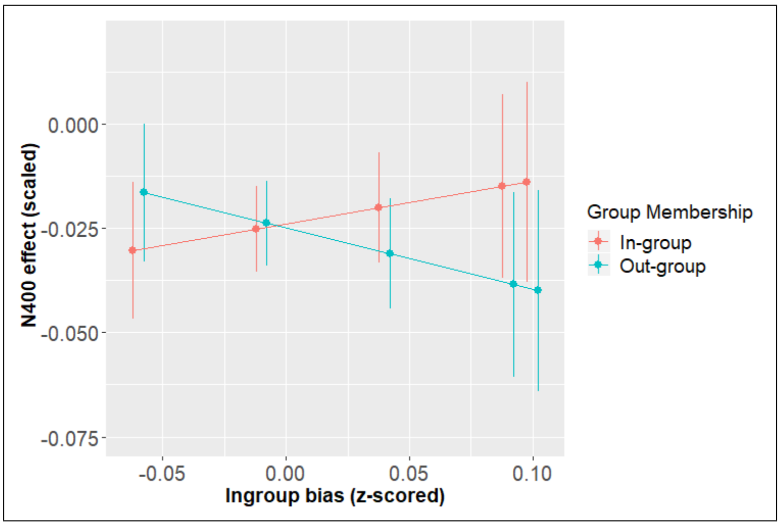

Figure 4.17: N400 effect (scaled) predicted by Group Membership and In-group bias (z-scored). Error bars represent standard errors. The more negative the values, the larger the N400 effect.

Relationship between N400 effects and Source memory

We assessed whether individual differences in in-group bias affected how well people remembered the source memory of the labels learned from in-group and out-group speakers. We expected the in-group bias to positively predict participants' performance in the source memory test. Firstly, we obtained the individual measure related to participants' source memory (hereinafter referred to as Source Memory Performance) by calculating per-participant d-prime values from the source memory test for in-group and out-group labels separately.

In a linear mixed-effect model, we entered scaled N400 effects as the dependent variable, predicted by Source Memory Performance (centered), In-group Bias (centered), Group Membership (in-group vs out-group) and their interactions. A per-participant random intercept was added. Results showed no significant effects ($ps>.05$), suggesting that participants' source memory did not predict the magnitude of the N400 effect.

Relationship between N400 effects and Gadget-label memory

We assessed whether individual differences in in-group bias affected how well people remembered trained gadget-label pairs learned from in-group and out-group speakers. We expected the in-group bias to positively predict participants' performance for in-group pairs and to negatively predict the performance for out-group pairs. Firstly, we obtained the individual measure of label memory (Label Accuracy, hereinafter) by calculating mean accuracy in the gadget-label test per group membership.

Then, in a linear mixed-effect model, we entered scaled N400 effects as the dependent variable, predicted by Label Accuracy, In-group Bias (centered), Group Membership (in-group vs out-group) and their interaction. A per-participant random intercept was added as well. Results showed a significant main effect of Label Accuracy (beta=0.11, SE=0.04, t=2.55, p<.05). This finding suggests that those participants who exhibited the larger N400 effects performed the worst when asked to recognize the learned labels in the gadget-label memory test (see Figure 4.18). None of the other effects were significant (ps>.45).

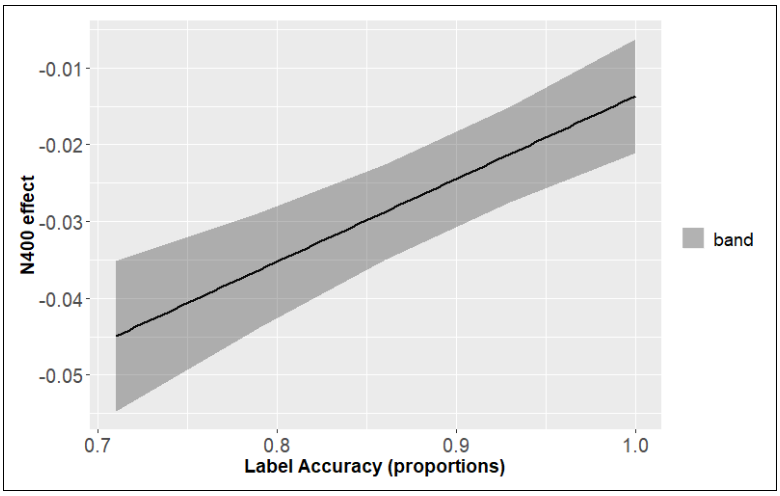

Figure 4.18: Effect of Label Accuracy on N400 effect (scaled). Shades represent standard errors.

Discussion

Previous evidence showed that speaker social identity is not filtered out when processing language, as abstractionist models assume (e.g., Gaskell & Marslen-Wilson, 1997; McClelland & Elman, 1986; Heyes, 1994), but can influence listeners' speech perception and linguistic predictions (e.g., Hay, Nolan, & Drager, 2006; Hay, Warren, & Drager, 2006; K. Johnson et al., 1999; Niedzielski, 1999; van Berkum et al., 2008).

In previous studies, we investigated whether learning new words from in-group vs. out-group speakers would influence the specificity of speaker-related information encoded in the novel representations. The results confirmed our predictions and demonstrated that not only is speaker group membership status encoded, but it interacts with learners' own in-group bias in determining how accurately they later remember the source of the information (i.e., speaker-specific details).

These findings gave rise to a new set of hypotheses that words learned by in-group speakers may be retrieved more effectively and therefore benefit from a processing benefit compared to words learned from out-group speakers. Such tendency would again be influenced by learners' own in-group bias. The current study was designed to test these predictions. We employed a word learning paradigm where participants learned a series of labels for uncommon gadgets. Crucially, half of the gadgets were labelled by in-group speakers and the other half by out-group speakers. On the following day, participants' EEG signal was recorded while they heard target labels spoken by the same speakers in relation to the gadgets. The labels differed in whether they were trained (i.e., matched the labels learned on the previous day) or non-trained (i.e., they were new labels, never heard before).

Results showed more positive N400 amplitudes for trained than for non-trained labels, confirming that the former were easier to process and integrate with the context than the latter. However, there were no differences between words learned from in-group speakers and words learned from out-group speakers since the N400 effects were of comparable magnitudes. Furthermore, learners' individual in-group bias did not modulate these patterns.

How can our results be reconciled with previous findings?

Previous evidence showed that in-group membership enhances information processing and associative learning (e.g., Enock et al., 2018; Moradi et al., 2015) and benefit from the allocation of additional cognitive resources (Meissner et al., 2005; Van Bavel & Cunningham, 2012). Furthermore, speaker group membership in tandem with individual in-group biases influences a) the accuracy of source memory of novel words (Iacozza et al., 2019b) and b) response bias during decision-making processes associated with information retrieval (Iacozza et al., 2019a). Therefore, we hypothesized that, due to the social salience ascribed to in-group speakers, in-group labels would have benefitted of a processing boost, compared to out-group labels. In the context of our design, this facilitation was expected to result in a larger N400 effect (non-trained vs trained labels) for in-group than for out-group words. Results showed that this was not the case and N400 effect for non-trained vs trained label was not modulated by speaker group membership.

Most likely, the label identity x group membership interaction on the N400 amplitudes may have been absent because of the extensive learning session participants went through that might have produced strong speaker-label associations – regardless of speaker social identity. As both in-group and out-group speaker-label associations were learned to ceiling in our learning session, initial group membership differences, if present at all, were eliminated by extensive training.

To test this hypothesis, we analyzed participant's performance in Task 3 of the learning session on day 1. In this task, participants learned to associate specific gadget-label pairs with either the in-group or the out-group speakers. Thus, we asked whether participants learned associations of labels and in-group speakers faster than associations of labels and out-group speakers. As reported above (Results section, Behavioral data), this turned out not to be the case: the analyses of the number of trials needed to learn the associations to criterion was not different for in-group vs. out-group speakers.

Nevertheless, we performed split-data exploratory analyses and looked at the two rounds of exposure separately to gain deeper understanding of the data. We found mild evidence that, in the first round, participants' behavior was influenced by a significant interaction between speaker group membership and their own in-group bias. Participants with stronger in-group bias showed a stronger tendency to commit fewer errors with in-group than with out-group associations. This was not the case in the second round, suggesting that initial differences driven by group membership, although mild, had been reduced with training.

While we can conclude that there were no strong differences in the learning trajectory of in-group and out-group labels, we cannot rule out the possibility that, by training both in-group and out-group speaker-label associations to ceiling on day 1, the likelihood of finding differences on the N400 on day 2 decreased. This explanation cannot be directly tested in the current study.

We designed the learning tasks of the current study with the intention of optimizing learning of both in-group and out-group words during the initial encoding of words. The reason was our interest in determining how group membership would affect the *processing* of newly learned words in later encounters.

We wanted to avoid differences in memory of the newly learned words, i.e., that in-group words would be learned better, and remembered more accurately, and that this could lead to differences in processing on the second day, when retrieving those words. Our results suggest that when speaker-label associations are learned to a similar extent, there are no differences in how such information is later processed.

To be able to test the hypothesis that training at ceiling speaker-label associations decreased the likelihood of finding differences in processing costs between in-group and out-group words, future studies could turn to a between-participant design where participants receive either similar training to the one received in the current study or are not forced to learn to an equal criterion and investigate whether differences in the amount of training affects the N400 effect and its relationship with speaker group membership and individual in-group bias in a similar EEG design.

On a different note, we investigated the resilience/stability of the representations of the trained labels after interfered by competing labels in the memory test that followed the EEG task. We explored the relationship between participants' individual accuracy rate in the gadget-label memory test and the N400 effect, as a measure of competitor interference. We found that the two measures were significantly qualified by a negative relationship. The larger the N400 effect (non-trained vs trained) experienced by individual participants, the worse they later performed in the memory test. Such relationship fits with the assumptions of the *Predictive Interactive Multiple Memory Systems* account (Henson & Gagnepain, 2010). This account states that feedback and feed-forward activity between top-down predictions (in our case, driven by learning of speaker-label associations) and bottom-up input interact to minimize prediction error so that the larger the prediction error experienced, the stronger the memory encoding (e.g., Greve, Cooper, Kaula, Anderson, & Henson, 2017).

Our findings suggest that the degree of difficulty experienced when processing unexpected, non-trained labels might have affected the resilience/stability of the representations of the newly learned labels. In other words, the more noticeable competitors were to a participant, the more they interfered with the learned labels in the subsequent memory test.

In this study, we did not find differences in how in-group and out-group labels were affected by prediction error, nor was this relation modulated by individual in-group bias. Nevertheless, it would be interesting to test the effects of prediction error on the resilience of novel lexical representations learned by in-group and out-group speakers by manipulating the priors of specific predictions, i.e. the constraint of the context. For instance, in the EEG session of the current study we did not introduce novel speakers or novel pictorial exemplars of the studied gadgets. However, one could hypothesize that if in-group words are indeed associated more strongly with speakers and pictorial referents, they could be less easily generalized to novel speakers and items. It would be interesting to see whether introducing novel elements in the task keeping the sensory evidence constant (the unexpected word) would affect participants' electrophysiological and mnemonic responses in a different way for in-group vs out-group words.

Finally, we looked at source memory of novel words and found that it was influenced by learners' in-group bias, but not by the group membership status of the speakers. The stronger the in-group bias, the more accurate participants' source memory of novel words, irrespective of the group membership of the speakers. The lack of modulation by group membership is not surprising given that labels were fully associated with a group and not just an individual, since both members of each group produced a label (contrary to the previous studies). Therefore, remembering the source of in-group labels corresponded to remembering the source of out-group labels. Indeed, any label not produced by the in-group members was necessarily produced by the out-group speakers. The significant main effect of in-group bias can be interpreted as reflecting the greater motivation for participants with strong in-group bias to avoid incorrectly attributing in-group labels to out-group speakers and vice versa (cft. Iacozza et al., 2019a, for effects of group membership and individual in-group bias on response bias).

To sum up, the results reported in the current study suggest that, while the specificity of speaker-related information encoded in novel lexical representations can be affected by group membership status and learners' in-group bias, as shown in previous studies (in Iacozza et al., x), on-line word processing may be less sensitive to speaker group membership manipulation, especially when differences during learning are intentionally minimized.

In light of previous evidence, our results could be interpreted in two different ways. On the one hand, we hypothesized that differences in the encoding of novel words learned from in-group and out-group speakers, reflected in the strength or level of detail of a memory representation, would lead to differences in the on-line processing of such words. The lack of modulation of group membership status on N400 effects reported in this study could be a consequence of the design we employed, in which we intentionally minimized initial differences during learning and strength of representation by providing extensive training. In everyday life listeners are not forced to learn at ceiling both in-group and out-group linguistic variants. On the contrary, in everyday life, where interactions with in-group members (e.g., friends, colleagues, relatives, partners) are arguably more frequent than interactions with out-group members, in-group bias effects may be stronger than in the current experiment.

On the other hand, the lack of these effects may result exactly from real-life experience, where people have to understand the language of speakers perceived as out-group members and still comprehend what they say. If processing differences exist when comprehending in-group and out-group language, they may not relate to how easily people integrate the words of in-group and out-group speakers into the given context, but maybe in how trustworthy or knowledgeable they believe those speakers to be, and/or in how deeply they process the meaning of what is said. Along this line, Mangardich and Sabbagh(2018) have recently showed that six-year-old children can equally recognize labels learned by knowledgeable and ignorant speakers, but the semantic processing of those labels differ. Therefore, if out-group speakers are perceived as less knowledgeable or trustworthy (e.g., Slessor, Phillips, Ruffman, Bailey, & Insch, 2014; Williams, 2001), people could process the meaning of what they say more shallowly and less correctly (see Lev-Ari & Keysar, 2012, for shallower processing of information delivered by foreign-accented vs native speakers), confirming, in turn, false believes and feeding back existing implicit in-group biases.

Here, we wanted to isolate processing-related effects so we minimized initial memory differences. Also, we used a novel word learning paradigm to control for previous exposure to the stimuli and avoid interference between our manipulations and world-knowledge top-down effects. What would have happened if we had used existing words that participants had already associated to speakers with a given social identity, like in the case of the word *sandcastles* produced by a child speaker? Or if the speakers we employed had had a particular accent to which participants ascribed a social weight, influencing their degree of belonging to an in- or out-group? The aim of this study was to isolate processing differences between words that solely differed for having been learned from in-group or out-group speakers, controlling for potentially other interfering factors (both social and otherwise). All the questions raised above invite further research to gain deeper insight into the exact influence of social factors during language processing.

5 Lack of in-group biases on orientation of spatial attention in a Posner cueing task

Abstract

In a very rich visual environment, some stimuli are perceived as highly salient and therefore preferentially attract attention. Recently, it has been proposed that information related to others conceived as in-group members might be perceived as highly salient and bias people's attention. Nevertheless, the existing evidence in favor of a larger attentional capture by in-group stimuli vs out-group stimuli appeared to be mixed. In two studies, we used a Posner cueing paradigm to test whether people's spatial attention was more effectively directed by in-group vs out-group faces, which were used as cues for a target to be classified for its orientation. Results showed that participants were generally faster at target classification when a facial cue correctly anticipated the location of the target than when it did not. However, the group membership status of the faces did not modulate the magnitude of this facilitatory effect. Our results shed light on potential boundaries of the in-group bias in relation to spatial attention. Theoretical and methodological implications of the findings are discussed.

Introduction

Our visual environment is extremely rich and therefore the ability to orient attention towards stimuli that are relevant for us helps to navigate the surroundings without reaching informational overload. A particular type of information that has been proposed to receive high attentional priority is information related both to the self and to others conceived as in-group members (e.g., Sui & Humphreys, 2015; Symons & Johnson, 1997). Nevertheless, empirical evidence supporting the hypothesis that in-group information attracts more attention than out-group information appears to be mixed. We carried out two studies in which we used a Posner cueing paradigm (Posner, 1980) to test whether in-group cues indeed attract attentional resources to a higher degree than out-group cues, and whether such extra attention would boost or hinder individuals' performance on the categorization of cued targets.

Before turning to the current studies, we review relevant literature. First, we report findings that show that information related to the self and to the in-group membership receives processing and memory benefits. Secondly, we review evidence in support of the hypothesis that high attentional saliency may be the underlying mechanisms of enhanced processing of self-related materials and extend this hypothesis to the case of in-group membership. Finally, we show that evidence for attention mediating in-group biases is actually mixed and not definitive. We end by presenting two studies we designed to test whether the hypothesis of attention being modulated by group membership finds is furtherly supported.

The special treatment given to information related to the self and to in-group members results in well-studied effects, including enhanced perceptual processing and boosted memory performance (e.g., S. J. Cunningham, Brebner, Quinn, & Turk, 2014; S. J. Cunningham, Turk, Macdonald, & Macrae, 2008;

Enock et al., 2018; Hugenberg et al., 2010; Leshikar, Dulas, & Duarte, 2015; Lou et al., 2004; Moradi et al., 2015; Rogers, Kuiper, & Kirker, 1977; Sui & Zhu, 2005; Sui, He, & Humphreys, 2012; Symons & Johnson, 1997; Van Bavel et al., 2008).

For example, memory for self-related material is typically better than for material related to other people (e.g., Conway, 2005; Conway & Pleydell-Pearce, 2000; S. J. Cunningham et al., 2008). Self-relevance has also been shown to strengthen and speed up the learning of arbitrary associations between previously neutral stimuli (e.g., geometrical shapes) and written labels referring to the self (i.e., you), compared to associations not involving the self (Sui et al., 2012). These results are consistent with the hypothesis that stimuli that are related to the self are perceived as highly socially salient and therefore receive enhanced perceptual processing (e.g., Bargh, 1982; Humphreys & Sui, 2016).

Importantly, many social psychological theories on how social identity is formed similarly predict that the salience ascribed to self-related information should extend to individuals that are categorized as in-group members (e.g., Brewer, 1999; Turner & Tajfel, 1986) since the former and the latter share psychological representations ("identity fusion," see Swann, Jetten, Gómez, Whitehouse, & Bastian, 2012).

Indeed, processing and memory advantages elicited by stimuli related to the self were robustly shown also in the case of in-group members (Enock et al., 2018; Frable & Bem, 1985; Hugenberg et al., 2010; Moradi et al., 2015; Van Bavel et al., 2008; Wilder, 1990). For instance, in-group faces are recognized more accurately than out-group faces (e.g., Van Bavel et al., 2008; Hugenberg et al., 2010) and the source of information (i.e., who said what) tends to be remembered better when the information is delivered by in-group vs out-group members (e.g., Frable & Bem, 1985; Greenstein et al., 2016; Wilder, 1990). Furthermore, as with self-related arbitrary associations, people are also faster and more accurate when learning to associate previously neutral stimuli (e.g., geometrical shapes) and their in-group memberships (e.g., the logo of their favorite football club) than when they learn associations involving an out-group membership (Enock et al., 2018; Moradi et al., 2015).

Attention has been proposed as one of the cognitive mechanisms supporting self and in-group biases (Bargh, 1982; Humphreys & Sui, 2016; Van Bavel & Cunningham, 2012). The proposal is that due to the high social salience attributed, stimuli related to the self and to in-group members receive additional attentional allocation that enhances processing and memory of materials. Such a hypothesis is supported by a large body of behavioral, electrophysiological and neuroimaging evidence. We review first evidence for attentional benefit for self-related stimuli.

Several behavioral studies demonstrated that self-related stimuli can act as powerful distractors, harder to ignore than other-related information (e.g., Bargh, 1982; Brédart, Delchambre, & Laureys, 2006; Moray, 1959). Similarly, cues like own first names and self-faces, but also arbitrary stimuli temporarily associated to the self (e.g., color-coded arrows), were reported to direct spatial attention to peripheral targets to a greater extent than stimuli not related to the self (e.g., Alexopoulos, Muller, Ric, & Marendaz, 2012; Liu, He, Rotsthein, & Sui, 2016; Sui, Liu, Wang, & Han, 2009). For instance, Alexopoulus and colleagues (2012) carried out a series of studies to test whether the own-name effect (the attentional bias for one's own first name) relies indeed on the automatic recruitment of attentional resources. They showed that own names used as cues to subsequent targets produced larger cueing effects than other names. This was true when the cues were presented shortly before the target (SOA of 235 ms), but also when they were unconscious, or contrasted with familiar names of known people (yet not friends').

Similarly, in Liu et al. (2016), images of self- or other-faces were used as cues in a Posner cueing task. In this task, faces oriented to either side of the screen, directing participants' attention towards the cued location. Subsequently, with a short time delay, peripheral targets were shown and participants had to classify them based on their orientation. Self-faces were reported to function as stronger cues than other-faces, producing larger cueing effects (i.e., the RTs difference when classifying validly and invalidly cued targets). This was true especially when the inter-stimulus interval (ISI) between the cue and the target was short (i.e., 50 ms vs 150 ms).

Along the same lines, numerous EEG studies reported that self-faces, compared to others-faces, elicited larger electrophysiological responses within the first 250 ms from stimulus onset (e.g., Caharel et al., 2002; Liu et al., 2016; Sui, Zhu, & Han, 2006). Such findings support the hypothesis that self-faces might benefit from boosted perceptual processing and extra attention allocation. Additional support for this interpretation was provided by neuroimaging data collected by Sui and colleagues (2013) who showed that areas previously linked to self-representations and social attention (namely, vmPFC and LpSTS) were more active for self as compared to friend or stranger related stimuli. Importantly, dynamic causal modelling analyses confirmed that neural activation spread from vmPFC to LpSTS in a top-down fashion. These findings indicate that the attentional system was primed by early activation of self-representations to enhance the perceptual processing of self-related stimuli (Sui et al., 2013).

Based on these findings, Humphreys and Sui (2016) put forward a theoretical model, the *Self-Attentional Network*, which formally describes how the brain might support the enhanced processing of self-related information via attention prioritization. The model proposes a three-component neural network where the dorsolateral prefrontal cortex (DLPFC) and the intraparietal sulcus (IPS) would be responsible for recruiting goal-oriented attention to self-relevant stimuli in a top-down fashion, sending excitatory or inhibitory signals to the vmPFC. From the vmPFC, excitatory signals would reach the pSTS so that the processing of self-related stimuli would be perceptually enhanced. Crucially, Humphreys and Sui (2016) suggest that these cognitive mechanisms, and brain circuits, might extend to the processing of information related to our in-group membership as well (see Van Bavel & Cunningham, 2012, for a model on attentional benefit for in-group membership).

While this is a plausible hypothesis that follows naturally from similarities between self and in-group biases on perceptual processing and memory, the evidence supporting the relation between attentional capture and group membership status is mixed. Additionally, it is not clear how this relation is qualified, i.e., whether extra attention is allocated to in-group or to out-group members. In the next section, we review studies attention is allocated to a greater extent to in-group members. In the following section, instead, we present conflicting evidence that suggests that out-group members are those that received extra attention. We discuss potential limitations of the existing studies and explanations for these seemingly contrasting results.

There are cases where results seem to corroborate the view that attention is preferentially oriented to in-group members (e.g., Chauhan et al., 2017; Van Bavel & Cunningham, 2012; Park et al., 2016, Experiment 2). For instance, in Chauhan et al. (2017), participants performed a Posner cueing task in which the eye gaze on friends' or unfamiliar faces acted as cues to subsequent, peripheral targets. Participants had to react to the targets with either saccadic eye movements or manual responses. Overall, participants were slower when cued by friends' faces, independently of the validity of the cues, and additionally showed a smaller cue validity effect. However, this effect was observed only in trials with a short SOA of 100 ms (as compared to a longer SOA of

200 ms), and when participants responded with saccadic eye movements. The authors interpreted their results as indicating that it was easier for individuals to disengage from unfamiliar faces than from friends' faces, so that their responses were faster after the unfamiliar face. However, this effect was observed relatively early and was short-lived, suggesting that when more time between the cue and target is given unfamiliar and friends' faces produce comparable attentional shifts.

Similarly, in a study using a minimal group paradigm where group membership was created by randomly assigning participants to the blue or green team (Park et al., 2016, Experiment 2), participants were worse at detecting a target letter when superimposed on in-group faces than when superimposed on out-group faces. This was true only in a high-load condition, where the target was hard to detect. The authors claimed that in circumstances where greater cognitive control is needed because of the difficulty of the task, task-irrelevant in-group faces were harder to ignore than out-group faces were and therefore disrupted selective attention to targets.

The evidence provided so far is in line with the hypothesis that in-group members could capture individuals' attention to a higher degree than out-group members do. Nevertheless, there exist some studies reporting results that conflict with such a claim where out-group faces were stronger cues than in-group faces (e.g., Brosch & Van Bavel, 2012; Park et al., 2016, Experiment 2). For instance, Brosch and Van Bavel (2012) used a dot probe task in which participants saw two facial stimuli appearing simultaneously one the screen. One of the two stimuli was either an in-group or an out-group face, and the other stimulus was always a neutral face (i.e., someone who had not been previously categorized as either in-group or out-group member). After a short delay, one of the two faces was replaced by a target that participants have to categorize based on its orientation (i.e., vertically or horizontally located), while the other face remained on screen. When the target replaced the neutral face, participants were slower if the remaining face was that of an out-group member as compared to when it was an in-group member. This was only true in trials with a short SOA of 100 ms (vs 500 ms). According to the authors, these findings fit within an emotional attention framework (Brosch, Sander, Pourtois, & Scherer, 2008; W. A. Cunningham, Van Bavel, & Johnsen, 2008) where stimuli that are perceived as emotionally salient would receive more attention. In the case of their study, they claimed that out-group faces were considered as threatening because of the evolutionary contrast existing between in-group and out-group membership and therefore attracted attention to a larger degree than in-group faces did (Brosch & Van Bavel, 2012).

Considered these reported mixed findings, it is unclear whether attention is indeed modulated by group membership and whether in-group faces exert high attentional saliency, compared to out-group faces. While self-related stimuli, compared to neutral ones, are arguably given higher saliency and positive valence, due of the special status ascribed to the self, information related to in-group and out-group members could be seen as equally salient, just with a difference in valence (in-group seen as positive vs out-group seen as negative).

Note that stimuli and tasks used greatly differed across the mentioned studies, so non-convergent findings may be due to differences in stimuli and tasks. Moreover, one of the biggest limitations of the previous studies is that, since they did not have a neutral cue condition as a reference level, they were not able to tease apart facilitatory effects of valid cues from inhibitory effects of invalid cues. When valid and invalid cues are directly compared (i.e., without reference to a neutral baseline), both facilitation and inhibition contribute to the observed effects (Neely, 1991). Therefore, the mixed results observed in the literature might as well reflect confusion between facilitation and inhibition, and make it hard to understand the nature of the findings.

Overview of the current studies

To shed further light on this issue, we designed two studies where we used a Posner cueing task to test for potential effects of group membership on attentional capture. We used the same paradigm as in Liu et al. (2016) to be able to make clearer comparisons between self and in-group biases, and opted for similar stimuli: dynamically orienting faces of in-group or out-group members that acted as cues to subsequent, peripheral targets that needed to be categorized for orientation (i.e., being upright or inverted). To disentangle facilitatory and inhibitory components, we added a neutral condition, i.e., a face that did not cue any location. Furthermore, following from Liu et al. (2016), two inter-stimulus-interval (ISI) conditions (short 50 ms ISI vs long 150 ms ISI) were used to clarify the temporal dynamics of group membership effects on attentional shift.

The orientation of the facial cues in relation to the location of the subsequent target was, in an equal number of trials, either a) congruent (i.e., valid-cue trials), b) incongruent (i.e., invalid-cue trials), or c) neutral, when the faces were frontally oriented and did not cue any location (i.e., neutral-cue trials). The design made possible to tease apart potential facilitatory and inhibitory effects by comparing valid-cue trials and invalid-cue trials to neutral ones. We predicted that in-group cues would attract more attention than out-group cues. Such a behavior was expected to be reflected in valid in-group cues (vs neutral cues) facilitating participants' performance, both in response accuracy and response times, compared to valid out-group cues. Similarly, we predicted that invalid in-group cues would hinder performance more than out-group cues would, since in-group faces, in comparison, should be harder to disengage from. Furthermore, since previous studies showed effects of group membership on attention only with a short interval between cue and target, we expected differences between in-group and out-group faces to be larger in the short ISI than in the long ISI.

The two studies differed only in the way participants learned about the faces' group membership (i.e., social learning task) before performing the Posner cueing task: in Experiment 1, they listened to the speakers mentioning their student affiliation, whereas, in Experiment 2, they performed a simple associative task between the student faces and their university logos.

Experiment 1

In Experiment 1, we asked participants first to listen to two female speakers referring to their habits as university students, while seeing their faces displayed. The utterances clearly indicated that one speaker attended the participants' own university (in-group speaker), whereas the other speaker attended a different university in the north of the Netherlands (out-group speaker). Afterwards, participants carried out the Posner cueing task in which the same student faces acted as cues.

Methods

Participants

Fifty-six Dutch female speakers (age range: 20-28, mean=23.22, SD=1.81) participated in the study. Data from twelve participants were excluded because participants reported not belonging to the in-group university in the debriefing phase. One additional participant was tested but excluded because she had difficulties learning during the social learning task. The remaining 43 participants were all students or recent graduates of Radboud University Nijmegen.

Materials

Materials for the social learning task

Speakers: Two fictitious speakers were created by pairing frontally oriented faces, that were taken for the purpose of this study by recruiting two women from the MPI subject pool, with the voices of native Dutch female speakers recorded in our laboratory. Prior to the experiment, voices were matched for perceived typicality and attractiveness (paired d-tests, ps>.05) via a norming on-line survey in which twenty different participants participated. Each speaker was a unique combination of one face and one voice, consistent across participants.

Facts: After a norming procedure (see Appendix 1, for details), we had a selection of twenty-four sentences about student life and recognizable landmarks or actions. Twelve sentences referred to Radboud University and twelve to Groningen University. This was achieved by having twelve unique sentential frames in which key words explicitly related to either Radboud or Groningen University were embedded. All statements were audio-recorded by each of the female voices described above.

Materials for the Posner cueing task

Eleven images of each of the two Dutch women from the previous task were taken at different angle orientations. One photo (used also in the social learning task) depicted frontally orienting faces; the remaining ten photos depicted the faces at angles ranging from 15° to 90°, in each direction, with equal steps. In all images, faces had neutral facial expressions and straight eye gazes.

Procedure

The experimental session included the social learning task and the Posner cueing task always administrated in the same order. The session lasted approximately one hour.

Social learning task

The goal of the social learning task was to familiarize the participants with the student faces and to ensure they knew their group membership status before performing the Posner cueing task. Participants were told the task was part of a joint project of Radboud (i.e., in-group university) and Groningen University (i.e., out-group university). In each trial, participants saw a fixation cross (500 ms) followed by the display of a photo of the speaker (500 ms). Then, a recording of an utterance was played along with the photo. Two exposure rounds were administered. Trials appeared in random order, but each sentence was repeated twice, once per round. Six sentences per speaker were presented during the exposure phase.

Following the exposure, participants' memory was tested. In each trial, the photo of one of the speakers appeared with four written sentences. Participants indicated which of the sentences was uttered by the displayed speaker and received feedback. The four displayed sentences were: the correct response; an incorrect sentence previously produced by the other speaker; and two incorrect sentences that had not been heard in the exposure, which referred to either the same or the other university to which the speaker belonged. Feedback was provided after each selection. Participants were specifically trained to distinguish between the two affiliations, therefore, if they selected a statement referring to the incorrect university, the trial repeated at the end of the task until a response with the correct affiliation was provided. Students' affiliations, and sentences selection, were randomized across participants.

Posner cueing task

The Posner cueing task was the main task of the study and was used to test potential in-group biases on spatial attentional shifts.

In each trial, participants saw a central fixation cross and two peripheral empty boxes (500 ms). The fixation cross was then replaced by one of the faces.

For two thirds of the trials, the face turned from a frontal orientation to either the left or right side in a dynamic way, finishing at 90 degrees. The dynamic movement was achieved by showing five images at different face angles, for 30 ms each (total duration of the perceptual movement 180 ms). In the remaining trials, participants saw a still, frontally oriented face for the total 180 ms. Afterwards, the initial display with the fixation cross and the peripheral empty boxes was shown again, for either 50 or 150 ms. A target was then displayed. It consisted of a letter "T", upright or inverted, embedded in a cross pattern displayed on one of the sides of the display for a time ranging between 800 to 1200 ms (see Figure 5.1 for the trial sequence).

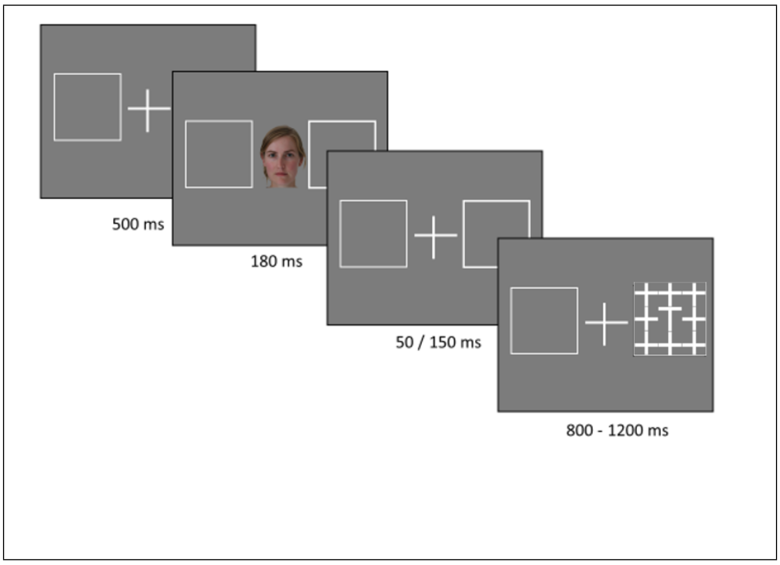

Figure 5.1: Example of the trial sequence for the Neutral-cue condition.

In one third of the trials, the location of the target was correctly predicted by the face orientation (i.e., valid-cue trials), in another third of the trials, it was incorrectly predicted (i.e., invalid-cue trials), and, in the remaining trials, the frontal orientation of the faces did not provide any spatial cue (i.e., neutral-cue trials). Participants categorized the target letter for being either upright or inverted by pressing the corresponding key on a button box. The instruction emphasized both response speed and accuracy. In a randomized presentation, each participant performed 432 experimental trials after 48 practice trials. There were 36 trials per condition.

Results

All analyses were performed with mixed-effects modelling as implemented in the lme4 package (Bates et al., 2014) in R (R Core Team, 2016). The models' random structures were determined following Bates et al. (2015).

Social learning task

Participants were accurate in 99.08% of the trials (SD=2.84), showing that they learned the speakers' affiliations very well.

Posner cueing task

To test if participants' performance was affected by our manipulations, we looked at both response accuracy and RTs.

Firstly, we ran a generalized linear mixed-model with accuracy as the dependent variable. We included fixed effects for Group Membership (Out-group vs In-group; reference level: In-group), Cue Validity (Valid vs Neutral; Invalid vs Neutral; reference level: Neutral) and ISI (150 vs 50; reference level: 50), as well all the possible interactions. Furthermore, we added Trial Number as covariate to control for potential training effects. Per-participant random intercepts, as well by-participant random slopes for ISI and Cue Validity were included.

Results showed a significant inhibitory effect of invalid cues leading to more errors than neutral cues (mean=90.9%, SD=28.2, and mean=92.6%, SD=26.2, respectively; beta= -0.43, SE=0.15, z=-2.96, p=.003), whereas no facilitatory effect of valid cues over neutral cues was found (beta= -0.09, SE=0.16, z=-0.60, p=.55). Importantly the inhibitory effect of invalid cues was marginally modulated by Group Membership (beta=0.33, SE=0.19, z=1.74, p=.08). Even though the interaction was only marginally significant, we followed up by performing simple effect analyses, using the function *emmeans* from the homonymous R package (Lenth, 2018), that confirmed that the size of the effect was larger for in-group (beta=-0.41, SE=0.11, z=-3.86, p=.0003) than for out-group cues (beta=-0.28, SE=0.10, z=-2.66, p=0.02). As Figure 5.2 shows, the difference in effect size was mainly driven by participants being more accurate after in-group neutral cues than after out-group neutral cues, and not by group membership differences in invalid-cue trials. These results corroborated our prediction that in-group cues would be harder to disengage from and would produce a larger inhibitory effect than out-group cues.

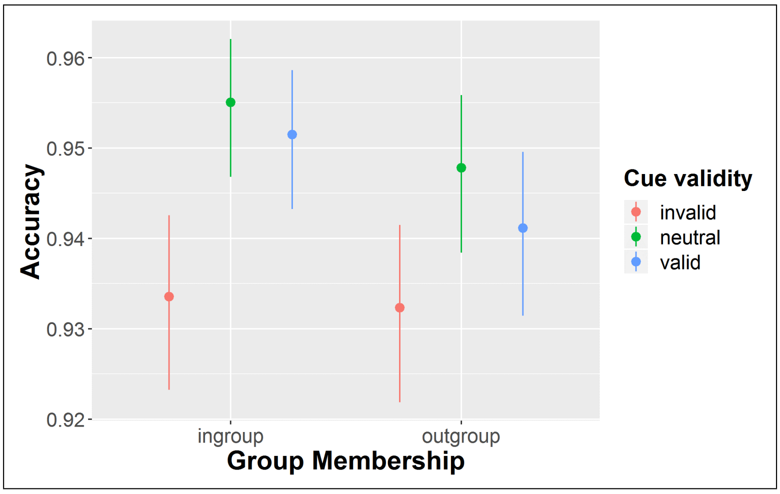

Figure 5.2: The effects of Cue Validity and Group Membership on accuracy in Experiment 1. Error bars represents the standard error.

All the other main effects or interactions, except from Trial Number (beta= 0.0006, SE=0.0002, z=2.96, p<0.05), which suggests participants got better across the experiment, were not significant (see Table 5.1 and Figure 5.3 for details).

Table 5.1: Glmer output for Experiment 1- Accuracy

	Beta	**St.Err**	**z-value**	**p**
Fixed effects				
(Intercept)	2.95	0.21	14.31	<.0001
Group Membership (Out vs Ingroup)	-0.21	0.14	-1.5	.13
Cue Validity1 (Invalid vs Neutral)	-0.43	0.15	-2.96	.003
Cue Validity2 (Valid vs Neutral)	-0.09	0.16	-0.6	.55
ISI (150 vs 50)	-0.07	0.15	-0.44	
.66				
Trial Number	0.0006	0.0002	2.96	.003
Group Membership x Cue Validity1	0.33	0.19	1.74	.08
Group Membership x Cue Validity2	0.07	0.2	0.37	.71
Group Membership x ISI	0.11	0.2	0.56	.58
Cue Validity1 x ISI	0.03	0.19	0.17	.87
Cue Validity2 x ISI	0.03	0.2	0.14	.89
Group Membership x Cue Validity1 x ISI	-0.39	0.27	-1.47	.14
Group Membership x Cue Validity2 x ISI	-0.24	0.28	-0.88	.38

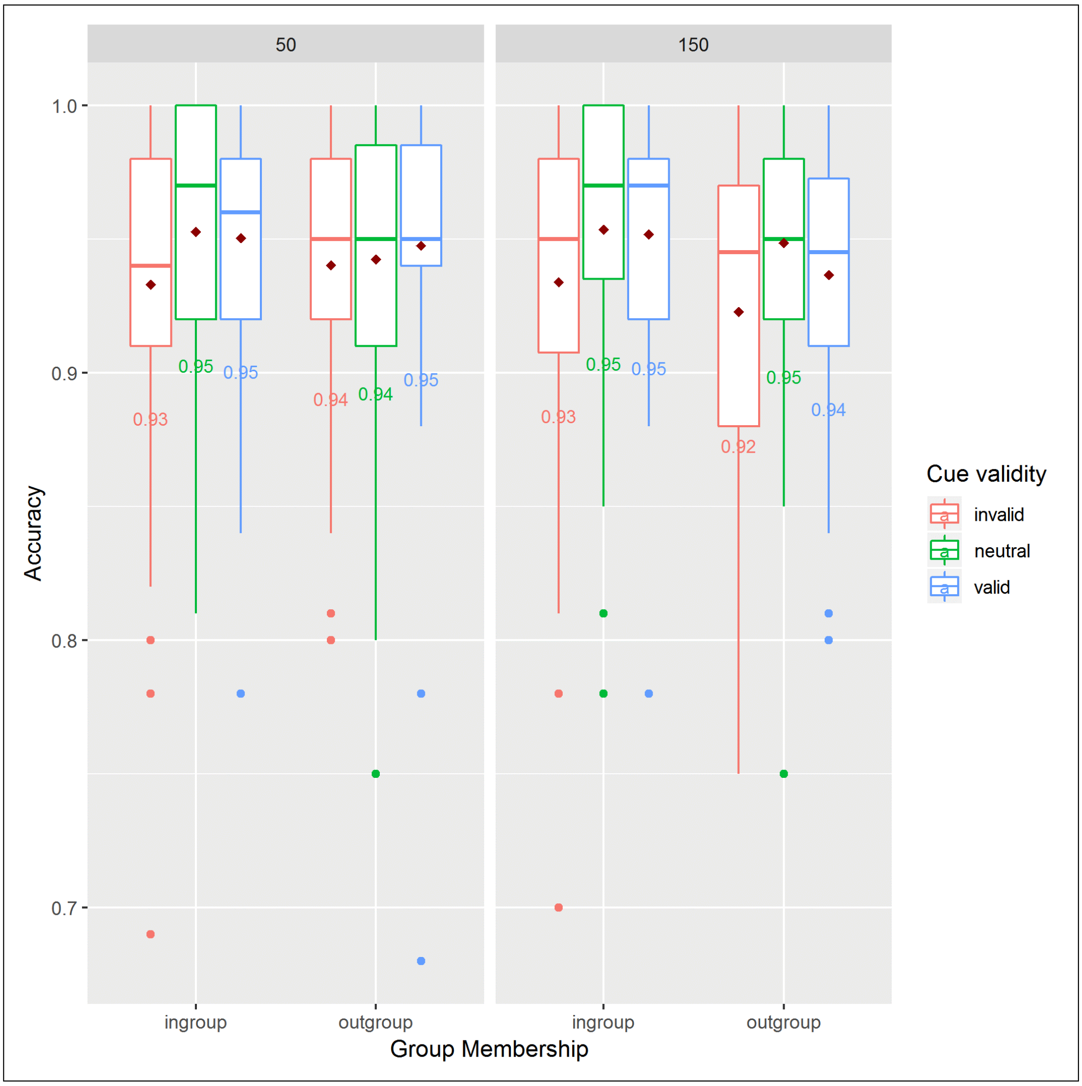

Figure 5.3: Accuracy in Experiment 1 as a function of Cue Validity (color-coded), Group Membership and ISI (facet). Solid lines represent the medians, whereas the red dots represent means, also visualized in digits. Error bars (whiskers) represent upper and lower quartiles of the distribution.

Secondly, we analyzed participants' reaction times with a linear mixed-effect model using the lmer function. First, incorrect trials and trials where RTs were 2.5SD away from the mean were excluded (9.66% of the total trials), RTs were log10-transformed and used as the dependent variable. We included fixed effects for Group Membership (Out-group vs In-group; reference level:

In-group), Cue Validity (Valid vs Neutral; Invalid vs Neutral; reference level: Neutral) and ISI (150 vs 50; reference level:50), as well all the possible interactions. Furthermore, we added Trial Number to control for potential confounding training effects. Per-participant random intercepts, as well by-participant random slopes for ISI and Cue Validity were included.

Results (see Table 5.2 and Fig 5.4 for details) showed a significant facilitatory effect with shorter RTs after Valid cues than after Neutral cues (raw mean=487ms, SD=93.2, and raw mean=514ms, SD=9173, respectively; beta= -0.03; SE=0.003; t=-9.70; p<.0001). The difference between Invalid vs Neutral was not significant (t=0.26; p=.80) suggesting that participants responded equally fast when the face was orienting toward the wrong direction, or remained frontally oriented.

Furthermore, results showed a significant effect for ISI (beta=-0.007, SE= 000.4, t=-2.81, p=.005), with the shorter responses after the ISI of 150ms than after the ISI of 50 ms (raw mean=503 ms, SD=95.9, and raw mean=508 ms, SD=93.7, respectively). No main or interactive effects of Group Membership were significant (ts<0.75, ps>.46), suggesting that participants' RTs were not sensitive to our social manipulation.

Table 5.2: Lmer output for Experiment 1- RTs

	Beta	**St.Err**	**z-value**	**p**
Fixed effects				
(Intercept)	2.71	0.005	545.67	<.0001
Group Membership (Out vs Ingroup)	0.0001	0.002	0.06	.95
Cue Validity1 (Invalid vs Neutral)	0.0007	0.003	0.26	.80
Cue Validity2 (Valid vs Neutral)	-0.03	0.003	-9.70	<.0001
ISI (150 vs 50)	-0.007	0.004	-2.81	.005
Trial Number	0.00004	0.000003	-10.19	<.0001
Group Membership x Cue Validity1	-0.002	0.003	-0.75	.46
Group Membership x Cue Validity2	0.0003	0.003	0.08	.93
Group Membership x ISI	0.001	0.003	0.30	.76
Cue Validity1 x ISI	-0.001	0.003	-0.39	.70
Cue Validity2 x ISI	0.005	0.003	1.68	.09
Group Membership x Cue Validity1 x ISI	0.003	0.005	0.64	.52
Group Membership x Cue Validity2 x ISI	0.001	0.005	0.25	.81

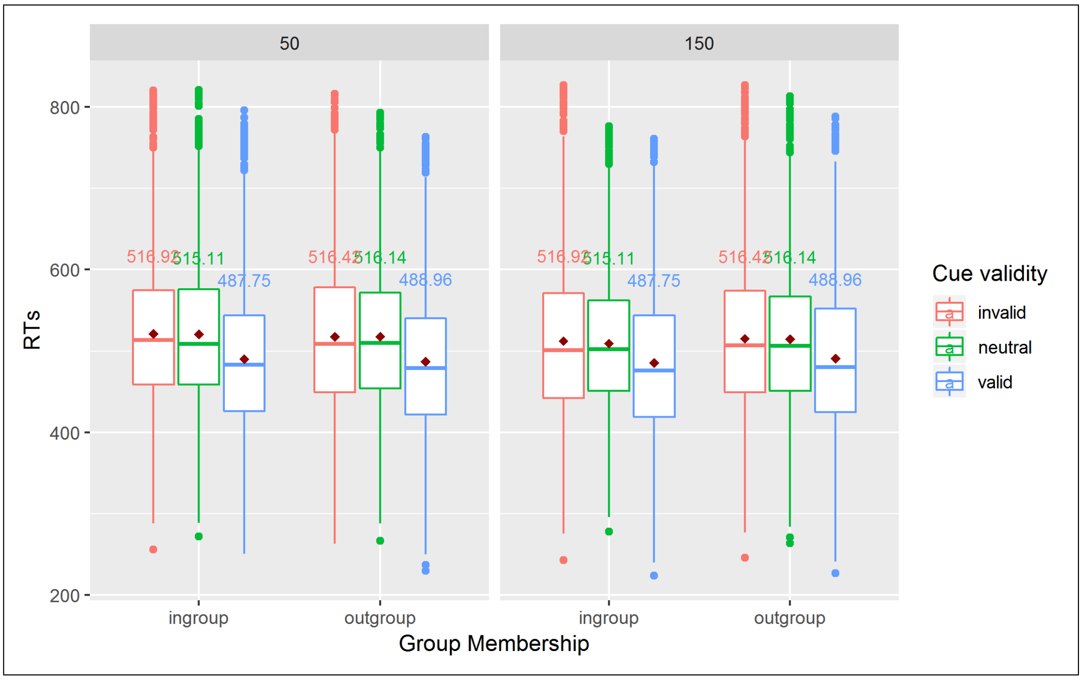

Figure 5.4: RTs in Experiment 1 as a function of Cue Validity (color-coded), Group Membership and ISI (facet). Solid lines represent the medians, whereas the red dots represent means, also visualized in digits. Error bars (whiskers) represent upper and lower quartiles of the distribution.

Interim discussion

To sum up, the results from the accuracy analyses are consistent with previous findings showing that compared to neutral cues, invalid in-group cues were harder to disengage from than invalid out-group cues (e.g., Chauhan et al., 2017). Including a neutral reference level allowed us to tease apart facilitatory from inhibitory effects. Our results revealed that, in this paradigm, inhibitory effects, but not facilitatory ones, might be more susceptible to in-group biases.

On the other hand, the analyses over RTs showed that participants' responses were facilitated when the cue was valid, as compared to neutral cues, indicating a facilitatory effect. However, we did not find an inhibitory effect (participants responded equally fast to both neutral and invalid trials) or modulation due to the Group Membership.

The lack of group membership modulation in RTs replicates a previous study where differences between in-group and out-group members were only found in accuracy rates, but not in reaction times (Park et al., 2016). However, these results conflict with other findings in which group membership also modulated RTs (e.g., Chauhan et al., 2017; Brosch & Van Bavel, 2012). Note that since none of the studies that we mentioned provided a reference level for valid and invalid cues, it is unclear how to reconcile our results with previous, conflicting findings.

Experiment 2

To verify the results in Experiment 1, the second experiment was conducted. Experiment 2 was identical to Experiment 1, except for the initial social learning task. In Experiment 2, participants learned the group membership status of the faces via a simple associative task where they associated each of the faces with the logo of their corresponding university. The rationale behind this procedural change was to maximize in-group biases: previous literature had showed that when participants were motivated to individuate in-group and out-group members in equal measures, they showed weaker in-group biases as compared to when no motivation instructions were given (e.g., Van Bavel & Cunningham, 2012). In Experiment 1, participants learned the two students' group membership by listening to their habits as university students. Their high performance in the memory test of the social learning task (99.08%) indicates that they encoded individual information about both speakers very attentively, potentially leading to equally strong individuation for both speakers. Such a strong individuation of the out-group speaker could have decreased the strength of participants' in-group bias making the group membership effect on participants' RTs harder to detect. Therefore, in Experiment 2, we tried to rule this account out by having a task that required participants to simply categorize the student faces as either in-group or out-group members. A third face, which was always associated to a third university, was also added. This was done to increase the number of stimuli participants had to process, and maximize the chance to activate in-group biases to favor the in-group student over the others.

Methods

Participants

Sixty-one Dutch native speakers (age range: 19-26, mean=22.58, SD=1.7, 37 female) participated in the study. All participants were students or recent graduates of Radboud University Nijmegen.

Materials

Materials for the social learning task

The same frontally orienting faces as in Experiment 1 were used. An additional frontally oriented face was selected from the Chicago Face Database (Ma et al., 2015) to slightly increase the difficulty of the task.

Black and white images of the logos of Radboud (in-group) Groningen (out-group), and Tilburg (neutral) universities were used.

Materials for the Posner cueing task

Same materials as in Experiment 1 were used.

Procedure

The experimental session included the social learning task and the Posner cueing task always in the same order and lasted approximately 45 mins.

Social learning task

In order to have participants learn who belonged to which university, they performed a perceptual matching task where they learned to associate the faces with their affiliations. In an initial exposure phase, the three frontally oriented faces were individually presented with the logo of either the in-group university –Radboud university- or the out-group university- Groningen University, or a third university-Tilburg University. The association involving the third university was included to increase the difficulty of the task. Each association was presented ten times. Then, participant performed a practice block (24 trials), followed by two blocks of 120 experimental trials each where associations either match or mismatch what they learned in the exposure phase. In 50% of the trials, the faces appeared with the respective logo, whereas in the remaining trials, faces were shown with the incorrect logos. In both practice and test trials, a fixation cross (500 ms) followed by a blank interval (1000-2000 ms) and the simultaneous presentation of logo and face (600 ms). Participants had 1500 ms to judge the accuracy of the pairings. Feedback was given only during practice.

Posner cueing task

The same procedure from Experiment 1 was followed.

Results

Social learning task

Participants were accurate in 95.06% of the matching trials (SD=21.66), showing that they successfully learned speakers' affiliations.

Posner cueing task

Firstly, we ran a logistic mixed-effect model with accuracy as the dependent variable and Group Membership (Out-group vs In-group; reference level: In-group), Cue Validity (Valid vs Neutral; Invalid vs Neutral; reference level: Neutral), ISI (50 vs 150, reference level: 50) and their interactions as fixed effects. This model, including the interactions of interest failed to converge, possibly due to ceiling effects. Participants in this experiment were numerically more accurate than the participants in the previous experiment (Experiment 1: mean=94.17%; Experiment 2: mean= 95.35%). As next step, we decided to run a model for each ISI condition separately. Unfortunately, both models failed to converge. A model without Group Membership x Cue Validity interaction is not informative. Therefore, we move to analyses on RTs and plot the raw data in Figure 5.5 for visual inspection.

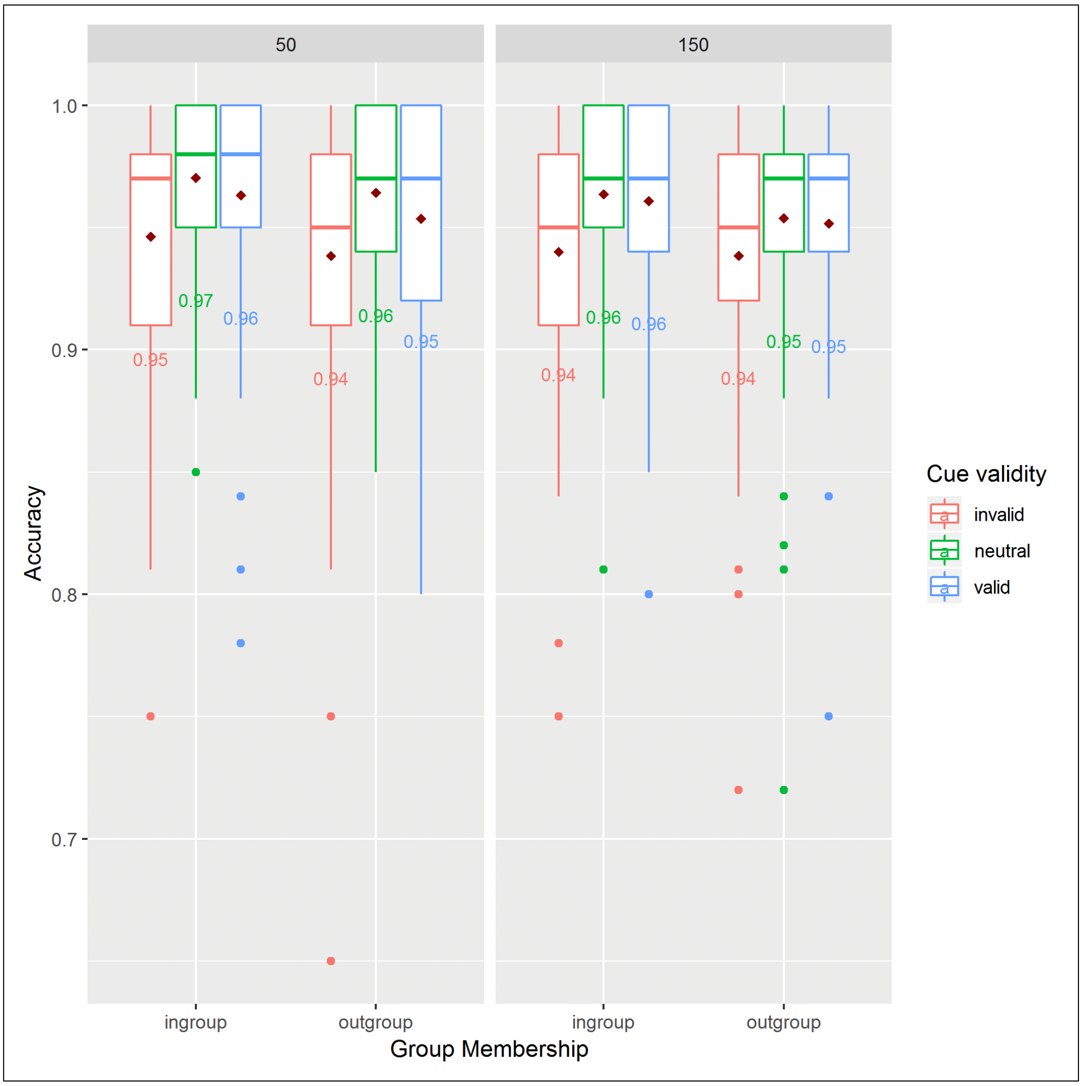

Figure 5.5: Accuracy in Experiment 2 as a function of Cue Validity (color-coded), Group Membership and ISI (facet). Solid lines represent the medians, whereas the red dots represent means, also visualized in digits. Error bars (whiskers) represent upper and lower quartiles of the distribution.

Incorrect trials and trials where RTs were 2.5SD away from the mean were excluded (14% of the total trials), RTs were log10-transformed and used as the dependent variable in a linear mixed-effect model. We included fixed effects for Group Membership (Out-group vs In-group; reference level: In-group), Cue Validity (Valid vs Neutral; Invalid vs Neutral; reference level: Neutral) and

ISI (50 vs 150, reference level: 50), as well all the possible interactions. Furthermore, we added Trial Number to control for potential confounding training effects. Per-participant random intercepts, as well by-participant random slopes for ISI and Cue Validity were included.

Results (see Table 5.3 and Figure 5.6 for details) showed a significant facilitatory effect with shorter RTs after Valid cues than after Neutral cues (raw mean=486 ms, SD=92.7, and raw mean=513 ms, SD=91.3, respectively; beta= -0.02.; SE=0.003; t=-8.68; p<.0001). The difference between Invalid vs Neutral was not significant (t=0.61; p=.54) suggesting that participants responded equally fast when the face was orienting toward the wrong direction, or remained frontally oriented. Trial Number was also a significant predictor, suggesting a practice effect (beta=-0.0007, SE=0.000004, t=-18.43, p<.0001). Furthermore, results showed a significant effect for ISI (beta=-0.005, SE=0.002, t=-2.08, p=.04), with the shorter responses after the ISI of 150 ms than after the ISI of 50 ms (raw mean=502, SD=95.6, and raw mean=507, SD=93.5, respectively). Results showed a marginal significant interaction between ISI and the inhibitory effect (beta=-0.006, SE=0.003, t=-1.9, p=.06), with simple effects analyses suggesting that the inhibitory effect was non-existent in the short ISI (z=-0.03, p=.99) but became marginally significant in the long ISI (z=2.16, p=.08).

Results showed a marginal interaction between Group Membership and the facilitatory effect (beta=-0.005, SE=0.003, t=-1.73, p=.08). While follow-up analyses showed non- significant simple effects (p>.05), the size of the facilitatory effect was slightly larger for out-group cues than for in-group cues (beta=0.029, SE=0.002 vs beta=0.025, SE=0.002, respectively).

To sum up, the results clearly showed that participants' responses were facilitated when the cue was valid, as compared to neutral cues, indicating a facilitation effect. However, we found neither a significant inhibitory effect (participants responded equally faster to both neutral and invalid trials), nor a significant modulation due to the Group Membership.

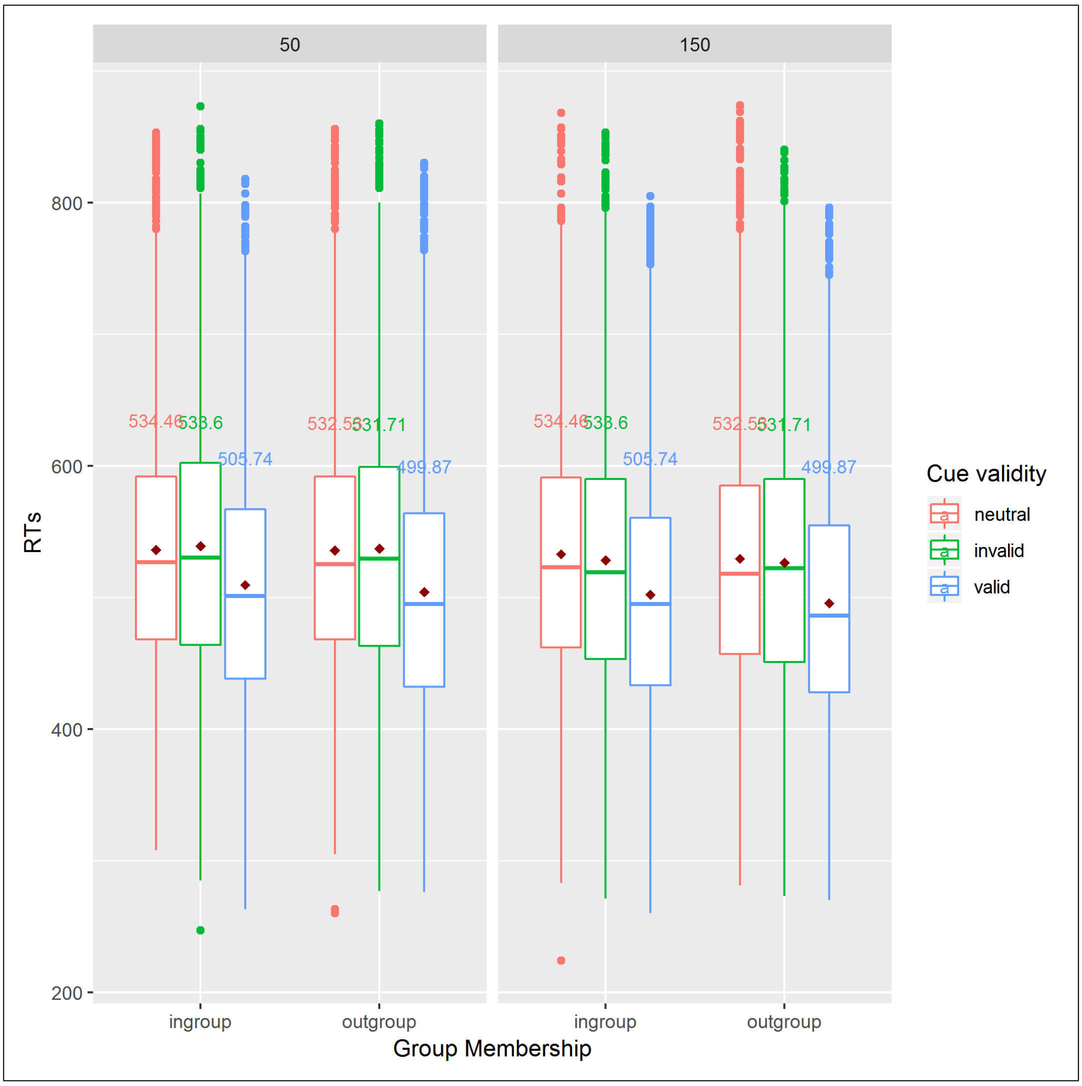

Figure 5.6: RTs in Experiment 2 as a function of Cue Validity (color-coded), Group Membership and ISI (facet). Solid lines represent the medians, whereas the red dots represent means, also visualized in digits. Error bars (whiskers) represent upper and lower quartiles of the distribution.

Table 5.3: Lmer output for Experiment 2- RTs

Fixed effects	**Beta**	**St.Err**	**z-value**	**p**
(Intercept)	2.74	0.006	433.07	<.0001
Group Membership (Out vs Ingroup)	-0.0005	0.002	-0.27	.79
Cue Validity1 (Invalid vs Neutral)	0.002	0.002	0.61	.54
Cue Validity2 (Valid vs Neutral)	-0.02	0.003	-8.68	<.0001
ISI (150 vs 50)	-0.005	0.002	-2.08	.04
Trial Number	-0.00007	0.000004	-18.43	<.0001
Group Membership x Cue Validity1	-0.003	0.003	-0.75	.35
Group Membership x Cue Validity2	-0.005	0.003	-1.73	.08
Group Membership x ISI	0.001	0.003	-0.51	.61
Cue Validity1 x ISI	-0.001	0.003	-1.90	.06
Cue Validity2 x ISI	0.005	0.003	-0.81	.42
Group Membership x Cue Validity1 x ISI	0.003	0.004	0.74	.46
Group Membership x Cue Validity2 x ISI	0.001	0.004	0.54	.59

General discussion

We designed two studies to test whether people's spatial orientation in a Posner cueing task was affected by the group membership status of facial cues. The two studies differed only in the social learning task preceding the Posner cueing task: while in Experiment 1 participants learned students' academic affiliations by listening to their habits as university students, in Experiment 2, they performed a simple associative task between the students and their university logos.

Overall, we predicted that, in both accuracy and RTs, facial cues correctly anticipating the location of the target would facilitate its categorization, whereas faces orienting to the incongruent location would hinder it. In Experiment 1, across both Group Membership conditions, we found that participants were significantly worse at classifying targets when cued by a face orienting to an incongruent location (vs neutral face), but no statistical differences were found in terms of accuracy between valid and neutral cues. Although in Experiment 2 converge issues prevented us from running statistical analyses, the patterns visualized in Figure 5.5 suggest that this tendency held in both experiments. The patterns suggest that participants' spatial attention was indeed driven by the orientation of the cues so that when the cue orientation was incongruent, participants produced a larger number of incorrect automatic responses, as compared to the other conditions.

In terms of RTs, instead, the results revealed that participants were overall significantly faster after a cue that correctly predicted the target location, and no differences were found between invalid and neutral cues. The differential effect of Cue Validity observed in accuracy and reaction times seems to reflect a dissociation in behavior that is explained by the fact that while the accuracy analyses were performed over the whole dataset, the RTs analyses referred only to a subset of it, namely, the correct trials.

Crucially, we expected in-group cues to produce greater facilitatory and inhibitory effects as compared to out-group faces. Furthermore, following from previous findings, we expected to find group membership modulations especially with a short ISI of 50 ms. Results from Experiment 1 showed a marginal interaction between Cue Validity and Group Membership on response accuracy, with participants performing worse after in-group faces than after out-group faces. In the RTs, instead, the Group Membership of the cue did not modulate participants' behavior. These results held in both the short and long ISI conditions.

Descriptive data from Experiment 2, visualized in Figure 5.5 suggest that the general inhibitory effect of invalid vs neutral cues over response accuracy was replicated in Experiment 2 as well, but there were no numerical differences between Group Membership conditions. Similarly, RTs results replicated both the general facilitatory effect of valid vs neutral cues, but confirmed that Group Membership did not affect participants' performance.

Overall, the results of the two studies (analyses over accuracy and RTs for Experiment 1 and over RTs for Experiment 2) suggested that in-group and out-group faces equally attracted participants' spatial attention in the paradigm we used. While in Experiment 1 there was a marginal interaction between Cue Validity and Group Membership, this interaction, which was not statistically tested due to convergence issues, was not numerically replicated in the second study. This might suggest that its effect is not very robust.

How can we reconcile our findings with previous ones?

How can we reconcile our findings with previous ones, which showed that attention was indeed affected by group membership (e.g., Chauhan et al., 2017;

Van Bavel & Cunningham, 2012; Park et al., 2016, Experiment 2). There are three sets of potential explanations behind the null effects observed in our studies:

a) *Participants did not identify enough with the in-group student.* We used previously unknown faces that only shortly before the Posner cueing task were categorized as in-group or out-group students. It might be the case that participants did not create shared representations between the in-group student and the self, which was indicated as the main mechanism through which self-reference is extended to in-group members ("identity fusion," see Swann et al., 2012). This explanation would suggest that longer training with increased exposure to the in-group and out-group faces could maximize the identification with the in-group student. This could for instance be achieved by having the academic affiliations learned a day before the Posner cueing task (see Chapters 3 and 4 of this dissertation for a similar procedure). Nevertheless, note that previous studies that reported in-group biases on attentional capture also used unfamiliar faces randomly assigned either to the participant's group or to an out-group just before the attention task (e.g., Park et al., 2016). Therefore, the lack of familiarity with the faces per se does not seem to explain entirely our null results.

b) *Self-biases on attention do not extend to in-group members.* While enhanced perceptual processing and better memory performance were robustly shown to extend from the self to in-group members (Enock et al., 2018; Hugenberg et al., 2010; Moradi et al., 2015; Van Bavel et al., 2008), it might be the case that in-group stimuli do not capture attention in the same way as self-related stimuli. One way to directly address this possibility is to make use of a procedure recently developed by Woźniak, Kourtis, and Knoblich (2018). In this study, a simple associative task, similar to the one we used in Experiment 2 as social learning task, was used to arbitrarily associate unfamiliar faces with concepts of the self, a friend (i.e., an in-group member), and a stranger (i.e., an out-group member). Via this procedure, the authors controlled for familiarity of the stimuli and were able to confirm that self-biases are easily obtained even with arbitrary faces (Woźniak et al., 2018). Even though the focus of the study was on self-biases, the results showed some evidence suggesting that associating unfamiliar faces to a friend or a stranger affected participants' behavioral and electrophysiological responses. A future study could include newly self-associated faces, in addition to in-group and out-group faces, and test whether they linearly modulate the orientation of spatial attention in a Posner cueing task. This experiment could directly test for similarities between self and in-group membership by assessing if participants' behavior to self and in-group cues co-vary.

c) *Attention orientation might not be affected by group membership.* In the current studies, we measured participants' spatial attention. However, attention orientation is only one of the sub-components of attention (Petersen & Posner, 2012). Previous studies on the relation between self-bias and attention orientation showed that self-related stimuli attracted more attention than other-related stimuli (Liu et al., 2016; Sui et al., 2009). However, the mixed evidence that we reviewed in this chapter shows that the relation between group membership and attention orientation is not very robust. It might be that attention is generally affected by group membership, as proposed by Humphreys and Sui (2016), but attention orientation might not the component that is mostly affected by it. For instance, a larger degree of sustained attention allocated when processing information from in-group members might explain enhanced memory performance for in-group vs out-group related information. More sustained attention to in-group members may allow individuals to encode episodic representations with more details, making these episodes easier to recall. Future studies using paradigms that target different attentional components could help elucidate the relationship between in-group biases and attentional capture and shed light its potential limits.

To conclude, results from our two studies show that there is no evidence for in-group biases modulating attentional shifts in a Posner cueing task. We suggested some explanations, which are not mutually exclusive and deserve further investigation.

6 | Summary and general discussion

Although people learn about how to communicate their thoughts and refer to the physical world by interacting with others (e.g., Tomasello, 2000; Lev-Ari, 2016; Levinson, 1995), most psycholinguistic research has focused on understanding the basic cognitive processes underlying language without accounting for the social context that shapes interactions. Many models of language processing assume that information related to the social context (e.g., speakers' identity) is filtered out when the linguistic input is processed (e.g., Gaskell & Marslen-Wilson, 1997; McClelland & Elman, 1986; Norris, 1994).

Recently, there has been a growing consensus about the need to model language processing in a way that can account for the encoding of both linguistic and contextual information (e.g., Cai et al., 2017; Hay, Nolan, & Drager, 2006; Kapnoula & Samuel, 2019; Nielsen, 2011; Münster & Knoeferle, 2018; Sumner et al., 2014; Drager & Kirtley, 2016). This consensus is supported by a similarly growing body of empirical evidence that shows that speakers' social identity, for example, is not filtered out but can influence how language is processed (e.g., Boland & Shana'e, 2014; Hay, Nolan, & Drager, 2006; Hay, Warren, & Drager, 2006; K. Johnson et al., 1999; Kim, 2016; Martin et al., 2016; van Berkum et al., 2008; Walker & Hay, 2011).

Among the many social factors that could be examined, I have chosen to investigate the potential role of speakers' social identity on linguistic processes. Specifically, I tested for differences in how listeners encode and process novel linguistic information, and in how they prioritize attentional allocation as a function of the group membership status of the speakers. The reason behind this choice is the extensive evidence, mostly from social psychology and social neuroscience, indicating that group membership status is a very salient social dimension. People exhibit a general automatic tendency to divide others into *us* or *them* (D. E. Brown, 1991; Turner & Tajfel, 1986), and this distinction has been shown to influence the way in which information from and about those others is processed and remembered (see in-group biases, e.g., Frable & Bem, 1985; Hugenberg et al., 2010; Van Bavel et al., 2008; Wilder, 1990). Yet, despite its relevance, there is little understanding of how group membership status can influence basic cognitive processes underlying language processing, making this a very promising field of study.

In this doctoral thesis, I sought to test the potential effects of speakers' group membership status and individual in-group biases on the representation and the processing of novel words, as well as the role of attention as the mechanism underlying in-group biases. In **Chapter 1**, I introduced three main working hypotheses according to which linguistic information learned from in-group speakers is (a) individuated to a larger extent and (b) retrieved and integrated in context in a faster and/or easier way than information learned from out-group speakers. Such differences in representation and processing between in-group and out-group information were predicted to be modulated by learners' individual in-group bias: The stronger the bias, the larger the in-group benefits should be. Finally, I suggested that (c) attentional prioritization may be the leading mechanism underlying in-group advantages in input processing. These hypotheses were tested throughout this thesis.

The current chapter summarizes the main findings of the preceding empirical chapters in relation to the hypotheses, discusses empirical challenges and theoretical implications, and highlights potential avenues for future research.

Summary of main findings

Hypothesis 1: Greater individuation of in-group information

Chapters 2 and 3 tested the hypothesis that in-group information is more individual-specific than out-group information and that this difference is modulated by the learners' own in-group bias. To test this hypothesis, I used a word learning paradigm in combination with a source memory test. In the learning phase, participants learned labels for uncommon gadgets from in-group (i.e., supposed students from the participants' own university) or out-group speakers (i.e., supposed students with a different academic affiliation). After the learning phase, a surprise memory test was administered to assess participants' source memory of the novel words. Speaker-label pairs were displayed and participants decided whether the pairs matched with what they learned in the previous phase, i.e., whether the speaker had produced the label. The source memory test, originally developed in the memory field to investigate context-dependent details encoded in memory traces, was used here to explore the degree of individuation of the speaker-related information encoded in the representations of the novel in-group and out-group words.

In **Chapter 2**, I validated the methods and demonstrated that the social identity of the speakers was encoded in the novel representations. This tendency was reflected in participants' source memory confusion: the likelihood of misattributing novel labels to incorrect speakers was higher when the incorrect speaker shared the group affiliation with the correct speaker than when they belonged to different groups. Crucially, I found no compelling evidence that the level of detail of speaker-specific information depended on the group membership status of the speaker or the individual in-group bias of the learners. Unexpectedly, results showed that it was not participants' sensitivity that was affected by the interaction between speakers' group membership and learners' in-group bias, but their response bias. The stronger the in-group bias, the more conservative participants were in their decisions, and this particularly applied to in-group related decisions. This unpredicted finding was interpreted as related to the *overexclusion hypothesis* (Leyens & Yzerbyt, 1992), which states that people prefer to not categorize uncertain targets as in-group members. The dataset from Chapter 2 seems to support this hypothesis. These differences in response bias between in-group and out-group related decisions as a function of learners' in-group bias might reflect asymmetries in the activation of the inhibitory system during decision processes in the recognition task (see Windmann et al., 2002, for the role of the inhibitory system on decision biases).

Despite the novel contribution of Chapter 2, the original hypothesis of greater individuation for in-group words was not supported. The outcomes of this chapter can be interpreted as lack of evidence for group membership effects on the *encoding* of speaker-related information. However, there is one aspect of the design that could have contributed to the null result: the group membership of the speakers was indicated by the logo of the academic affiliation, which was presented visually with the speakers' faces. This means that multiple pieces of information were potentially encoded and retrieved from memory at the same time (e.g., the speaker's affiliation, the speaker's personal identity, a label's source, the referred item) increasing the difficulty of the task at hand. Consequently, to reduce the number of details to be learned (i.e., only group affiliation instead of group affiliation *plus* personal identity), participants may have relied on the very salient visual cue about group membership in equal measures for both in-group and out-group members and to a greater extent than they would normally do.

To shed further light on this issue, in **Chapter 3**, I improved the experimental design of the previous chapter designing a 2-day experiment. On day

1, participants learned the group membership status of speakers in a more natural way (by listening to speakers referring to their university lives). On day 2, they performed the word learning task, followed by the surprise memory test. By having participants learn the speakers' affiliation in an earlier session, I eliminated the need for the overtly presented logo and, at the same time, simplified the task by reducing the number of details to be concurrently encoded. For the same purpose, I additionally decreased the number of speakers from whom participants learned the new labels from eight to six.

This procedure led to the following results: The predicted effect of speakers' group membership was not found at the group level; however, I did find the expected interaction between group membership and individual in-group bias. That is, the stronger the in-group bias, the less individuated out-group representations were as compared to in-group representations. The results are consistent with previous findings in favor of greater individual-specific encoding for in-group than for out-group members (e.g., Frable & Bem, 1985; Greenstein et al., 2016; Hugenberg et al., 2010; Wilder, 1990). I extend them to the case of word learning and show that, for newly learnt words, the extent to which people preferentially encode speaker-specific information in in-group vs. out-group words strongly depends on their individual in-group bias (see Amodio et al., 2003; Hein et al., 2010; Platow et al., 1990; Van Bavel et al., 2012, for effects of individual in-group bias).

Additionally, a comparison between the results reported in Chapter 2 and Chapter 3 suggests that (a) the lack of effects related to group membership in Chapter 2 could be attributable to the fact that multiple pieces of information were concurrently provided in the word learning phase. It also suggests that (b) the main effect of social categorization, reflected in source memory confusion, which was found only in Chapter 2, could depend on the great cognitive demands experienced because participants had to concurrently encode speakers' group affiliation and novel words. In situations of high cognitive load, people might be more likely to rely on social categories for representing linguistic information, regardless of the group membership status of the speaker, than they might be in low-cognitive load circumstances. Future research could test this hypothesis by keeping constant the way participants learn about the group membership affiliation of the speakers but manipulating the difficulty of the task at hand. For instance, one could have participants learn novel words from in-group and out-group speakers while performing an additional working-memory task where they need to keep in memory one (low load) or multiple (high load) digits until of the end of each learning trial. This experiment would help to explain the inconsistencies found in Chapters 2 and 3 and elucidate whether cognitive load may affect the likelihood of using social categories in representing novel linguistic information.

Hypothesis 2: Facilitated lexical retrieval for in-group words

To test the hypothesis that words learnt from in-group speakers are characterized by facilitated lexical retrieval, I used a word learning task in combination with a N400 paradigm in a 2-day experiment, which is described in **Chapter 4**.

Contrary to the studies in Chapters 2 and 3, which focused on the *encoding* of information in novel representations, the main goal of this chapter was to assess potential effects of speaker group membership status on the *on-line processing* of newly acquired words on later encounters, e.g., on the day following learning. To pinpoint on-line effects, encoding differences were minimized by providing participants with an intense training session of about 2.5 hours that ended only when participants demonstrated equally accurate in-group and out-group speaker-label-gadget associations, i.e., they were trained to reach the same criterion for in-group and out-group associations.

On the following day, participants' EEG signal was recorded while they were presented with the highly trained speaker-gadget sequences to generate specific expectations about the upcoming words. Crucially, the spoken words that participants heard either confirmed such expectations (i.e., being the *trained* label learned on day 1) or disconfirmed them (i.e., being a *non-trained*, competing, label).

I targeted N400 amplitudes for non-trained vs. trained labels and predicted that the difference between the two conditions (i.e., N400 effect) would be modulated by the group membership status of the speakers and by the individual in-group bias of the learners. Crucially, two potential, not mutually exclusive, outcomes could result in a larger N400 effect of in-group words as compared to out-group words. On the one hand, relative to non-trained words, the in-group words could elicit smaller N400 amplitudes than the out-group counterparts. Such a pattern would provide evidence for the hypothesis that lexical retrieval is facilitated for in-group words (see Bakker, Takashima, van Hell, Janzen, & McQueen, 2015, for lexical consolidation effects reflected in N400-amplitude modulations).

On the other hand, relative to trained out-group stimuli, in-group stimuli were expected to generate faster and/or stronger expectations about the upcoming words, eliciting more negative N400 amplitudes when such expectations were violated by presenting non-trained labels. This effect could be interpreted as showing that the strength of specific linguistic expectations generated by a context or cue (in this case the speaker's face followed by the gadget) is linked to the difficulty of integrating unexpected stimuli in that context or sequence (e.g., Altmann & Mirković, 2009; Dell & Chang, 2014; Federmeier, 2007; Huettig, 2015; Kutas, DeLong, & Smith, 2011; Pickering & Garrod, 2007, 2013).

Results showed that trained words were indeed more easily integrated in the speaker-gadget-word sequence than non-trained words, as indexed by an N400 effect, providing evidence for the lexicalization of those newly acquired words. However, there was no support for group membership effects on the magnitude of the N400 effect. Additional analyses tested for the contribution of learners' individual in-group bias on the observed patterns and showed that individual in-group biases did not influence participants' N400 effects for in-group vs out-group words either.

Many potential factors can be identified as responsible for the lack of group membership effects on semantic integration, as reflected in the N400. First, compared to the studies reported in Chapters 2 and 3, here participants did not learn *competing* labels for the same gadget from speakers belonging to different groups. Instead half the gadgets were uniquely labelled by in-group speakers, and the other half by out-group speakers. In addition, participants heard both speakers from each affiliation share the same labels, whereas in previous studies the source, i.e., the speaker, of each label was only one speaker. Therefore, the lack of competition in gadget-labelling, and the lack of speaker-specific encoding may have reduced participants' motivation and/or need to encode in-group and out-group information differently.

Another potential explanation for the lack of an effect of group membership is that the requirement to capture effects on on-line processing of newly acquired words had eliminated any potential retrieval or integration differences: limiting the *encoding* differences between in-group and out-group associations (via intensive training) could have curbed any potential in-group advantages on the on-line processing of newly acquired words. To support this intuition, I explored the behavioral data from the learning tasks and found some evidence that suggests that, during the first learning round, participants with stronger in-group bias needed fewer trials to learn the in-group associations as compared to out-group associations. Such an effect was absent in the second learning round. While this evidence is exploratory in nature and calls for further testing, it suggests that intense training to an equal criterion could have contributed to the lack of group membership effects during on-line processing.

This account would leave open the possibility that in an experimental setting where equal learning for in-group and out-group membership is not guaranteed, and where competing information is provided by speakers from different groups, on-line processing of newly acquired words could be affected by group membership. This is a hypothesis that remains to be tested. The results would help determine (a) whether in-group biases on on-line language processing are indeed absent or can arise under specific circumstances, (b) whether or not, in those circumstances, group membership effects on processing reflect cascading effects driven by encoding differences. Crucially, further evidence for a lack of group membership effects on on-line processing of well-established words would indicate that group membership might exclusively influence to what degree episodic (e.g., speaker-related) information is encoded when new words are learned, but not the ease with which the words are then consolidated and integrated in semantic memory. Although previous evidence indicated that talker-specific effects on novel word processing persisted for at least a week after learning (e.g., H. Brown & Gaskell, 2014), showing the stability of speaker-specific information in memory representations, it might be that biases due to the speakers' group membership status are less stable. They might depend on episodic encoding to a larger extent than talker-specific effects do. Therefore, they might quickly decay after off-line consolidation or need specific circumstances (e.g., repeated exposures across a longer period of time) to become stable.

Hypothesis 3: Attentional prioritization by in-group membership

In **Chapter 5**, I tested the third working hypothesis that predicted that attentional benefits may underlie in-group biases. The hypothesis was supported by social psychological theories on how social identity is formed. These theories claim that the social salience ascribed to stimuli related to the self should extend to material related to in-group members (e.g., Brewer, 1999; Turner & Tajfel, 1986) via shared psychological representations ("identity fusion," see Swann et al., 2012). There is extensive empirical evidence that self-related stimuli receive attentional advantages, relative to other-related stimuli (e.g., Alexopoulos et al., 2012; Bargh, 1982; Brédart et al., 2006; Liu et al., 2016; Moray, 1959; Sui et al., 2009), and that those advantages may underlie self-biases in processing and memory (Humphreys & Sui, 2016).

Because of shared representations between the self and in-group members, similar attentional priorities have been suggested to extend beyond the self, to favor in-group members as well (e.g., Bargh, 1982; Humphreys & Sui, 2016; Van Bavel & Cunningham, 2012). The two experiments described in Chapter 5 are not the first attempt to test this hypothesis. However, results from previous studies reported mixed findings. In some cases, participants' attention was directed preferentially to in-group members (e.g., Chauhan et al., 2017; Van Bavel & Cunningham, 2012; Park et al., 2016, Experiment 2), whereas in other cases out-group members were those who received such attentional benefits (e.g., Brosch & Van Bavel, 2012; Park et al., 2016, Experiment 2).

In order to shed light on the issue, I used an adapted version of a Posner cueing task (Posner, 1980), previously employed in Liu et al. (2016). In the version designed by Liu et al., faces dynamically turned to the sides of the screen and acted as valid/invalid cues to the location of subsequent targets, which had to be classified for their orientation. In Liu et al. (2016), participants' own faces oriented their spatial attention to a greater extent than other faces did, as indexed by a larger cueing effect (i.e., valid vs invalid condition). I used the same paradigm as in Liu et al. and included a neutral-cue condition (i.e., faces did not turn) as reference level to be able to tease apart potential facilitatory and inhibitory effects. The group membership of the face was expected to affect the magnitude of these effects with larger facilitatory and inhibitory effects being elicited by in-group as compared to out-group facial cues.

In a first phase, participants acquired the group membership status of two faces by listening to them refer to their university life (Experiment 1) or by association with academic logos- (Experiment 2). Then I measured their target classification behavior in the spatial orientation task. Across the two experiments, participants were consistently faster at target classification after valid cues, compared to neutral ones. However, this facilitatory effect was not modulated by the group membership status of the faces. In Experiment 1, there was an inhibitory effect on participants' accuracy, with responses after invalid cues being less accurate than after neutral ones. This effect was larger after in-group than after out-group faces. Convergence issues with statistical models in Experiment 2, potentially driven by ceiling effects, did not allow me to test the replicability of this pattern. No inhibitory effect was found in the RTs in either experiment, indicating that the neutral condition was comparable to the invalid condition.

The lack of group membership effects on attentional orientation can be interpreted in different ways. First, it may indicate that participants did not identify enough with the in-group member. According to the *identity fusion hypothesis* (see Swann et al., 2012), shared representations are needed to extend the self-reference effect to in-group members. Since the faces were only briefly introduced before the Posner cueing task, exposure may have not been strong enough to mediate the social salience effect on attention. Longer training, as provided in the studies described in Chapters 3 and 4, may have produced different results. A second explanation is that self-biases on attention do not extend to group members. Together with the well-known *publication bias*, the mixed findings reported in the literature could indicate that the effect may be less robust than one would expect. A meta-analysis on this issue would help us understand whether and how group-membership status affects attention prioritization, and how this relationship relates to the encoding and memory advantages reported in the literature.

Overall discussion: what have we learnt?

To sum up, in the current thesis I aimed to test three specific hypotheses. In Chapter 2 and 3, I tested the prediction that, relative to out-group words, words learned from in-group speakers would result in more individuated representations, and this difference in encoding would be modulated by learners' own in-group bias. The results did not support this prediction but showed that participants' response bias in the source memory task was jointly modulated by group membership and individual in-group bias, with strongly biased participants being more conservative in their decisions, especially when these decisions concerned input from members of their own in-group. This finding suggests that, for strongly in-group biased participants, the inhibitory system may have been activated to different extents during decision-making processes that involved in-group and out-group membership.

After making some adjustments to the experimental design, I ran the study described in Chapter 3. Crucially, results showed that in-group representations of newly acquired words were more individuated than out-group representations, but only for participants that exhibited a strong in-group bias.

In Chapter 4, I tested the hypothesis that, during on-line processing, newly acquired words learned from in-group speakers would be characterized by facilitated lexical retrieval and integration in the given context. These predictions were tested by employing a N400 paradigm where participants were shown highly trained face-gadget sequences followed by target words that either confirmed or disconfirmed the expectations generated by the sequence of preceding stimuli. That is, the target words were either the previously trained labels or a non-trained, novel labels. I expected a N400 effect as the difference in semantic integration for non-trained vs trained words. Crucially, I predicted that in-group cues would generate stronger expectations about the upcoming word, leading to a larger N400 effect for in-group than for out-group words. This difference was expected to be modulated by individual in-group biases. Results showed that, relative to non-trained stimuli, trained words were indeed more easily retrieved and integrated in the given sequences, as indexed by the predicted N400 effect. However, no effects of group membership or in-group bias were found to modulate the magnitude of the N400 effects.

Taken together, the findings suggest that, at least for strongly biased individuals, speakers' group membership status may influence how learners encode and organize novel linguistic information, with representations of in-group words being more individuated than the out-group counterparts. By contrast, these social biases seem not to influence the subsequent on-line processing of newly acquired words.

Finally, in Chapter 5 I tested the hypothesis that faces of in-group members would preferentially attract attentional resources in a spatial orientation task. I opted for a version of a Posner cueing task and used dynamically turning in-group and out-group faces that provided valid, invalid or no cues to the subsequent locations of targets that participants had to classify for its orientation. l operationalized attention allocation in terms of cueing effects on the target categorization task. I found no evidence that in-group cues attracted more spatial attention than out-group cues did. This finding contributes to the large body of mixed evidence, which suggests that stimuli pertaining to an in-group member may not reliably orient attention to a greater extent than stimuli pertaining to an out-group.

Importantly, the lack of a group membership effect on automatic attention orientation in Chapter 5 suggests that there may be other cognitive mechanisms underlying in-group biases in perception and memory. Relevant to this last point, in Chapter 2 I found that effects of group membership and in-group bias jointly affected participants' response biases. Response biases have been shown to depend on criterion setting functions of the prefrontal cortex (Schacter et al., 1998; Swick & Knight, 1999; Miller et al., 2001), which is involved in supporting high-level executive functions necessary for initiating, monitoring and controlling item-retrieval from memory (Buckner, 1996; Fletcher et al., 1998; Henson et al., 1999; Tomita et al., 1999; Wagner et al., 1998). The finding from Chapter 2 suggests that more strategic components involved in response bias processes may be sensitive to individual in-group bias and group membership status.

Taken together, the findings from Chapters 2 and 5 can be interpreted as showing that in tasks that tap into attention mechanisms that are more automatic and rely on bottom-up information processing, e.g., visual attention orientation, participants' behavior may be less sensitive to social top-down manipulations than it is in tasks that permit more strategic and deliberate processes. This might be especially the case when the affiliation to group membership is not based on features that are perceptually more salient and can bias behavior in a bottom-up fashion, e.g., race and gender.

Experimental challenges and future directions

During this PhD project, many new empirical questions arose, mainly due to the inconsistency of some of the results I reported in the current thesis. Because of time constraints, these questions were not explored further or assessed experimentally. In the next section, I consider some design limitations of the studies which limit our understanding of group membership effects on language processes. Additionally, I propose ideas for new experiments that could unravel the inconsistencies and provide informative answers.

Do cognitive demands influence the likelihood of relying on social categories at the expenses of individuation processes when encoding novel linguistic input?

To understand the contrasting results obtained in Chapters 2 and 3[1], one could argue that the participants tested in the main study of Chapter 2 performed a task characterized by higher cognitive demands, relative to the task performed in the study of Chapter 3. In fact, in the former experiment, participants had to learn speakers' group affiliation (i.e., logo) and personal identity (i.e., face) while also learning gadgets' labels and pictorial referents. On the other hand, in Chapter 3, participants learned the speakers' group membership and personal identity before performing the word learning task. Moreover, in Chapter 3, I also reduced the number of speakers from whom participants learned new words. An interesting novel hypothesis arising from these results is that cog-

[1]In Chapter 2, I found a general tendency for social categorization, reflected in source memory confusion, and this tendency was not modulated by group membership and in-group biases. In Chapter 3, instead, I found that there was no a general tendency for social categorization, but this effect was qualified in a 3-way interaction with group membership and in-group bias.

nitive load interacts with in-group bias to determine the encoding of novel words.

Empirical work on the formation of and reliance on stereotypes suggests that relying on broader social categories, and associated stereotypes, is cognitively less demanding than processing information in an individual-specific fashion, hence people tend to prefer category-based processing when their cognitive resources are limited (e.g., Fiske & Neuberg, 1989; Macrae, Milne, & Bodenhausen, 1994). Importantly, motivational factors have been reported to influence people's tendency to rely on individual- or group-based processing (see Fiske & Neuberg, 1989).

One could therefore argue that participants who exhibit stronger in-group bias may be more motivated to invest cognitive resources to encode in-group information in an individuated manner, as compared to less in-group biased participants. However, in a situation in which the cognitive load is high, for instance because of concomitant tasks or concurrently encoded pieces of information, those participants in favor of in-group individuation may experience cognitive overload. Two alternative outcomes can be predicted in those circumstances: either individuals process both in-group and out-group information in a category-based processing manner so that they can perform the other task at hand accurately (similar to what may have happened in Chapter 2), or they keep relying on individuating processes for in-group information, but their performance in the other task may be impaired. Presumably, the shift from one to the other processing manner depends on resource availability and motivational factors including in-group bias.

To test these predictions in the context of novel word learning, studies that vary cognitive load are required. One such study could use a between-participant cognitive load manipulation, keeping the way group membership affiliation is learnt constant across load conditions. For instance, participants would learn speakers' group membership on day 1 via listening to speakers' student habits, like in Chapter 3. On day 2, participants would perform the word learning task designed for Chapters 2 and 3 while carrying out an additional task, e.g., a working memory task. Crucially, the working memory task would vary in cognitive load, such that one group of participants would perform a simple version (e.g., holding in memory the result of a simple arithmetic operation until the end of the trial; low cognitive load condition), whereas another group of participants would perform a harder version (e.g., holding in memory the result of a complex arithmetic operation until the end of the trial; high cognitive load condition). Participants' source memory of novel words would then be tested, as well as their individual in-group bias, as in Chapters 2-4.

By contrasting the two groups in how they encode in-group vs out-group words (i.e., more category-based vs individual-based), and how well they perform the working memory task, we could determine whether cognitive load and group membership jointly influence the encoding of speaker-specific information in novel lexical representations. Additionally, by including individual in-group bias in the analysis, we could assess the hypothesis that the performance of participants exhibiting strong in-group bias would be more influenced by the cognitive load manipulation, relative to those participants who are less strongly biased.

Does extensive training overwrite potential effects of group membership on on-line processing of newly learnt words?

In Chapter 4, I aimed to test the hypothesis that words learned from in-group speakers are characterized by facilitated lexical retrieval, compared to words learned from out-group speakers. As a requirement to explore group membership effects on on-line processing, participants received intense training of ca.

2.5 hours until they demonstrated equally accurate in-group and out-group speaker-label-gadget associations. Crucially, I found no evidence that group membership and in-group biases modulate lexical retrieval of newly acquired words, as reflected by N400 effects of equal magnitude for in-group and out-group words. I proposed that the training procedure could have reduced the chances of detecting a group membership effect on the on-line processing of newly acquired words by minimizing any potential differences in strength of memory representations between in-group and out-group words.

Indeed, one might question whether group membership effects may affect the retrieval of well learned labels at all. In the study in Chapter 4, I tested whether listeners would process words differently because they were cued by in-group vs out-group face-gadget associations. Relative to out-group face-gadget sequences, I expected the face of in-group speakers, followed by the picture of a gadget, to activate the associated referring label more strongly, in line with the premise that in-group membership is perceived as highly socially salient, and therefore the in-group labels are attended more strongly and/or faster than out-group labels. As a result of this difference in activation, I predicted N400 effects of different magnitudes for in-group and out-group words. However, by training participants to reach the same criterion for in-group and out-group associations, I may have destroyed the chances to get group membership differences in on-line word processing.

To assess whether the training was indeed responsible for the lack of group membership effects on later on-line processing of newly acquired words, a new study needs to be carried out. I suggest an experimental design where participants are trained for a fixed, shorter period of time. After the training, the participants' source knowledge of the words is established. The EEG task is then performed exactly as described in Chapter 4. If again no differences are found in lexical retrieval for in-group and out-group words, reflected in similar N400 effects for in-group and out-group words, we would conclude that group membership does not affect on-line processing of words, but only word encoding. If an N400 effect is found, it would indicate that differences in the initial learning and encoding of ingroup vs. outgroup words percolate to the retrieval task on the next day. However, if there is an effect of group membership on on-line word processing, one might ask whether it reflects cascading effects driven by encoding differences between in-group and out-group words. Analyzing participants' source memory performance in relation to the N400 effects might shed light on the extent to which group membership-driven differences in the encoding of novel words influence the strength of activation of these words during later on-line word processing.

Another interesting follow-up question that would be worth assessing is: to what extent do group membership effects on on-line word processing depend on the activation of speaker-gadget-label associations learned during the training? One way to address this question would be to generate expectations about the newly acquired words without relying on face-gadget cues, which may be strongly dependent on the strength of the encoded associations. For instance, one could turn to a more classical N400 paradigm where target words are embedded in sentences, visually presented one word at the time (e.g., *This morning I wanted a lemonade and I used my new* ***citrus-schiller***). In this case, it is not the trained speaker-gadget sequence but the semantic content of the sentences that generates specific expectations for one of the newly acquired words. Relative to trained target words, non-trained words should still elicit larger N400 amplitudes. If group membership effects were to modulate the magnitude of N400 effects in this paradigm, this evidence would indicate that the effects are not strictly dependent on the reactivation of the trained associations, but are linked to the relative ease of integrating in-group vs out-group words into semantic contexts.

Are attentional resources prioritized when processing in-group related information?

The results of Chapters 2 and 5 do not allow us to decide whether or not in-group membership prioritizes attentional allocation. On the one hand, in Chapter 5, I did not find a group membership effect on spatial orientation. On the other hand, in Chapter 2 I reported that group membership and in-group bias jointly modulated participants' response bias, which may be indicative of differences in the activation of prefrontal cortex areas during decision-making processes (e.g., Windmann et al., 2002). Given the numerous methodological differences between the studies, any interpretations of this divergence in the results can only be tentative. Nonetheless, these results may suggest that strategic components involved in decision making and response bias may be more susceptible to social top-down manipulations than the orientation of visual attention, which is a more automatic and bottom-up process. Moreover, the mixed evidence in the literature for in-group biases on attention captures suggests that the effect of group membership may not be very robust, or it may depend on very specific circumstances, for instance when group membership is defined by perceptually salient features (e.g., race, or gender).

I suggest that in-group biases might affect participants' performance in tasks that allow more strategic behavior to take place, with in-group members attracting more attention than out-group members. To test this hypothesis, one could employ tasks that have been previously shown to be sensitive to top-down manipulations. For instance, selective attention tasks, like the dot-probe task, have been reported to be affected by top-down demands, such as goals and expectations (e.g., Yantis, 1996), emotional valence (e.g., Bradley, Mogg, & Lee, 1997; Yiend, 2010) and even group membership manipulation (e.g., Brosch & Van Bavel, 2012). Future studies could assess this prediction by making use of one of these tasks. This would elucidate whether group membership effects on attention modulation are task-dependent or not.

To what degree are the patterns of results I found linked to the specific group membership distinction utilized?

I manipulated group membership in my experiments by having different speakers supposedly attending participants' own university (i.e., in-group speakers) or belonging to a different academic affiliation (i.e., out-group speakers). One may argue that I selected a rather "weak" group membership distinction compared to more perceptually salient factors, like race or gender, or dimensions that are arguably more relevant for self-categorization, like one's political orientation or favorite football team. Indeed, I cannot exclude that one of the factors that could have contributed to the pattern of results obtained is the type of group membership manipulation used in this thesis. While I agree with this argument in principle, I would like to highlight that the aim was to pinpoint effects that could *solely* be explained by the group membership manipulation, and not by additional stereotypes and prejudices which intrinsically characterized other more salient dimensions.

I predict that a dimension that encompasses stronger and more consolidated social biases, in addition to group membership, would lead to larger effects of group membership. For instance, I predict that the effect of group membership found in Chapter 3 only for strongly biased participants could have been boosted by a stronger group membership distinction, like political orientation, and affected participants' group-level behavior. Similarly, using race, for instance, as group membership manipulation may actually influence attention orientation, due to its high perceptual saliency. These are hypotheses that could be tested by using group membership manipulations that are perceived as more socially relevant in real life.

In addition, group membership biases may have stronger effects in face-to-face interactions. One way to assess this intuition is to have highly trained confederates, characterized as either in-group or out-group members, contrastively labelling gadgets as part of an interactive task and then test participants' representations of these newly acquired in-group and out-group words. I predict the results from Chapter 3 would be replicated, and potentially be detected at the participant group level. i.e., independently of individual in-group bias.

Conclusions

I started this doctoral thesis introducing three main working hypotheses. I tested these hypotheses using a combination of behavioral and electrophysiological tasks that examined group membership effects and individual in-group biases on the **encoding** of novel lexical representations, on the **on-line processing** of the newly acquired words, and on **attention orientation**. Results from these studies provide compelling evidence only for the first working hypothesis showing that the level of detail of speaker-specific information encoded in the representations of novel words depends on the group membership status of the speakers and the individual in-group bias of the learners. In sum, this thesis contributes to our understanding of how speakers' social identity and learners' in-group biases influence the way linguistic input is represented. Furtherly, the thesis also discusses experimental challenges and theoretical implications, and highlights potential avenues for future research.

References

Alday, P. (2019). How much baseline correction do we need in erp research? extended glm model can replace baseline correction while lifting its limits. *Psychophysiology*(56:e13451).

Alexopoulos, T., Muller, D., Ric, F., & Marendaz, C. (2012). I, me, mine: Automatic attentional capture by self-related stimuli. *European Journal of Social Psychology*, *42*(6), 770-779.

Altmann, G. T., & Mirković, J. (2009). Incrementality and prediction in human sentence processing. *Cognitive science*, *33*(4), 583-609.

Amodio, D. M., Harmon-Jones, E., & Devine, P. G. (2003). Individual differences in the activation and control of affective race bias as assessed by startle eyeblink response and self-report. *Journal of personality and social psychology*, *84*(4), 738.

Aron, A., Aron, E. N., Tudor, M., & Nelson, G. (1991). Close relationships as including other in the self. *Journal of personality and social psychology*, *60*(2), 241.

Babel, M. (2012). Evidence for phonetic and social selectivity in spontaneous phonetic imitation. *Journal of Phonetics*, *40*(1), 177-189.

Bakker, I., Takashima, A., van Hell, J. G., Janzen, G., & McQueen, J. M. (2015). Changes in theta and beta oscillations as signatures of novel word consolidation. *Journal of Cognitive Neuroscience*, *27*(7), 1286-1297.

Bandura, A., & Walters, R. H. (1977). *Social learning theory (vol. 1)*. Englewood Cliffs, NJ: Prentice-hall.

Bargh, J. A. (1982). Attention and automaticity in the processing of self-relevant information. *Journal of Personality and Social Psychology*, *43*(3), 425.

Bargh, J. A., Schwader, K. L., Hailey, S. E., Dyer, R. L., & Boothby, E. J. (2012). Automaticity in social-cognitive processes. *Trends in Cognitive Sciences*, *16*(12), 593-605.

Bates, D., Kliegl, R., Vasishth, S., & Baayen, H. (2015). *Parsimonious mixed models. arxiv.* (preprint)

Bates, D., Maechler, M., Bolker, B., & Walker, S. (2014). lme4: Linear mixed-effects models using eigen and s4. *R package version*, *1*, 7.

Blascovich, J., Wyer, N. A., Swart, L. A., & Kibler, J. L. (1997). Racism and racial categorization. *Journal of Personality and Social Psychology*, *72*(6), 1364.

Boland, J. E., & Shana'e, N. C. (2014). Homophone disambiguation and vocal stereotypes. *Ann Arbor*, *1001*(48109).

Bradley, B. P., Mogg, K., & Lee, S. C. (1997). Attentional biases for negative information in induced and naturally occurring dysphoria. *Behaviour Research and Therapy*, *35*, 911-927.

Brewer, M. B. (1999). The psychology of prejudice: Ingroup love and outgroup hate? *Journal of social issues*, *55*(3), 429-444.

Brosch, T., Sander, D., Pourtois, G., & Scherer, K. R. (2008). Beyond fear: Rapid spatial orienting toward positive emotional stimuli. *Psychological science*, *19*(4), 362-370.

Brosch, T., & Van Bavel, J. J. (2012). The flexibility of emotional attention: Accessible social identities guide rapid attentional orienting. *Cognition*, *125*(2), 309-316.

Brown, D. E. (1991). *Human universals.* Philadelphia: Temple University Press.

Brown, H., & Gaskell, M. G. (2014). The time-course of talker-specificity and lexical competition effects during word learning. *Language, Cognition and Neuroscience*, *29*(9), 1163-1179.

Brown-Schmidt, S., Yoon, S. O., & Ryskin, R. A. (2015). People as contexts in conversation. In *Psychology of learning and motivation (vol* (p. 59-99). Academic Press: 62.

Brédart, S., Delchambre, M., & Laureys, S. (2006). Short article: one's own face is hard to ignore. *Quarterly Journal of Experimental Psychology*, *59*(1), 46–52.

Buckner, R. L. (1996). Beyond hera: Contributions of specific prefrontal brain areas to long-term memory retrieval. *Psychonomic Bulletin Review*, *3*(2), 149-158.

Cadinu, M. R., & Rothbart, M. (1996). Self-anchoring and differentiation processes in the minimal group setting. *Journal of personality and social psychology*, *70*(4), 661.

Caharel, S., Poiroux, S., Bernard, C., Thibaut, F., Lalonde, R., & Rebai, M. (2002). Erps associated with familiarity and degree of familiarity during face recognition. *International Journal of Neuroscience*, *112*(12), 1499-1512.

Cai, Z. G., Gilbert, R. A., Davis, M. H., Gaskell, M. G., Farrar, L., Adler, S., & Rodd, J. M. (2017). Accent modulates access to word meaning: Evidence for a speaker-model account of spoken word recognition. *Cognitive Psychology*, *98*, 73-101.

Castano, E., Yzerbyt, V., Bourguignon, D., & Seron, E. (2002). Who may enter? the impact of in-group identification on in-group/out-group categorization. *Journal of Experimental Social Psychology*, *38*(3), 315-322.

Chang, D. G., F., & Bock, K. (2006). Becoming syntactic. *Psychological review*, *113*(2), 234.

Chauhan, V., Visconti di Oleggio Castello, M., Soltani, A., & Gobbini, M. I. (2017). Social saliency of the cue slows attention shifts. *Frontiers in psychology*, *8*, 738.

Clark, H. H., & Carlson, T. B. (1981). Context for comprehension. *Attention and performance IX*, *313*, 30.

Clark, H. H., & Marshall, C. R. (1981). Definite reference and mutual knowledge. In W. B. Joshi A. & I. Sag (Eds.), *Elements of discourse understanding* (p. 10–63). Cambridge: Cambridge University Press.

Conway, M. A. (2005). Memory and the self. *Journal of memory and language*, *53*(4), 594–628.

Conway, M. A., & Pleydell-Pearce, C. W. (2000). The construction of autobiographical memories in the self-memory system. *Psychological review*, *107*(2), 261.

Cunningham, S. J., Brebner, J. L., Quinn, F., & Turk, D. J. (2014). The self-reference effect on memory in early childhood. *Child Development*, *85*(2), 808-823.

Cunningham, S. J., Turk, D. J., Macdonald, L. M., & Macrae, C. N. (2008). Yours or mine? ownership and memory. *Consciousness and cognition*, *17*(1), 312-318.

Cunningham, W. A., Van Bavel, J. J., & Johnsen, I. R. (2008). Affective flexibility: evaluative processing goals shape amygdala activity. *Psychological Science*, *19*(2), 152-160.

Dell, G. S., & Chang, F. (2014). The p-chain: Relating sentence production and its disorders to comprehension and acquisition. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *369*(1634), 20120394.

Dovidio, J. F., & Gaertner, S. L. (2000). Aversive racism and selection decisions: 1989 and 1999. *Psychological science*, *11*(4), 315-319.

Dovidio, J. F., Kawakami, K., Johnson, C., Johnson, B., & Howard, A. (1997). On the nature of prejudice: Automatic and controlled processes. *Journal of experimental social psychology*, *33*(5), 510-540.

Drager, K., & Kirtley, M. J. (2016). Awareness, salience, and stereotypes in exemplar-based models of speech production and perception. In M. Babel (Ed.), *Awareness and control in sociolinguistic research* (p. 1-24). Cambridge University Press: Cambridge.

Dunham, Y. (2018). Mere membership. *Trends in cognitive sciences*, *22*(9), 780-793.

Enock, F., Sui, J., Hewstone, M., & Humphreys, G. W. (2018). Self and team prioritisation effects in perceptual matching: Evidence for a shared representation. *Acta Psychologica*, *182*, 107-118.

Federmeier, K. D. (2007). Thinking ahead: The role and roots of prediction in language comprehension. *Psychophysiology*, *44*(4), 491-505.

Fiske, S. T., & Neuberg, S. L. (1989). Category-based and individuating processes as a function of information and motivation: Evidence from our laboratory. In K. A. S. W. Bar-Tal D. Graumann C.F. (Ed.), *Stereotyping and prejudice. springer series in social psychology.* (p. 83-103). New York: Springer.

Fletcher, P. C., Shallice, T., Frith, C. D., Frackowiak, R. S., & Dolan, R. J. (1998). The functional roles of prefrontal cortex in episodic memory. ii. retrieval. *Brain: a Journal of Neurology*, *121*(7), 1249-1256.

Frable, D. E., & Bem, S. L. (1985). If you are gender schematic, all members of the opposite sex look alike. *Journal of Personality and Social Psychology*, *49*(2), 459.

Frömer, R., Maier, M., & Abdel Rahman, R. (2018). Group-level eeg-processing pipeline for flexible single trial-based analyses including linear mixed models. *Frontiers in neuroscience*, *12*(48).

Garrod, G. C., S, & Pickering, M. (2014). Prediction at all levels: forward model predictions can enhance comprehension. *Language, Cognition and Neuroscience*, *29*(1), 46–48.

Gaskell, M. G., & Marslen-Wilson, W. D. (1997). Integrating form and meaning: A distributed model of speech perception. *Language and cognitive Processes*, *12*(5-6), 613-656.

Goldinger, S. D. (2007). A complementary-systems approach to abstract and episodic speech perception. *In Proceedings of the 16th international* *congress of phonetic sciences*, *16*, 49-54.

Greenstein, M., Franklin, N., & Klug, J. (2016). In-group versus out-group source memory. *Experimental psychology*, *63*(3), 150-158.

Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: attitudes, self-esteem, and stereotypes. *Psychological Review*, *102*(1), 4.

Greve, A., Cooper, E., Kaula, A., Anderson, M. C., & Henson, R. (2017). Does prediction error drive one-shot declarative learning? *Journal of memory and language*, *94*, 149-165.

Hay, J., Nolan, A., & Drager, K. (2006). From fush to feesh: Exemplar priming in speech perception. *The Linguistic Review*, *23*(3), 351-379.

Hay, J., Warren, P., & Drager, K. (2006). Factors influencing speech perception in the context of a merger-in-progress. *Journal of Phonetics*, *34*(4), 458-484.

Hein, G., Silani, G., Preuschoff, K., Batson, C. D., & Singer, T. (2010). Neural responses to ingroup and outgroup members' suffering predict individual differences in costly helping. *Neuron*, *68*(1), 149-160.

Henson, R., & Gagnepain, P. (2010). Predictive, interactive multiple memory systems. *Hippocampus*, *20*(11), 1315-1326.

Henson, R., Shallice, T., & Dolan, R. J. (1999). Right prefrontal cortex and episodic memory retrieval: a functional mri test of the monitoring hypothesis. *Brain*, *122*(7), 1367-1381.

Hewstone, M., Rubin, M., & Willis, H. (2002). Intergroup bias. *Annual review of psychology*, *53*(1), 575-604.

Heyes, C. (1994). Social learning in animals: categories and mechanisms. *Biological Reviews*, *69*(2), 207-231.

Heyes, C. (2011). Automatic imitation. *Psychological bulletin*, *137*(3), 463-483.

Heyes, C. (2012). What's social about social learning? *Journal of Comparative Psychology*, *126*(2), 193-202.

Heyselaar, E., Hagoort, P., & Segaert, K. (2017). In dialogue with an avatar, language behavior is identical to dialogue with a human partner. *Behavior research methods*, *49*(1), 46-60.

Huettig, F. (2015). Four central questions about prediction in language processing. *Brain research*, *1626*, 118-135.

Hugenberg, K., Young, S. G., Bernstein, M. J., & Sacco, D. F. (2010). The categorization-individuation model: An integrative account of the other-race recognition deficit. *Psychological Review*, *117*(4), 1168.

Humphreys, G. W., & Sui, J. (2016). Attentional control and the self: the self-attention network (san). *Cognitive neuroscience*, *7*(1-4), 5-17.

Iacozza, S., Meyer, A. S., & Lev-Ari, S. (2019a). How in-group bias influences source memory for words learned from in-group and out-group speakers. *Frontiers in human neuroscience*, *13*, 308.

Iacozza, S., Meyer, A. S., & Lev-Ari, S. (2019b). How in-group bias influences the level of detail of speaker-specific information encoded in novel lexical representations. *Journal of experimental psychology: Learning, memory, and cognition*.

Insko, C. A., Schopler, J., Hoyle, R. H., Dardis, G. J., & Graetz, K. A. (1990). Individual-group discontinuity as a function of fear and greed. *Journal of Personality and Social Psychology*, *58*(1), 68.

Insko, C. A., Schopler, J., & Sedikides, C. (1998). Personal control, entitativity, and evolution. In I. Cognition & I. Behavior (Eds.), *ed* (p. 109-20). C Sedikides, J Schopler, CA Insko, . Mahwah, NJ: Erlbaum.

Johnson, A. L., Crawford, M. T., Sherman, S. J., Rutchick, A. M., Hamilton, D. L., Ferreira, M. B., & Petrocelli, J. V. (2006). A functional perspective on group memberships: Differential need fulfillment in a group typology. *Journal of Experimental Social Psychology*, *42*(6), 707-719.

Johnson, K., Strand, E. A., & D'Imperio, M. (1999). Auditory–visual integration of talker gender in vowel perception. *Journal of Phonetics*, *27*(4), 359-384.

Jones, C. R., & Fazio, R. H. (2010). Person categorization and automatic racial stereotyping effects on weapon identification. *Personality and Social Psychology Bulletin*, *36*(8), 1073-1085.

Judd, C. M., & Park, B. (1988). Out-group homogeneity: Judgments of variability at the individual and group levels. *Journal of Personality and Social Psychology*, *54*(5), 778.

Kapnoula, E. C., & Samuel, A. G. (2019). Voices in the mental lexicon: Words carry indexical information that can affect access to their meaning. *Journal of Memory and Language*, *107*, 111-127.

Kim, J. (2016). Perceptual associations between words and speaker age. *Laboratory Phonology: Journal of the Association for Laboratory Phonology*, *7*(1).

Kinzler, K. D., Corriveau, K. H., & Harris, P. L. (2011). Children's selective trust in native-accented speakers. *Developmental Science*, *14*(1), 106-111.

Kleinschmidt, D. F., & Jaeger, T. F. (2015). Robust speech perception: recognize the familiar, generalize to the similar, and adapt to the novel. *Psychological review*, *122*(2), 148.

Kuperberg, G. R., & Jaeger, T. F. (2016). What do we mean by prediction in language comprehension? *Language, cognition and neuroscience*, *31*(1), 32–59.

Kutas, M., DeLong, K. A., & Smith, N. J. (2011). A look around at what lies ahead: Prediction and predictability in language processing. In M. Bar (Ed.), *Predictions in the brain: Using our past to generate a future.* Oxford: Oxford University Press.

Kutas, M., & Federmeier, K. D. (2011). Thirty years and counting: finding meaning in the n400 component of the event-related brain potential (erp). *Annual review of psychology*, *62*, 621-647.

Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmertest package: tests in linear mixed effects models. *Journal of Statistical Software*, *82*(13).

Lenth, R. (2018). Emmeans: Estimated marginal means, aka least-squares means. *R package version*, *1*(1).

Leshikar, E., Dulas, M., & Duarte, A. (2015). Self-referencing enhances recollection in both young and older adults. *Aging, Neuropsychology, and Cognition*, *22*(4), 388-412.

Lev-Ari, S., & Keysar, B. (2012). Less-detailed representation of non-native language: Why non-native speakers' stories seem more vague. *Discourse Processes*, *49*(7), 523-538.

Lev-Ari, S., & Peperkamp, S. (2014). An experimental study of the role of social factors in language change: The case of loanword adaptations. *Laboratory Phonology*, *5*(3), 379-401.

Levinson, S. C. (1995). Interactional biases in human thinking. In E. N. Goody (Ed.), *Social intelligence and interaction* (p. 221-260). Cambridge: Cambridge University Press.

Levinson, S. C. (2006). On the human "interaction engine". In S. C. Enfield N. J. Levinson (Ed.), *Roots of human sociality* (p. 39-69). Oxford: Berg Publishers.

Lev-Ari, S. (2016). How the size of our social network influences our semantic skills. *Cognitive science*, *40*(8), 2050-2064.

Lev-Ari, S., San Giacomo, M., & Peperkamp, S. (2014). The effect of domain prestige and interlocutors' bilingualism on loanword adaptations. *Journal* *of Sociolinguistics*, *18*(5), 658-684.

Leyens, J. P., & Yzerbyt, V. Y. (1992). The ingroup overexclusion effect: Impact of valence and confirmation on stereotypical information search. *European Journal of Social Psychology*, *22*(6), 549-569.

Liu, M., He, X., Rotsthein, P., & Sui, J. (2016). Dynamically orienting your own face facilitates the automatic attraction of attention. *Cognitive neuroscience*, *7*(1-4), 37-44.

Lou, H. C., Luber, B., Crupain, M., Keenan, J. P., Nowak, M., & Kjaer, S. H., T. W. ... Lisanby. (2004). Parietal cortex and representation of the mental self. *Proceedings of the National Academy of Sciences*, *101*(17), 6827-6832.

Ma, D. S., Correll, J., & Wittenbrink, B. (2015). The chicago face database: A free stimulus set of faces and norming data. *Behavior Research Methods*, *47*(4), 1122-1135.

Macmillan, N., & Creelman, C. (1991). *Detection theory: a user's guide.*

Macrae, C. N., Milne, A. B., & Bodenhausen, G. V. (1994). Stereotypes as energy-saving devices: A peek inside the cognitive toolbox. *Journal of personality and Social Psychology*, *66*(1), 37-47.

Mangardich, H., & Sabbagh, M. A. (2018). Children remember words from ignorant speakers but do not attach meaning: evidence from event-related potentials. *Developmental science*, *21*, 2.

Martin, C. D., Garcia, X., Potter, D., Melinger, A., & Costa, A. (2016). Holiday or vacation? the processing of variation in vocabulary across dialects. *Language, Cognition and Neuroscience*, *31*(3), 375-390.

McClelland, J. L., & Elman, J. L. (1986). The trace model of speech perception. *Cognitive psychology*, *18*(1), 1-86.

Meissner, C. A., Brigham, J. C., & Butz, D. A. (2005). Memory for own-and other-race faces: A dual-process approach. *Applied cognitive psychology*, *19*(5), 545-567.

Miller, M. B., Handy, T. C., Cutler, J., Inati, S., & Wolford, G. L. (2001). Brain activations associated with shifts in response criterion on a recognition test. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, *55*(2), 162.

Molenberghs, P., & Louis, W. R. (2018). Insights from fmri studies into ingroup bias. *Frontiers in Psychology*, *9*(1868).

Moradi, Z., Sui, J., Hewstone, M., & Humphreys, G. W. (2015). In-group modulation of perceptual matching. *Psychonomic Bulletin Review*, *22*(5), 1255-1277.

Moradi, Z., Sui, J., Hewstone, M., & Humphreys, G. W. (2017). In-group relevance facilitates learning across existing and new associations. *European Journal of Social Psychology*, *47*(6), 763-774.

Moray, N. (1959). Attention in dichotic listening: Affective cues and the influence of instructions. *Quarterly journal of experimental psychology*, *11*(1), 56-60.

Münster, K., & Knoeferle, P. (2018). Extending situated language comprehension (accounts) with speaker and comprehender characteristics: Toward socially situated interpretation. *Frontiers in Psychology*, *8*(2267).

Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In H. G. W. Besner D. (Ed.), *Basic processes in reading: Visual word recognition* (p. 264-336). Hillsdale, NJ: Erlbaum.

Niedzielski, N. (1999). The effect of social information on the perception of sociolinguistic variables. *Journal of Language and Social Psychology*, *18*(1), 62-85.

Nielsen, K. (2011). Specificity and abstractness of vot imitation. *Journal of Phonetics*, *39*(2), 132-142.

Norris, D. (1994). Shortlist: A connectionist model of continuous speech recognition. *Cognition*, *52*(3), 189-234.

Oostenveld, R., Fries, P., Maris, E., & Schoffelen, J. M. (2011). Fieldtrip: open source software for advanced analysis of meg, eeg, and invasive electrophysiological data. *Computational intelligence and neuroscience*, *2011*(1).

Ostrom, T. M., & Sedikides, C. (1992). Out-group homogeneity effects in natural and minimal groups. *Psychological bulletin*, *112*(3), 536-552.

Park, G., Van Bavel, J. J., Hill, L. K., Williams, D. P., & Thayer, J. F. (2016). Social groups prioritize selective attention to faces: How social identity shapes distractor interference. *PloS one*, *11*(8).

Petersen, S. E., & Posner, M. I. (2012). The attention system of the human brain: 20 years after. *Annual review of neuroscience*, *35*, 73-89.

Pickering, M. J., & Garrod, S. (2007). Do people use language production to make predictions during comprehension? *Trends in cognitive sciences*, *11*(3), 105-110.

Pickering, M. J., & Garrod, S. (2013). An integrated theory of language production and comprehension. *Behavioral and brain sciences*, *36*(4), 329-347.

Platow, M. J., McClintock, C. G., & Liebrand, W. B. (1990). Predicting intergroup fairness and ingroup bias in the minimal group paradigm. *European Journal of Social Psychology*, *20*(3), 221-239.

Posner, M. I. (1980). Orienting of attention. *Quarterly journal of experimental psychology*, *32*(1), 3-25.

Rogers, T. B., Kuiper, N. A., & Kirker, W. S. (1977). Self-reference and the encoding of personal information. *Journal of personality and social psychology*, *35*(9), 677.

Schacter, D. L., Norman, K. A., & Koutstaal, W. (1998). The cognitive neuroscience of constructive memory. *Annual Review of Psychology*, *49*(1), 289-318.

Slessor, G., Phillips, L. H., Ruffman, T., Bailey, P. E., & Insch, P. (2014). Exploring own-age biases in deception detection. *Cognition emotion*, *28*(3), 493-506.

Stephan, W. S., & Stephan, C. W. (2000). An integrated threat theory of prejudice. In S. Oskamp (Ed.), *In reducing prejudice and discrimination* (p. 23-46). Mahwah, NJ: Lawrence Erlbaum Associates.

Stolte, M., Humphreys, G., Yankouskaya, A., & Sui, J. (2017). Dissociating biases towards the self and positive emotion. *The Quarterly Journal of Experimental Psychology*, *70*(6), 1011-1022.

Sui, J., He, X., & Humphreys, G. W. (2012). Perceptual effects of social salience: evidence from self-prioritization effects on perceptual matching. *Journal of Experimental Psychology: Human Perception and Performance*, *38*(5), 1105.

Sui, J., & Humphreys, G. W. (2015). The integrative self: how self-reference integrates perception and memory. *Trends in Cognitive Sciences*, *19*(12), 719-728.

Sui, J., Liu, C. H., Wang, L., & Han, S. (2009). Short article: Attentional orientation induced by temporarily established self-referential cues. *Quarterly Journal of Experimental Psychology*, *62*(5), 844–849.

Sui, J., Rotshtein, P., & Humphreys, G. W. (2013). Coupling social attention to the self forms a network for personal significance. *Proceedings of the National Academy of Sciences*, *110*(19), 7607-7612.

Sui, J., & Zhu, Y. (2005). Five-year-olds can show the self-reference advantage. *International Journal of Behavioral Development*, *29*(5), 382-387.

Sui, J., Zhu, Y., & Han, S. (2006). Self-face recognition in attended and unattended conditions: an event-related brain potential study. *Neuroreport*, *17*(4), 423-427.

Sumner, M., Kim, S. K., King, E., & McGowan, K. B. (2014). The socially weighted encoding of spoken words: A dual-route approach to speech perception. *Frontiers in Psychology*, *4*(1015).

Swann, J. W., Jetten, J., Gómez, , Whitehouse, H., & Bastian, B. (2012). When group membership gets personal: A theory of identity fusion. *Psychological review*, *119*(3), 441.

Swick, D., & Knight, R. T. (1999). Contributions of prefrontal cortex to recognition memory: electrophysiological and behavioral evidence. *Neuropsychology*, *13*(2), 155.

Symons, C. S., & Johnson, B. T. (1997). The self-reference effect in memory: a meta-analysis. *Psychological bulletin*, *121*(3), 371.

Tomasello, M. (2000). First steps toward a usage-based theory of language acquisition. *Cognitive linguistics*, *11*(1), 61-82.

Tomita, H., Ohbayashi, M., Nakahara, K., Hasegawa, I., & Miyashita, Y. (1999). Top-down signal from prefrontal cortex in executive control of memory retrieval. *Nature*, *401*(6754), 699.

Turner, J. C., & Tajfel, H. (1986). The social identity theory of intergroup behavior. *Psychology of intergroup relations*, *5*, 7-24.

Van Bavel, J. J., & Cunningham, W. A. (2012). A social identity approach to person memory: Group membership, collective identification, and social role shape attention and memory. *Personality and Social Psychology Bulletin*, *38*(12), 1566-1578.

Van Bavel, J. J., Packer, D. J., & Cunningham, W. A. (2008). The neural substrates of in-group bias: a functional magnetic resonance imaging investigation. *Psychological Science*, *19*(11), 1131-1139.

Van Bavel, J. J., Swencionis, J. K., O'Connor, R. C., & Cunningham, W. A. (2012). Motivated social memory: Belonging needs moderate the own-group bias in face recognition. *Journal of Experimental Social Psychology*, *48*(3), 707-713.

van Berkum, J., Van den Brink, D., Tesink, C., Kos, M., & Hagoort, P. (2008). The neural integration of speaker and message. *Journal of Cognitive Neuroscience*, *20*(4), 580-591.

Wagner, A. D., Desmond, J. E., Glover, G. H., & Gabrieli, J. D. (1998). Prefrontal cortex and recognition memory. functional-mri evidence for context-dependent retrieval processes. *Brain: a Journal of Neurology*, *121*(10), 1985-2002.

Walker, A., & Hay, J. (2011). Congruence between 'word age'and 'voice age'facilitates lexical access. *Laboratory Phonology*, *2*(1), 219-237.

Wilder, D. A. (1990). Some determinants of the persuasive power of in-groups and out-groups: Organization of information and attribution of independence. *Journal of Personality and Social Psychology*, *59*(6), 1202.

Williams, M. (2001). In whom we trust: Group membership as an affective context for trust development. *Academy of management review*, *26*(3), 377-396.

Windmann, S., Urbach, T. P., & Kutas, M. (2002). Cognitive and neural mechanisms of decision biases in recognition memory. *Cerebral Cortex*, *12*(8), 808-817.

Woźniak, M., Kourtis, D., & Knoblich, G. (2018). Prioritization of arbitrary faces associated to self: An eeg study. *PloS one*, *13*(1).

Xiao, Y. J., Coppin, G., & Van Bavel, J. J. (2016). Perceiving the world through group-colored glasses: A perceptual model of intergroup relations. *Psychological Inquiry*, *27*(4), 255-274.

Yantis, S. (1996). Attentional capture in vision. In A. F. Kramer, M. G. H. Coles, & G. D. Logan (Eds.), *Converging operations in the study of visual selective attention* (p. 45-76). Washington, DC: American Psychological Association.

Yiend, J. (2010). The effects of emotion on attention: A review of attentional processing of emotional information. *Cognition and Emotion*, *24*(1), 3-47.

Yzerbyt, V. Y., Leyens, J. P., & Bellour, F. (1995). The ingroup overexclusion effect: Identity concerns in decisions about group membership. *European Journal of Social Psychology*, *25*(1), 1-16.

Appendices

Appendix 1 - Norming of labels and images for word learning

The aim of the norming study was to select the object images and labels for the word learning task, such that target objects would have at least two equally good labels. To do so, a two-phase norming study was carried out. In each phase, a different group of Dutch university students completed an online survey hosted on LimeSurvey (LimeSurvey Project Team and Schmitz, 2012). Participants received either university-credit or monetary compensation.

1st phase: Written picture naming

Eighteen volunteers (12 female; age mean= 20.89, SD= 1.71) participated. They saw 120 images of uncommon gadgets collected from the internet (e.g., a corn peeler) and provided potential labels for them. After data collection, three Dutch native speakers corrected misspelt forms and excluded redundant forms, as well as regionalisms, non-Dutch forms, and inaccurate names (i.e., if the object was confused with something else). The majority of the elicited labels were nominal compounds with a semantic head (on the right, as common in Dutch), which refers to the action or the purpose of the gadget (e.g., peeler), and another noun (or compound) referring to the object that receives the action (e.g., corn). To ensure competing labels were comparable, labels not following this structure were excluded at this stage.

At the end of this process, 116 out of the initial 120 items were selected. Each gadget elicited between three and fourteen alternative labels.

2nd phase: Goodness-of-fit ratings

During the second phase, 40 university students (33 female; age mean=19.67, SD= 1.96) were recruited for a new online survey. The aim of the survey was to test how well the given labels described the gadgets. The 116 objects selected in the previous phase were pre-randomized and divided into three blocks. Six different versions of the survey were created in order to counterbalance the presentation order of the blocks. Each object was individually presented along with all the given labels. Participants had to rate how well each label described the item on a 1-to-7 scale (1 representing inappropriate, i.e., *ongeschikt*, and 7 being the perfect name, i.e., *perfecte benaming*). Participants were instructed that several labels for a specific item could receive the same value if they were equally good.

Next, paired t-tests were performed on the ratings of all possible pairwise combinations of labels for the same objects. Once we identified pairs which did not statistically differ in ratings (p¿.05), we compared the frequency of their nominal constituents using SUBTLEX-NL (Keuleers, Brysbaert Boris, 2010), to ensure they do not significantly differ in frequency either. For fillers, we selected items with one label that was rated as significantly better than all the other options (paired t-tests, $ps<.05$).

Forty-two target and 43 filler items fulfilled the requirements.

Appendix 2: Perceptual Matching Task results-Chapter 3

Analyses over RTs

Prior to analyses, trials with incorrect responses or with RTs faster than 200 ms or slower than 2100 ms were excluded. For these confirmatory analyses, we selected only matching trials (i.e., in which the logo of the university was displayed with the associated geometrical shape) which referred to the in-group university and the out-group university used in the study (i.e., Groningen University). We then performed an outlier removal procedure by removing trials with RTs 2.5 SDs or higher from the mean per condition, per participant. The resulting dataset was analyzed using linear mixed-effect model in which log(10)-transformed RTs were predicted by the fixed effect for Group Membership (In-group vs Out-group, reference level: In-group). We added per-participant random intercept and by-participant random slope for Group Membership. Results confirmed the usual trend for this task: participants were faster at recognizing in-group-related associations (M=672 ms, SD=150) than out-group-related associations (M=735 ms, SD=176) (beta=0.08, SE=0.01, t=7.37, $p<.0001$).

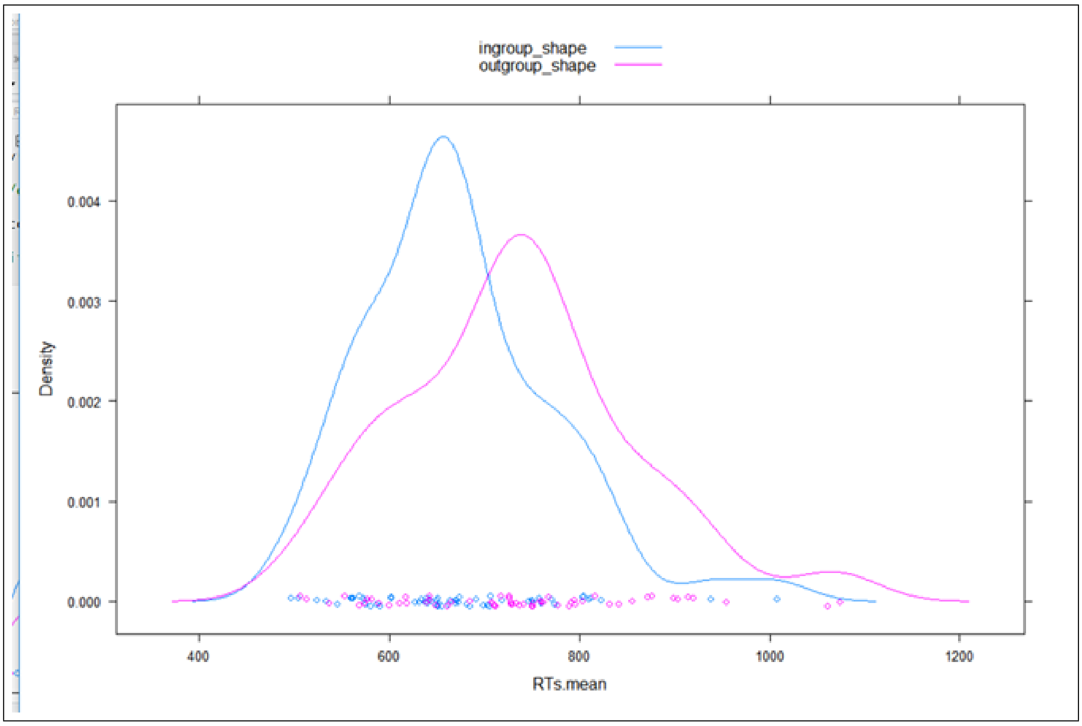

Figure 6.1: Density plot showing RT means per group membership.

Analyses over Accuracy

Prior to analyses, we selected only matching trials (i.e., in which the logo of the university was displayed with the associated geometrical shape) which referred to the in-group university and the out-group university used in the study (i.e., Groningen University). Responses' accuracy on a single trial were analyzed using logistic mixed-effect model with a fixed effect for Group Membership (In-group vs Out-group, reference level: In-group). We added per-participant random intercept and by-participant random slope for Group Membership. Results confirmed the usual trend for this task: participants were better at recognizing in-group-related associations (M=96%, SD=20) than out-group-related associations (M=89%, SD=31) (beta=-1.18, SE=0.29, t=-4.73, p<.0001).

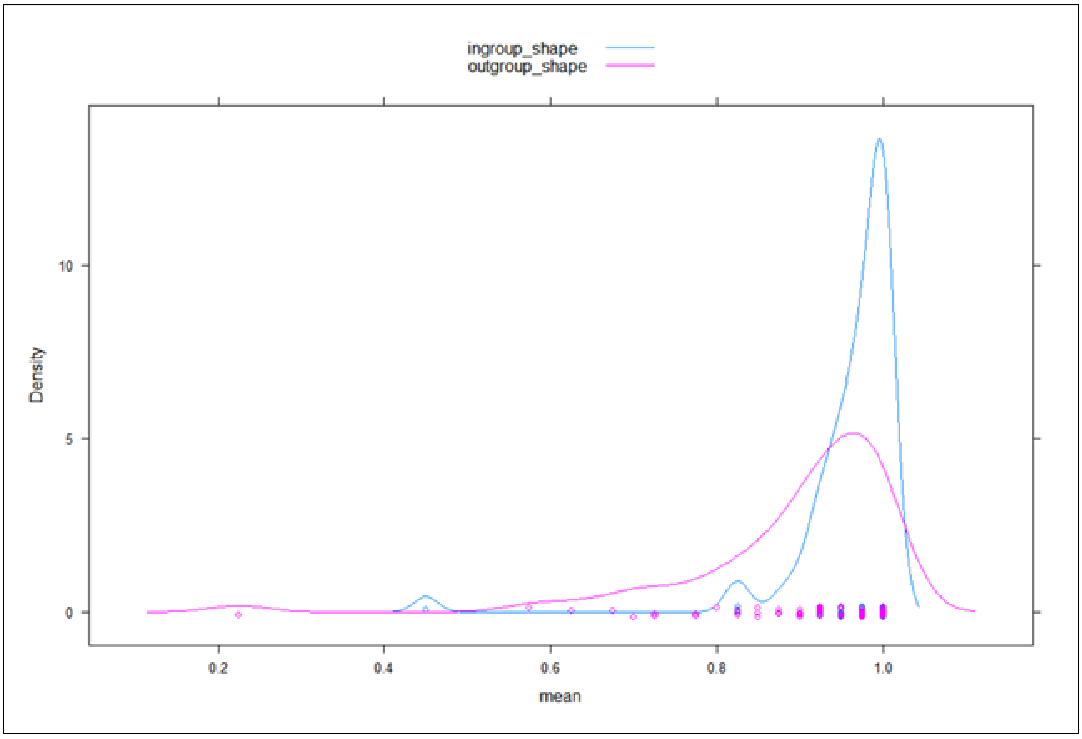

Figure 6.2: Density plot showing accuracy rates per group membership.

Distribution of the individual in-group bias

NB.The in-group bias measure is z-scored, therefore the values around 0 do not correspond to actual 0.

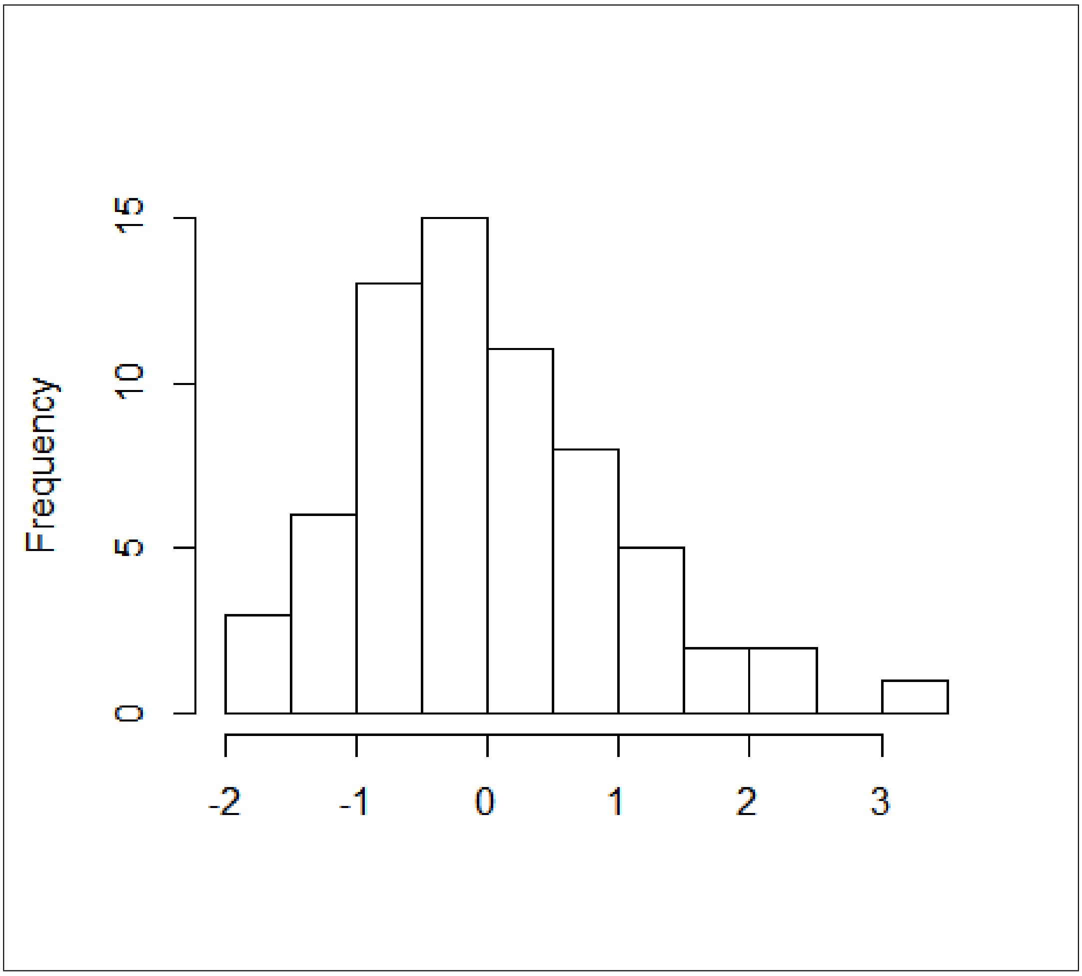

Figure 6.3: Distribution of the individual in-group bias measure.

List of selected target items, in pairs

boterhamhouder	sandwichhouder
tandpastaperser	tubeperser
tassendrager	handvatdrager
kabelhouders	snoerhouders
bellenblaasborstel	afwasbellenblaas
citroenspuit	citroenspray
kruidenschaar	kruidenknipper
cakesnijder	cakezaag
melk-klopper	melkschuimer
hamburgervorm	hamburgerpers
avocadolepel	avocadosnijder
kruidenschraper	kruidenstripper
bananenschiller	bananenpeller
citroenknijper	citroenperser
citroendispenser	citroenpipet
drumstickvingers	drumstickklauwen
hotdogmaker	hotdogmachine
plantendouche	flesgieter
lepelklem	spatelhouder
slaaphelm	reiskussen
snorkelband	duikmasker
bladzijdenspreider	paginaspreider
pantoffelhouder	slipperhouder
lampsloffen	nachtsloffen
handvliezen	vingervliezen
bootfiets	fietsschip
verfpipet	verfhersteller
ballenschieter	ballengooier
maiskolfpeller	maiskolfrasp
dweilsloffen	dweilslippers
eidooierscheider	eidooieropvanger
babydraagband	babydraagzak
avocadoschiller	avocadopeller
maisgrill	maisklem
citruspeller	citrus- schiller
glazuurpen	decoratiepen
ijsesmaker	ijsmallen
deegschep	deegschraper
drinkkraantje	mondspoeler
oliedispenser	sauskwast
kraanadapter	straalverbreder
saladeschaar	groenteknipper

English summary

Human communication is intrinsically social: people interact with one another to share thoughts and beliefs and talk about the surrounding world.

Yet, many theoretical models that describe how people produce and comprehend language assume that only the linguistic signal is salient, and that other aspects of interaction, such as the identity of the interlocutor, are irrelevant. In these accounts, such factors do not affect interpretation and are not stored in memory. But research provides us with many examples in which the identity of the speaker proves surprisingly salient for how language is processed and remembered. For example, people can be tricked into hearing stereotypical pronunciations of certain sounds when they are told that the speaker they will listen to is from one region or another. Similarly, it was shown that listeners expect different words depending on the identity of the speaker. For instance, people listening to a child expect to hear words that are more commonly produced by children, but not from adults. If these expectations are then violated by the child commenting on her habit of drinking *wine* before going to sleep, listeners' brains have a hard time coping with it.

Taken together, available evidence confirms that aspects outside of the linguistic signal, such as the identity of the speaker, are processed and stored in memory when comprehending language, and they can guide inferences in future similar experiences. What is still unknown though is what the precise mechanisms underlying this phenomenon are and how they interact with other aspects of social nature that are known to influence human interactions.

One important factor that influences social dynamics is whether people regard others as being members of their own social group (i.e., in-group member) or not (i.e., out-group member). People who are students at the same university, who have the same gender, or support the same sports club can, for instance, be considered to belong to the same social group. Many studies have shown that individuals have a strong bias towards members of their own group (i.e., in-group bias) which lead them to prefer in-group members over out-group members, and to better process and remember information about the former compared to the latter. Yet, we don't know whether these preferences can also influence how language produced by in-group and out-group speakers is processed.

The central focus of this thesis was to explore the possibility that **people may treat language differently depending on the group membership status of the speaker**. To test this account, I used a word learning paradigm in three different experiments. In this paradigm, participants heard speakers, who were students from either their own university (i.e., in-group speakers) or a different one (i.e., out-group speakers), that labelled novel, uncommon gadgets (e.g., spatula-holder). In this way, I could test for effects of group membership on learning new words, controlling for their prior exposure. I tested a series of hypotheses:

1. When people learn labels for new, uncommon gadgets, they pay more attention to the speaker when she is an in-group member than when she is an out-group one. This differential attention can be later reflected in better **source memory** (i.e., remembering who said what) for the in-group words.

2. The in-group advantage during learning new words can influence how easily people comprehend those words in later encounters, perhaps be-

cause words learned from in-group members receive a **processing boost**, as compared to those learned from out-group ones.

3. People may differ in how strongly they prefer their in-group vs out-group membership. Therefore, individuals with a stronger in-group bias should show larger source memory and processing advantages for in-group vs out-group words.

4. In-group biases may be due to the **allocation of extra attention** when processing information that is related to the in-group membership.

To test this last hypothesis, I used a different task where human faces, belonging to in-group and out-group members, preceded the display on screen of a target that participants had to categorize as quickly and as accurately as possible. If in-group faces receive more attention than out-group faces, they should slow down people's responses when the target appears at a different screen location. On the other hand, in-group faces should speed responses up when the target appears at the same location.

The results of my experiments show that indeed people that have a strong in-group bias tended to better remember the exact in-group speaker from whom they learned new words. When the speaker was an out-group member these participants with a stronger in-group bias did not remember the exact speakers but only the out-group membership. This finding supports the first and the third hypotheses about better source memory for words spoken by in-group members and in-group individual differences. On the other hand, I did not find convincing evidence in favor of the other two hypotheses about processing benefit and extra attention.

Taken together, these findings suggest that people may store information differently when listening to in-group speakers. However, such initial differences may not depend on the allocation of extra attention, and they may not affect how easily people comprehend newly learned words in the future. Hence, the theoretical story in which only the signal matters is not the whole story. Social characteristics of the speakers can influence how we treat new information and what we store in memory. These findings have implications also for child language learning, but also for social politics and marketing fields.

Nederlandse samenvatting

Communicatie tussen mensen is intrinsiek sociaal. Mensen gaan gesprekken met elkaar aan om gedachtes en overtuigingen te kunnen delen en om over de wereld om zich heen te praten.

Veel theoretische modellen die omschrijven hoe mensen taal produceren en begrijpen, nemen echter aan dat alleen het taalkundige signaal betekenisvol is. Andere aspecten van de interactie, zoals de identiteit van de gesprekspartner, beschouwen ze als irrelevant. Volgens deze theorieën hebben factoren als identiteit geen invloed op de interpretatie van een boodschap en worden deze sociale kenmerken niet in ons geheugen opgeslagen. Uit onderzoek blijkt echter dat de identiteit van de spreker juist erg belangrijk is voor de manier waarop taal wordt verwerkt en onthouden. Mensen kunnen bijvoorbeeld gaan denken dat ze de stereotype uitspraak van een bepaalde klank horen wanneer hen verteld wordt dat de spreker uit een bepaalde regio komt. Daarnaast werd eerder aangetoond dat luisteraars verschillende woorden verwachten op basis van de identiteit van de spreker. Mensen die luisteren naar een kind, verwachten bijvoorbeeld woorden die vaker door kinderen en minder vaak door volwassenen worden gebruikt. Als deze verwachtingen worden doorbroken, bijvoorbeeld door een kind dat vertelt over haar gewoonte om *wijn* te drinken voor het slapengaan, kost het onze hersenen meer moeite om dit te verwerken.

Samengevat suggereert beschikbaar onderzoek dat aspecten buiten het linguïstische signaal, zoals de identiteit van de spreker, worden verwerkt en opgeslagen in het geheugen terwijl we taal interpreteren. Dit kan interpretaties in toekomstige, vergelijkbare situaties beïnvloeden. Het is echter nog onbekend wat de precieze onderliggende mechanismes zijn die dit fenomeen mogelijk maken en hoe de wisselwerking is tussen deze mechanismes en andere sociale aspecten die de menselijke interactie beïnvloeden.

Een belangrijke factor die effect heeft op sociale dynamiek, is of mensen elkaar zien als lid van hun eigen sociale groep (*in-group* leden), of niet (*out-group* leden). Studenten aan dezelfde universiteit, mensen van hetzelfde geslacht, of fans van dezelfde sportclub bijvoorbeeld, kunnen we beschouwen als leden van dezelfde groep. Veel studies tonen aan dat individuen een sterke voorkeur hebben voor leden van hun eigen groep (in-group bias). Hierdoor verwerken en onthouden ze makkelijker informatie over leden van hun eigen groep dan over mensen uit de out-group. We weten echter nog niet of deze voorkeur ook de verwerking van taal die geproduceerd wordt door in-group en out-group sprekers kan beïnvloeden.

Het doel van dit proefschrift was om te onderzoeken **of de manier waarop mensen omgaan met taal afhankelijk is van de groep waar de spreker toe behoort**. Om dit te testen, gebruikte ik drie verschillende experimenten waarin nieuwe woorden geleerd moesten worden. Proefpersonen luisterden naar sprekers die aan hun eigen universiteit (in-group), of aan een andere universiteit (out-group) studeerden. De sprekers gaven een naam aan nieuwe, bijzondere voorwerpen (bijv. spatelhouder). Omdat ik in de experimenten kon bepalen welke woorden door welke spreker werden aangeleerd, kon ik de invloed van sociale groep op het leren van nieuwe woorden onderzoeken. Ik testte verschillende hypotheses:

1. Mensen luisteren aandachtiger naar een speker die tot hun in-group behoort dan naar iemand die tot de out-group behoort wanneer ze nieuwe namen voor ongebruikelijke voorwerpen leren. Het verschil in aandacht

uit zich later in een betere **source memory** (wie heeft welke woorden benoemd?) voor in-group woorden.

2. 2. Het in-group voordeel dat proefpersonen hebben bij het leren van nieuwe woorden heeft invloed op hoe makkelijk ze de woorden begrijpen wanneer ze deze woorden later nog een keer tegenkomen. Misschien komt dit doordat woorden die ze leren van in-group leden een **verwerkingsboost** krijgen ten opzichte van woorden die ze leren van mensen die behoren tot een out-group.

3. Het is mogelijk dat er verschillen zijn in hoe sterk mensen een voorkeur hebben voor in-group sprekers ten opzichte van out-group sprekers. Individuen met een sterkere voorkeur voor de in-group, zouden een betere source memory voor in-group woorden vs. out-group woorden moeten hebben en zouden het makkelijker moeten vinden om in-group woorden te verwerken.

4. 4. Voorkeuren voor de in-group zouden kunnen ontstaan omdat we **aandachtiger** zijn tijdens het verwerken van informatie die te maken heeft met de in-group.

Om deze laatste hypothese te testen, gebruikte ik een andere taak. Hierbij liet ik gezichten van mensen uit de in- of out-group zien. Daarna verscheen er een symbool op het scherm, op dezelfde of een andere positie op het scherm. Proefpersonen moesten zo snel en zo accurraat mogelijk selecteren welk symbool dit was (2 opties). Als in-group gezichten meer aandacht krijgen dan out-group gezichten, zou de proefpersoon sneller moeten kunnen reageren als het symbool op dezelfde postie verschijnt als het in-group gezicht. Omgekeerd zouden in-group gezichten de reactie moeten vertragen als het symbool op een andere positie dan het gezicht verschijnt. Out-group gezichten zouden minder effect op de reactiesnelheid moeten hebben.

De resultaten van mijn experimenten tonen aan dat mensen die een sterke voorkeur hebben voor in-group leden, inderdaad makkelijker konden onthouden van welke in-group spreker ze een nieuw woord hadden geleerd. Als de spreker lid was van een out-group, wisten de deelnemers met een sterke in-group voorkeur niet meer van welke spreker ze een woord hadden gehoord, maar alleen dat het iemand was die tot de out-group behoorde. Deze uitkomst ondersteunt de eerste hypothese over betere source memory voor woorden die aangeleerd worden door in-group leden en de derde hypothese over individuele verschillen in voorkeur voor de in-group. Ik vond echter geen overtuigend bewijs ter ondersteuning van de andere hypotheses over verwerkingsvoordelen en extra aandacht.

Samen suggereren deze bevindingen dat we informatie wellicht anders opslaan wanneer we naar in-group sprekers luisteren. Dit is echter misschien niet afhankelijk van hoeveel (extra) aandacht we aan de spreker besteden, en heeft mogelijk geen invloed op hoe makkelijk mensen nieuwe woorden begrijpen.

Een theoretisch verhaal waarin alleen het taalkundige signaal telt, is dus niet het volledige verhaal. De sociale eigenschappen van sprekers kunnen beïnvloeden hoe we nieuwe informatie verwerken en wat we opslaan in ons geheugen. Deze bevindingen hebben ook implicaties voor onderzoek naar hoe kinderen taal leren, sociale politicologie en marketing.

Acknowledgements

None of the work written in this dissertation would have been possible without the support, patience and love of the amazing people I had the luck to walk with during these last four years. One feeling always accompanied me on the way: deep gratitude.

First of all, I would like to thank my promoter, Antje, who gave me the chance to explore my research interests, and my limits, and who never denied me her support in whichever new challenge I pushed myself into.

Secondly, I would like to thank all my supervisors. To Shiri, for teaching me to be precise and never approximate in my thoughts and my reasoning. To Alastair, for reminding me to be true and compassionate towards myself. To Monique, for being critical, but in a constructive way, professional, but also sweet, and for being such a great help in the ups and downs of the last year. I am really grateful for having had you!

I would like to thank all the members of my reading committee, Asli, Jos and Atsuko, for reading this thesis and sending encouraging and positive words back. Thank you!

A special thank you goes also to the IMPRs people, Kevin in primis, for giving me the possibility to grow as a researcher, and as a human being.

Then, I would like to thank all my beautiful colleagues and friends without whom I wouldn't have finished this thesis. Thank you, Nina, for becoming a true friend, someone I could turn to in any circumstances, for listening to me and teaching me how to listen to my heart. Thank you, Markus, for endless conversations about every single topic I can think of and for walking my mind around in the forest. Thank you, Amie, for being, from the beginning to the end, such a sweet office-, RLadies-, climbing-, booksharing-, mindsharing-mate. Thank you, Joe (and Mischa, of course) for your friendship, the great food we cooked and ate together, the laughters and the bitter gossip. Thank you, Eirini, for your kind soul and the endless chats about culture, literature, RLadies, open-science, long-distance relationships and future jobs. Grazie, Federica, per le mille sigarette scroccate e i caffe' buoni, per credere in un'Italia migliore ed aiutarmi con il mio italiano pessimo. Thank you, Laurel, for being always available, encouraging, sweet and helpful in every single matter. Thank you, Caitlin, for your kindness and your deep heart and for helping me program my very first experiment. Thank you, Aitor, for bringing fresh air into our evenings and in my mind. Thank you, Annelies, for helping out with pretty much everything and sharing entertaining lunch breaks with joy and enthusiasm. Thank you to all my POL colleagues, Jeroen, Merel, Merel, Saoradh, Miguel, Hans Rutger, Florian, Falk, Andrea, Phillip, Sophie, for being always supportive and available whenever it was needed.

Thank you to my non-POL friends, Yingying, Julia, Sophie, Laura, Elliot, Marvin, Birgit, Anne, Hannah, Linda, Lot, Ellen, Fabian, Ine, Midas, Ella, James, Xiaochen, and too many others, who made my PhD life much easier with their help, coffee and support.

Special thank you, René, for your wonderful friendship, and all the time we spent together chatting, gossiping, playing music and having coffees.

Grazie ai miei italiani del cuore, Paolo e Arianna, ed Ileana (e Coen) che mi hanno fatto riscoprire la bellezza di essere italiana, e, soprattutto, di mangiare cibo italiano.

Grazie, Valeria, per essere *tu* e per avermi accompagnata in questa avventura colossale con filosofia, jazz e amicizia.

Gracias a Amaia, Sarita, Varun y Gabi, por haberme mantenido vivo el Castellano y por haber alimentado mi alma con charlas, amistad, bromas, paseos, capuccinos, carreras, comida rica...vamos, por todo!

Thank you to Cherry and the cakes, Tim, Christian, Giel, Dan, Tom, Ilse, René, Fenja, Elliot and Greta, for having enriched my life with good vibes and great concerts!

Thank you to all my climbing friends who were always there to support my next step :)

Finally, thank you to my wonderful, incredible friends and paranymphs, Limor, Greta, and Giuseppe. Thank you, Limor, for your unconditional love and for sharing with me your heart and your brain. I couldn't have wished for better friend and supporter throughout all these years.

Thank you, Greta, for being an amazing friend, office-mate and music-mate, and for sharing smiles, tears and laughters together. Your friendship gave the daily strength I needed to reach the end.

Grazie, Giusep, di tutto l'amore che in questi anni mi ha riempito il cuore.

Grazie alla mia famiglia. A Simone, per aver speso un intero viaggio in macchina Ciampino-Fondi ascoltando i dettagli preliminari di questo progetto, e per avermi sempre ascoltata con interesse. A Silvia, per essersi armata di santa pazienza e aver disegnato e realizzato la miglior cover del mondo per questa tesi. Ed in particolare alla mia mamma, per avermi sempre sostenuta ed amata in tutte le pazze avventure della vita, compresa questa qui. Senza il tuo esempio e il tuo coraggio, mammina mia, non avrei potuto farcela.

Curriculum Vitae

Sara Iacozza was born in 1988 in Rome, Italy. She obtained her bachelor's degree in Foreign Languages and Cultures from Universitá degli Studi di Urbino (Italy) in 2011. This was followed by a Master's degree in English Literature and Linguistics at the Universidad de Granada (Spain) in 2013, and another Master's degree in Cognitive Neuroscience of Language at the Basque Center on Cognition, Brain and Language in San Sebastian (Spain) in 2015. In October 2015, Sara began her PhD project in the Psychology of Language Department at the Max Planck Institute for Psycholinguistics.

Sara is currently working as data analyst for Do Big Good (Seattle).

Publications

Iacozza, S., Costa, A., Duñabeitia, J. A. (2017). What do your eyes reveal about your foreign language? Reading emotional sentences in a native and foreign language. *PloS one, 12*(10), e0186027.

Iacozza, S., Meyer, A. S., Lev-Ari, S. (2019). How in-group bias influences source memory for words learned from in-group and out-group speakers. *Frontiers in human neuroscience, 13*, 308.

Iacozza, S., Meyer, A. S., Lev-Ari, S. (2019). How in-group bias influences the level of detail of speaker-specific information encoded in novel lexical representations. Journal of experimental psychology. *Learning, memory, and cognition.*

MPI Series in Psycholinguistics

1. The electrophysiology of speaking: Investigations on the time course of semantic, syntactic, and phonological processing. *Miranda I. van Turennout*
2. The role of the syllable in speech production: Evidence from lexical statistics, metalinguistics, masked priming, and electromagnetic midsagittal articulography. *Niels O. Schiller*
3. Lexical access in the production of ellipsis and pronouns. *Bernadette M. Schmitt*
4. The open-/closed class distinction in spoken-word recognition. *Alette Petra Haveman*
5. The acquisition of phonetic categories in young infants: A self-organising artificial neural network approach. *Kay Behnke*
6. Gesture and speech production. *Jan-Peter de Ruiter*
7. Comparative intonational phonology: English and German. *Esther Grabe*
8. Finiteness in adult and child German. *Ingeborg Lasser*
9. Language input for word discovery. *Joost van de Weijer*
10. Inherent complement verbs revisited: Towards an understanding of argument structure in Ewe. *James Essegbey*
11. Producing past and plural inflections. *Dirk J.Janssen*
12. Valence and transitivity in Saliba: An Oceanic language of Papua New Guinea. *Anna Margetts*
13. From speech to words. *Arie H. van der Lugt*
14. Simple and complex verbs in Jaminjung: A study of event categorisation in an Australian language. *Eva Schultze-Berndt*
15. Interpreting indefinites: An experimental study of children's language comprehension. *Irene Krämer*
16. Language-specific listening: The case of phonetic sequences. *Andrea Christine Weber*

17. Moving eyes and naming objects. *Femke Frederike van der Meulen*
18. Analogy in morphology: The selection of linking elements in Dutch compounds. *Andrea Krott*
19. Morphology in speech comprehension. *Kerstin Mauth*
20. Morphological families in the mental lexicon. *Nivja Helena de Jong*
21. Fixed expressions and the production of idioms. *Simone Annegret Sprenger*
22. The grammatical coding of postural semantics in Goemai (a West Chadic language of Nigeria. *Birgit Hellwig*
23. Paradigmatic structures in morphological processing: Computational and cross-linguistic experimental studies. *Fermín Moscoso del Prado Martín*
24. Contextual influences on spoken-word processing: An electrophysiological approach. *Danielle van den Brink*
25. Perceptual relevance of prevoicing in Dutch. *Petra Martine van Alphen*
26. Syllables in speech production: Effects of syllable preparation and syllable frequency. *Joana Cholin*
27. Producing complex spoken numerals for time and space. *Marjolein Henriëtte Wilhelmina Meeuwissen*
28. Morphology in auditory lexical processing: Sensitivity to fine phonetic detail and insensitivity to suffix reduction. *Rachèl Jenny Judith Karin Kemps*
29. At the same time...: The expression of simultaneity in learner varieties. *Barbara Schmiedtová*
30. A grammar of Jalonke argument structure. *Friederike Lüpke*
31. Agrammatic comprehension: An electrophysiological approach. *Marijtje Elizabeth Debora Wassenaar*
32. The structure and use of shape-based noun classes in Miraña (North West Amazon). *Frank Seifart*
33. Prosodically-conditioned detail in the recognition of spoken words. *Anne Pier Salverda*
34. Phonetic and lexical processing in a second language. *Mirjam Elisabeth Broersma*
35. Retrieving semantic and syntactic word properties: ERP studies on the time course in language comprehension. *Oliver Müller*

36. Lexically-guided perceptual learning in speech processing. *Frank Eisner*
37. Sensitivity to detailed acoustic information in word recognition. *Keren Batya Shatzman*
38. The relationship between spoken word production and comprehension. *Rebecca Özdemir*
39. Disfluency: Interrupting speech and gesture. *Mandana Seyfeddinipur*
40. The acquisition of phonological structure: Distinguishing contrastive from non-constrative variation. *Christiane Dietrich*
41. Cognitive cladistics and the relativity of spatial cognition. *Daniel Haun*
42. The acquisition of auditory categories. *Martijn Bastiaan Goudbeek*
43. Affix reduction in spoken Dutch: Probabilistic effects in production and perception. *Mark Plumaekers*
44. Continuous-speech segmentation at the beginning of language acquisition: Electrophysiological evidence. *Valesca Madalla Kooijman*
45. Space and iconicity in German sign language (DGS). *Pamela M. Perniss*
46. On the production of morphologically complex words with special attention to effects of frequency. *Heidrun Bien*
47. Crosslinguistic influence in first and second languages: Convergence in speech and gesture. *Amanda Brown*
48. The acquisition of verb compounding in Mandarin Chinese. *Jidong Chen*
49. Phoneme inventories and patterns of speech sound perception. *Anita Eva Wagner*
50. Lexical processing of morphologically complex words: An information-theoretical perspective. *Victor Kuperman*
51. A grammar of Savosavo: A Papuan language of the Solomon Islands. *Claudia Ursula Wegener*
52. Prosodic structure in speech production and perception. *Claudia Kuzla*
53. The acquisition of finiteness by Turkish learners of German and Turkish learners of French: Investigating knowledge of forms and functions in production and comprehension. *Sarah Schimke*
54. Studies on intonation and information structure in child and adult German. *Laura de Ruiter*
55. Processing the fine temporal structure of spoken words. *Eva Reinisch*

56. Semantics and (ir)regular inflection in morphological processing. *Wieke Tabak*
57. Processing strongly reduced forms in casual speech. *Susanne Brouwer*
58. Ambiguous pronoun resolution in L1 and L2 German and Dutch. *Miriam Ellert*
59. Lexical interactions in non-native speech comprehension: Evidence from electroencephalography, eye-tracking, and functional magnetic resonance imaging. *Ian FitzPatrick*
60. Processing casual speech in native and non-native language. *Annelie Tuinman*
61. Split intransitivity in Rotokas, a Papuan language of Bougainville. *Stuart Payton Robinson*
62. Evidentiality and intersubjectivity in Yurakaré: An interactional account. *Sonja Gipper*
63. The influence of information structure on language comprehension: A neurocognitive perspective. *Lin Wang*
64. The meaning and use of ideophones in Siwu. *Mark Dingemanse*
65. The role of acoustic detail and context in the comprehension of reduced pronunciation variants. *Marco van de Ven*
66. Speech reduction in spontaneous French and Spanish. *Francisco Torreira*
67. The relevance of early word recognition: Insights from the infant brain. *Caroline Mary Magteld Junge*
68. Adjusting to different speakers: Extrinsic normalization in vowel perception. *Matthias Johannes Sjerps*
69. Structuring language: Contributions to the neurocognition of syntax. *Katrien Rachel Segaert*
70. Infants' appreciation of others' mental states in prelinguistic communication: A second person approach to mindreading. *Birgit Knudsen*
71. Gaze behavior in face-to-face interaction. *Federico Rossano*
72. Sign-spatiality in Kata Kolok: How a village sign language of Bali inscribes its signing space. *Connie de Vos*
73. Who is talking? Behavioural and neural evidence for norm-based coding in voice identity learning. *Attila Andics*
74. Lexical processing of foreign-accented speech: Rapid and flexible adaptation. *Marijt Witteman*

75. The use of deictic versus representational gestures in infancy. *Daniel Puccini*

76. Territories of knowledge in Japanese conversation. *Kaoru Hayano*

77. Family and neighbourhood relations in the mental lexicon: A cross-language perspective. *Kimberley Mulder*

78. Contributions of executive control to individual differences in word production. *Zeshu Shao*

79. Hearing speech and seeing speech: Perceptual adjustments in auditory-visual processing. *Patrick van der Zande*

80. High pitches and thick voices: The role of language in space-pitch associations. *Sarah Dolscheid*

81. Seeing what's next: Processing and anticipating language referring to objects. *Joost Rommers*

82. Mental representation and processing of reduced words in casual speech. *Iris Hanique*

83. The many ways listeners adapt to reductions in casual speech. *Katja Pöllmann*

84. Contrasting opposite polarity in Germanic and Romance languages: Verum Focus and affirmative particles in native speakers and advanced L2 learners. *Giuseppina Turco*

85. Morphological processing in younger and older people: Evidence for flexible dual-route access. *Jana Reifegerste*

86. Semantic and syntactic constraints on the production of subject-verb agreement. *Alma Veenstra*

87. The acquisition of morphophonological alternations across languages. *Helen Buckler*

88. The evolutionary dynamics of motion event encoding. *Annemarie Verkerk*

89. Rediscovering a forgotten language. *Jiyoun Choi*

90. The road to native listening: Language-general perception, language-specific input. *Sho Tsuji*

91. Infants'understanding of communication as participants and observers. *Gudmundur Bjarki Thorgrímsson*

92. Information structure in Avatime. *Saskia van Putten*

93. Switch reference in Whitesands. *Jeremy Hammond*
94. Machine learning for gesture recognition from videos. *Binyam Gebrekidan Gebre*
95. Acquisition of spatial language by signing and speaking children: A comparison of Turkish sign language (TID) and Turkish. *Beyza Sumer*
96. An ear for pitch: On the effects of experience and aptitude in processing pitch in language and music. *Salomi Savvatia Asaridou*
97. lncrementality and Flexibility in Sentence Production. *Maartje van de Velde*
98. Social learning dynamics in chimpanzees: Reflections on (nonhuman) animal culture. *Edwin van Leeuwen*
99. The request system in Italian interaction. *Giovanni Rossi*
100. Timing turns in conversation: A temporal preparation account. *Lilla Magyari*
101. Assessing birth language memory in young adoptees. *Wencui Zhou*
102. A social and neurobiological approach to pointing in speech and gesture. *David Peeters*
103. Investigating the genetic basis of reading and language skills. *Alessandro Gialluisi*
104. Conversation electrified: The electrophysiology of spoken speech act recognition. *Rósa Signý Gisladottir*
105. Modelling multimodal language processing. *Alastair Charles Smith*
106. Predicting language in different contexts: The nature and limits of mechanisms in anticipatory language processing. *Florian Hintz*
107. Situational variation in non-native communication. *Huib Kouwenhoven*
108. Sustained attention in language production. *Suzanne Jongman*
109. Acoustic reduction in spoken-word processing: Distributional, syntactic, morphosyntatic, and orthographic effects. *Malte Viebahn*
110. Nativeness, dominance, and the flexibility of listening to spoken language. *Laurence Bruggeman*
111. Semantic specificity of perception verbs in Maniq. *Ewelina Wnuk*
112. On the identification of FOXP2 gene enhancers and their role in brain development. *Martin Becker*

113. Events in language and thought: The case of serial verb constructions in Avatime. *Rebecca Defina*

114. Deciphering common and rare genetic effects on reading ability. *Amaia Carrión Castillo*

115. Music and language comprehension in the brain. *Richard Kunert*

116. Comprehending Comprehension: Insights from neuronal oscillations on the neuronal basis of language. *Nietzsche H.L. Lam*

117. The biology of variation in anatomical brain asymmetries. *Tulio Guadalupe*

118. Language processing in a conversation context. *Lotte Schoot*

119. Achieving mutual understanding in Argentine Sign Language. *Elizabeth Manrique*

120. Talking Sense: the behavioural and neural correlates of sound symbolism. *Gwilym Lockwood*

121. Getting under your skin: The role of perspective and simulation of experience in narrative comprehension. *Franziska Hartung*

122. Sensorimotor experience in speech perception. *Will Schuerman*

123. Explorations of beta-band neural oscillations during language comprehension: Sentence processing and beyond. *Ashley Lewis*

124. Influences on the magnitude of syntactic priming. *Evelien Heyselaar*

125. Lapse organization in interaction. *Elliott Hoey*

126. The processing of reduced word pronunciation variants by natives and foreign language learners: Evidence from French casual speech. *Sophie Brand*

127. The neighbors will tell you what to expect: Effects of aging and predictability on language processing. *Cornelia Moers*

128. The role of voice and word order in incremental sentence processing. *Sebastian Sauppe*

129. Learning from the (un)expected: Age and individual differences in statistical learning and perceptual learning in speech. *Thordis Neger*

130. Mental representations of Dutch regular morphologically complex neologisms. *Laura de Vaan*

131. Speech production, perception, and input of simultaneous bilingual preschoolers: Evidence from voice onset time. *Antje Stoehr*

132. A holistic approach to understanding pre-history. *Vishnupriya Kolipakam*

133. Characterization of transcription factors in monogenic disorders of speech and language. *Sara Busquets Estruch*

134. Indirect request comprehension in different contexts. *Johanne Tromp*

135. Envisioning Language - An Exploration of Perceptual Processes in Language Comprehension. *Markus Ostarek*

136. Listening for the WHAT and the HOW: Older adults' processing of semantic and affective information in speech. *Juliane Kirsch*

137. Let the agents do the talking: on the influence of vocal tract anatomy on speech during ontogeny and glossogeny. *Rick Janssen*

138. Age and hearing loss effects on speech processing. *Xaver Koch*

139. Vocabulary knowledge and learning: Individual differences in adult native speakers. *Nina Mainz*

140. The face in face-to-face communication: Signals of understanding and non-understanding. *Paul Hömke*

141. Person reference and interaction in Umpila/Kuuku Ya'u narrative. *Clair Hill*

142. Beyond the language given: The neurobiological infrastructure for pragmatic inferencing. *Jana Bašnáková*

143. From Kawapanan to Shawi: Topics in language variation and change. *Luis Miguel Rojas-Berscia*

144. On the oscillatory dynamics underlying speech-gesture integration in clear and adverse listening conditions. *Linda Drijvers*

145. Understanding temporal overlap between production and comprehension. *Amie Fairs*

146. The role of exemplars in speech comprehension. *Annika Nijveld*

147. A network of interacting proteins disrupted in language-related disorders. *Elliot Sollis*

148. Fast speech can sound slow: Effects of contextual speech rate on word recognition. *Merel Maslowski*

149. Reason-giving in everyday activities. *Julija Baranova*

150. Speech planning in dialogue - Psycholinguistic studies of the timing of turn taking. *Mathias Barthel*

151. The role of neural feedback in language unification: How awareness affects combinatorial processing. *Valeria Mongelli*
152. Exploring social biases in language processing. *Sara Iacozza*