Morphology in Minds with Machines: Representational Similarity Analysis of HD-EEG and Models of Morphological Decomposition in Word Processing

by

Tyson Jordan

(Under the Direction of Dustin Chacón and Jin Sun)

Abstract

Neuroimaging studies provide mixed evidence for morphological decomposition in the brain, with results spanning reading and listening. This study examines how the brain processes complex, pseudocomplex, and simplex words using high-density EEG (HD-EEG) with source reconstruction. Subjects completed a lexical decision task in four modalities: text, audio, audio-video (AV), and silent video. We applied a representational similarity analysis (RSA) to compare neural activation with theoretical models of morphological structure and semantic representations derived from distributional word embeddings. Our findings suggest the activation of (pseudo)stems during reading, supporting a model of full decomposition. We report additional effects of decomposition in the audio and AV conditions, following the morphological disambiguating point. No morphological effects were observed during lipreading.

Index words: Computational neuroscience, Cognitive neuroscience, Neurolinguistics, Morphology, Morphological Decomposition

Morphology in Minds with Machines: Representational Similarity Analysis of HD-EEG and Models of Morphological Decomposition in Word Processing

by

Tyson Jordan

B.A. Cognitive Science, University of Georgia, 2023

B.S. Psychology, University of Georgia, 2023

A Thesis Submitted to the Graduate Faculty

of The University of Georgia in Partial Fulfillment

of the

Requirements for the Degree

Master of Science

Athens, Georgia

2025

Morphology in Minds with Machines: Representational Similarity Analysis of HD-EEG and Models of Morphological Decomposition in Word Processing

by

Tyson Jordan

Approved:

Major Professors:	Dustin Chacón Jin Sun
Committee:	Frederick Maier

Electronic Version Approved:

Ron Walcott
Dean of the Graduate School
The University of Georgia
August 2025

Dedication

To Mom, Noeli, and the folks on Whisperwind.

Acknowledgments

There are many people who supported me in completing this thesis. First, I would like to thank Dr. Dustin Chacón, who welcomed me into his lab and provided me with formal training in EEG. This work has made me a better scientist. I must also thank Dr. Jin Sun, an expert in computer vision who accepted me as a student before I had ever written a line of code in Python. You taught me everything I know about deep learning. I am also grateful to my committee member, Dr. Frederick Maier, who has tirelessly offered his mentorship since I was an undergraduate student.

I give special thanks to Donald Dunagan, who spent countless hours helping me collect data and troubleshoot issues. I am also extremely grateful to my family, friends, and labmates, who dedicated their time to participate in my experiments on a volunteer basis (all N=21 of you). I thank Victoria Taylor, our stimuli actress, and Harsha Veena Tadavarthy, who assisted in preparing the stimuli for this project.

I acknowledge my mentors: Drs. Anna Abraham, Harrison Chapman, and Adam Goodie. Finally, I thank Dr. John T. Hale for his contributions to the lab, and Dr. Linnaea Stockall for her foundational work. This project was funded by the UGA Office of Research (Chacón) and the National Science Foundation (1903783, Hale).

Contents

1 Introduction **1**

1.1 Contributions 8

1.2 Thesis Structure 10

2 Related Work **11**

2.1 Theoretical and Behavioral Foundations of Morphological Decomposition 12

2.2 Neural Evidence of Morphological Decomposition in Reading 15

2.3 Decomposition Across Modalities 17

2.4 Probing Lexical Representations with Natural Language Processing Models 21

3 Experiments **24**

3.1 Rationale and Hypotheses 24

3.2 Materials 27

3.3 Participants 29

3.4 Procedure 30

3.5 HD-EEG Processing 32

4 Representational Similarity Analysis (RSA) 35

4.1 Theoretical Model RDMs for Morphological Decomposition 36

4.2 Lexical Semantic Association with RSA 40

4.3 Statistical Analyses on RSA Correlation Coefficients 45

5 Results 48

5.1 Behavioral Results 48

5.2 RSA 50

6 Discussion 55

6.1 Limitations and Future Directions 61

7 Conclusion 63

A Word List by Condition 78

B Anatomical Labels used in Constrained Source Space Analyses 83

List of Figures

3.1 The multimodal stimuli used in the experiments. Audio icons are included for demonstration and were not present during the actual experiment. 30

4.1 Example Binary Models of Morphological Decomposition. When crafting the representational dissimilarity matrices (RDMs), Similar = 0 and Dissimilar = 1. 38

4.2 Cosine similarity values of GPT-2 representations of our grammatical stimuli. As shown, L2 normalization does not mitigate the issue of anisotropy. 42

4.3 Distribution of cosine similarity values among GloVe representations of our real word stimuli. 42

5.1 Behavioral results for the lexical decision tasks in the Text, AV, Silent Video, and Audio modalities. For each modality, we plot mean classification accuracy for each key condition. 49

5.2 Significant statistical results on the RSA correlations for Text, Audio, and AV. Spatiotemporal cluster-based permutation t-tests indicate significant differences from a null distribution, while repeated-measures one-way ANOVAs identify differences between models. FDR-corrected results are reported as q-values. 52

List of Tables

3.1 Summary of stimulus conditions used in the experiments, grouped by lexical status. 28

4.1 Summary of parameters used for cluster-based permutation tests across different stimulus modalities and model types. 46

Chapter 1

Introduction

How are words processed by and represented in the brain? This foundational question concerns the structure of the mental lexicon. One view holds that words are stored holistically, represented as single units in the mind. On this account, any resemblance between the activation of a word like 'baker' and its stem, 'bake', arises purely from semantic similarity or analogical structure. An opposing theory posits that words are stored compositionally, organized by their morphology. This view motivates the theory of **morphological decomposition**: the idea that complex words are parsed into their constituent components during processing. Under this premise, the word 'baker' is decomposed into its stem, 'bake', and affix, 'er' ($baker \rightarrow bake + er$). As morphemes are the smallest atomic units of language that maintain meaning, **morphologically complex** words are semantically related to their stems; a baker is someone who bakes. While the concept of morphological decomposition is generally accepted for such transparent cases, it remains unclear how the brain handles

morphologically pseudocomplex words, like 'whisper', which appear complex orthographically but do not maintain a semantic relationship with their stems. There is evidence that morphological decomposition overapplies: 'whisper' decomposes into 'whisp' + 'er' (Rastle et al., 2004; Gwilliams and Marantz, 2018), suggesting morphological decomposition is first attempted on the basis of the visual form of the word, rather than true morphological form. EEG experiments (Morris and Stockall, 2012) relate a similar priming effect to an N250, a negative event-related potential (ERP) that peaks ∼250 ms post-stimulus onset. Meanwhile, MEG studies have linked early decomposition to activity in the visual word form area (VWFA) ∼170 ms post-stimulus onset, known as the M170 component (Zweig and Pylkkänen, 2009; Solomyak and Marantz, 2010). The VWFA is a region in the left anterior fusiform gyrus believed to be involved in the recognition of written words and letters (Cohen et al., 2000). Crucially, Zweig and Pylkkänen (2009) show a larger M170 for morphologically plausible strings read in isolation (*singer, refill*) than for **simplex** words that appear complex but do not contain existing stems (*winter, reckon*), suggesting that the initial check is morphological. The idea of form-based parsing is taken a step further by the full decomposition model (Stockall and Marantz, 2006), which argues for a blind decomposition where even irregular forms are decomposed (*taught* → *teach*). Critically, this is followed by a stage of root activation, indexed by an M350 (∼350 ms post-stimulus onset), localized to the left temporal cortex. Later still, semantic interpretation and whole-word composition engage higher-level frontal areas such as the left orbitofrontal cortex (Stockall et al., 2019).

These findings raise critical questions about the temporal sequence associated

with processing complex, pseudocomplex, and simplex words. First, we note that evidence of early morpho-orthographic decomposition is implicated by both M170 and N250 effects. Whether this discrepancy arises from different imaging techniques or experimental paradigms, further research is needed. Moreover, Zweig and Pylkkänen (2009) yield inconsistent evidence regarding whether pseudocomplex words (*whisper*) are decomposed, suggesting that these opaque derived forms pattern with complex words in one analysis and diverge from complex words in two other analyses. As indicated, this does not coincide with additional accounts of early morphological processing, and thus, a follow-up is warranted. Finally, we question how the morpho-orthographic priming of pseudostems (Rastle et al., 2004), which indicate the access of 'whisp' when presented with 'whisper', coincides with the full decomposition model, which proposes post-decomposition processes that result in semantic interpretation at later stages following decomposition. While a pseudocomplex word might be blindly parsed during the initial stages of processing, the parse might be abandoned altogether once the conflict between the whole-word and its stem is recognized (Gwilliams and Marantz, 2018; Stockall et al., 2019). Thus, we question *if* and *when* the (pseudo)stems of (pseudo)complex words are accessed during whole-word processing.

While most studies of morphological decomposition have focused on visual word recognition, a growing body of work has demonstrated that decomposition occurs in other modalities. Reading is unique in that all aspects of morphological form are visible at once. However, speech provides morphemes sequentially, phoneme-by-phoneme. Behavioral studies utilizing auditory masked priming and primed lexical

decision tasks report priming effects for semantically transparent (*treat, treatment*) and semantically opaque (*pig, pigment*) prime-target pairs, indicating that listeners retrieve the meanings of (pseudo)stems even when they are not semantically related to the whole word (Creemers et al., 2020; Creemers and Embick, 2021; Creemers et al., 2023). An experiment using MEG collected during Arabic auditory word recognition (Gwilliams and Marantz, 2015) found that neural activity in the superior temporal gyrus (STG) is modulated by morpheme-level surprisal (the unexpectedness of a phoneme given the word's root consonants; e.g., the surprisal of *n-b-t* given *n-b*), rather than whole-word surprisal (the unexpectedness of a phoneme given all preceding phonemes; e.g., the surprisal of *nabata* given *naba*), with emerging effects as early as ∼130 ms post-onset of the critical root consonant. This provides additional evidence that root morphemes are accessed during spoken word recognition. Taken together, one might argue that the serial nature of speech seemingly triggers pseudostem activation, with separate effects of morphology occurring following a disambiguating point. Whether the initial activation lingers or is corrected once the true morphological structure is known is a question this thesis aims to answer.

Interestingly, lexical access during auditory processing has been shown to improve when auditory stimuli are integrated with visual cues (i.e., a speaking face) (Sumby and Pollack, 1954; Munhall et al., 2004). Campbell (2008) proposes that audiovisual (AV) speech processing involves both complementary and correlated (redundant) neural modes, supported by distinct neural regions: the posterior superior temporal sulcus (pSTS) for integrating auditory and visual cues, and ventral visual areas for analyzing visible speech features. If audiovisual integration strengthens word iden-

tification, what can be said about lexical access during silent lipreading? Auditory speech comprehension is not a sufficient predictor of lipreading skill (Mohammed et al., 2006). While it is generally accepted that lipreading recruits cortical regions typically active during auditory processing, such as the primary auditory cortex (A1) (Calvert et al., 1997; Molholm and Foxe, 2005), others argue that lipreading is mediated by distinct visual pathways (Bernstein and Liebenthal, 2014). MEG research on lipreading indicates that the auditory cortex synchronizes with the rhythm of the *missing speech* sound, not the available visual cues, suggesting that the brain might internally reconstruct an auditory representation in the absence of auditory input (Bourguignon et al., 2020). Though these findings entail a critical role of visual inputs during speech processing, there is ultimately a shortage of audiovisual and lipreading research particularly interested in lexical access from a linguistic perspective. As part of this study, we question how effects of morphological decomposition manifest during the audiovisual processing of spoken language. Furthermore, we aim to address whether morphological form can be decoded via lipreading, and thus, whether morphological complexity is encoded within sequences of visemes.

Meanwhile, in the field of natural language processing, artificial intelligence models have demonstrated a remarkable capacity for extracting and storing linguistic information from text. Previous efforts have leveraged machine learning techniques to reveal convergence in the internal representations of deep artificial neural networks and biological neural networks during language processing (Schrimpf et al., 2021; Caucheteux and King, 2022), suggesting similar semantic associations or distances between words in the neural and vector representation spaces. By establishing linear

relationships between the representations of distributional word embedding models or large language models (LLMs) and neural data recorded during language processing, these approaches can successfully probe spatiotemporal neural activity associated with semantic processing in EEG, MEG, functional magnetic resonance imaging (fMRI), and electrocorticography (ECoG). Looking beyond text-based language representations, state-of-the-art (SOTA) multimodal language models are equipped with cross-modal understanding, allowing linguistic interpretation of both text and video channels (Zhang et al., 2023; Li et al., 2023; Maaz et al., 2023). While these architectures provide a powerful framework for representing multimodal language stimuli, our work solely relies on static text embedding models.

To investigate how (pseudo)morphemes of morphologically complex, pseudocomplex, and simplex words are stored and processed in the brain, we conducted an experiment with a 128-channel high-density Electroencephalogram (HD-EEG), utilizing the single-word lexical decision paradigm. We curated a list of 360 (non)words, organized into multiple conditions, including morphologically **Complex** words ('baker'), morphologically **Pseudocomplex** words ('whisper'), **Simplex** words ('monster'), and the (pseudo)stems of these whole word forms. Driven to uncover traces of decomposition across modalities, we recorded an actress's face as she spoke each word from the stimuli list. This allowed us to prepare and present each word across four modalities: **Text, Audio, Audio-Video (AV), and Silent Video**. Subjects' neural data were recorded while they viewed stimuli in each modality and decided if each word was grammatical English or not. To localize the recorded HD-EEG activity, we applied source localization (reconstruction) with standardized low res-

olution electromagnetic tomography (sLORETA) (Pascual-Marqui et al., 2002). As EEG only measures electrical activity on the scalp, source reconstruction can be used to estimate the activity of source points in the brain. This method is commonly paired with MEG, although whether it is an appropriate method for EEG has been extensively questioned (Whittingstall et al., 2003; Grech et al., 2008). Nevertheless, others have validated the accuracy and reliability of source localization for EEG with proper optimization (Mikulan et al., 2020). The use of sLORETA with EEG has been shown to successfully localize neural markers of language processing in the visual word form system, with results cross-validated by fMRI localization (Brem et al., 2009). Our study builds on this foundation by leveraging sLORETA with HD-EEG to explore modality-dependent neural representations of morphological structure, demonstrating its utility for localizing language-related brain activity across diverse stimuli.

To analyze the neural data collected during the lexical decision tasks, we used a representational similarity analysis (RSA). The RSA examines abstracted relationships between neural activation patterns of different stimuli, rather than absolute activity levels, which can be difficult to compare across subjects (Kriegeskorte et al., 2008). RSA requires two representational dissimilarity matrices (RDMs) for comparison: a "model" RDM and a "brain" RDM. The RDMs serve as representations storing vital relationships between channels of information across time. In our study, RSA served two purposes. First, we tested whether the brain's response patterns aligned with theoretical models of morphological decomposition, leveraging the inherent property that words should share representational similarity with their

(pseudo)stems if decomposition occurs. Second, we extended RSA to examine lexical semantic associations, by comparing dissimilarities of words in the neural space with the differences of representations derived from a computational natural language processing model: GloVe. The goal of this approach is to probe semantic processing of whole words and, if words are decomposed, their stems. By correlating EEG-derived representational dissimilarity matrices (RDMs) with those from morphological and semantic models, we aimed to determine whether the brain processes words in a decompositional or holistic manner.

1.1 Contributions

This work provides several key contributions to the fields of neurolinguistics and computational neuroscience.

- We conduct a robust investigation of morphological decomposition across four lexical decision experiments, each distinguished by the modality in which stimuli are presented. To our knowledge, this is the first study to investigate morphological decomposition using audiovisual and silent video stimuli.

- We provide novel evidence that morphological decomposition 'overapplies' during visual word reading. With the use of a representational similarity analysis, we demonstrate that the neural signatures associated with processing **Complex** and **Pseudocomplex** derived forms (*baker*, *whisper*) overlap with the signatures of those words' (pseudo)stems (*bake*, *whisp*). This suggests full or blind decomposition, probing an apparent stem activation effect that aligns

with the lexeme-lookup for semantically transparent and opaque words ∼330 ms post-stimulus onset in the left temporal region. Replicating the Text RSA in source space, we identified a potential cluster of source points where the alignment of collected neural data and a model supporting full decomposition diverges from the alignment of other semantic models, with a positive trend localizing to the left anterior temporal lobe (LATL) beginning at ∼300 ms. This reveals a significant difference between semantic models.

- We find near-significant (n.s.) effects of morphological decomposition during audio processing in sensor and source space. Importantly, findings emerge *after* the average morpheme boundary or disambiguating point. These results allude to the decomposition of **Complex** and **Pseudocomplex** words. An ANOVA in source space supports these findings, distinguishing these theoretical decomposition models over a model that also decomposes **Simplex** cases. This distinction localizes to the right anterior temporal lobe (RATL), potentially signaling semantic interpretation in addition to morphological effects.

- We note strong effects of decomposition in response to our audio-visual stimuli in sensor and source space, following the disambiguating point. The occipital-temporal region displays the highest alignment with theoretical models that support the decomposition of **Pseudocomplex** words. An ANOVA indicated that the alignment of these models extended beyond that of a veridical decomposition model, which implies decomposition for **Complex** words only. These results suggest that the decomposition of **Pseudocomplex** words might be a critical aspect of the decomposition process in audio-visual word recognition.

1.2 Thesis Structure

The rest of this thesis is structured as follows:

- Chapter 2 contains a literature review on the study of morphological decomposition in visual word reading experiments. We consider approaches that investigate decomposition and lexical access in different modalities, and look towards neuroimaging research for further insights. We also discuss efforts in natural language processing (NLP) that leverage computational models to probe neural activity associated with language processing.

- Chapter 3 details the rationale, hypotheses, methodology, materials, and participants involved in our EEG experiments.

- Chapter 4 provides a thorough summary of the representation similarity analyses (RSAs), which we applied to our collected neural data. We also discuss the statistical tests applied to the outputs of the RSAs.

- Chapter 5 provides the behavioral and neural results for each modality-specific experiment.

- Chapter 6 includes a summative discussion of our work, highlighting limitations and proposing potential avenues for future research.

- Chapter 7 concludes this work, summarizing our efforts and key contributions.

Chapter 2

Related Work

In this section, we review key literature that motivates this thesis. First, we introduce historical theories of lexical organization, contextualizing the study of morphemes in the brain within broader debates about the structure of the mental lexicon. We cover foundational behavioral experiments that provide evidence of morphological decomposition during reading. Later, we discuss neuroimaging research, before highlighting studies that investigate decomposition in modalities beyond reading. Finally, we consider recent advances in machine learning and natural language processing, focusing on how static embedding and language models have been used to explore the relationship between computational representations and human brain activity.

2.1 Theoretical and Behavioral Foundations of Morphological Decomposition

Understanding how morphologically complex words are stored in the brain is central to uncovering how language is represented and processed in the mind. Accordingly, this topic has been the center of discourse for decades. On one end of this debate are the proponents of the full-listing or single-unit hypothesis, arguing that simple and morphologically complex words are stored holistically within the mental lexicon (Butterworth, 1983). Evidence of full-listing comes from early behavioral experiments that found no difference between processing affixed (*melting*) and non-affixed (*sister*) words (Manelis and Tharp, 1977). Some have taken the full-listing hypothesis a step further, maintaining that complex forms are stored in the lexicon but are connected through a network of phonological and semantic relationships (Bybee, 1985, 1988). In this connectionist view, the lexicon is not organized by morphemes, rather, morphology is a mere byproduct of word usage and similarity (Bybee, 1995). Pinker (1998) proposed a dual-mechanism model where past tense irregular verbs (*went, broke*) are stored as whole lexical entries, as they do not conform to regular past-tenses that are formed by rules (e.g., adding '-ed'). Arguably, this approach assumes that irregular forms are not compositional, potentially overlooking linguistic evidence that both irregulars and regulars can be decomposed into morphosyntactic and phonological components (Embick and Marantz, 2005).

The full-listing hypothesis is distinct from but often compatible with lexicalist theory, which proposes holistic lexical entries that store form-based, semantic, and

syntactic information for unique words. The presumption that words are stored as atomic units (at the word level) while sentences are formed compositionally is commonly accepted in neurolinguistics, though there are arguments against this trend (Krauska and Lau, 2023). Notwithstanding full-listing accounts, there is an abundance of evidence that morphologically complex words are decomposed during language processing. Early work in English visual word reading (Murrell and Morton, 1974) showed that after learning a word list, recognition of a target (*sees*) was facilitated by morphologically related primes (*seen*) beyond orthographically similar primes (*singe*), supporting a model in which morphemes serve as recognition units. This aligns with later findings in morphologically rich languages like Hebrew (Frost et al., 1997), where root morphemes (*zmr*, related to singing) primed lexical access of morphologically related words (*tizmoret* "orchestra"), while semantically related but morphologically unrelated roots (*ngn*, to play an instrument) did not. These studies reinforce a model of lexical organization grounded in morphological structure.

Studies adopting different experimental paradigms have yielded similar findings. For instance, several experiments utilizing the *lexical decision task*, which requires subjects to classify stimuli as grammatical or not, revealed that nonword stems belonging to prefixed words (*juvenate*) take longer to process than nonwords that are *not* stems of grammatical prefixed words (*pertoire*) (Taft and Forster, 1975). An interpretation of this finding is that prefixes (e.g., 're-') are stripped *prelexically* during visual word reading, before a subsequent "look-up" of the (pseudo)stem. Similarly, a study in Italian showed that plausibly decomposable nonwords (*cantevi*, "sing" + a valid affix) result in slower reaction times and lower accuracies compared to

partially decomposable nonwords (*cantovi*, where *ovi* is not a valid affix), which are themselves more difficult to classify than nonwords containing no grammatical morphemes (*canzovi*) (Caramazza et al., 1988). Reaction times for the lexical decision task are faster for Dutch monomorphemic words with larger morphological families (groups of morphologically related words sharing a common root) (Schreuder and Baayen, 1997). This suggests a **shared activation** for target words *and related words* when only the target is presented.

Research using the masked priming paradigm (Forster and Davis, 1984), in which a briefly presented prime is flashed between a visual mask (######) and a target word, supports the idea that complex words activate shared representations with their stems (Rastle and Davis, 2003). In a landmark study, Rastle et al. (2000) found robust masked priming effects for semantically transparent, morphologically related pairs (*departure–depart*), even at the shortest stimulus onset asynchrony (SOA). These effects were statistically greater than those elicited by semantic or orthographic similarity alone. Interestingly, a follow-up study from Rastle et al. (2004) confirmed robust priming for semantically opaque pseudocomplex forms (*brother*), but not for words that do not appear to have an affix (*brothel*), lending support to early, form-based morphological decomposition during visual word recognition.

2.2 Neural Evidence of Morphological Decomposition in Reading

Neurolinguists have leveraged neuroimaging techniques to isolate brain regions involved in morphological processing. In a key study, Zweig and Pylkkänen (2009) conducted two MEG experiments using a single-word lexical decision task to compare responses to morphologically complex words and matched controls. In the first experiment, they observed a larger M170 component for suffixed ('-er') morphologically complex words (*baker*) compared to monomorphemic orthographic controls (simplex words like *monster*) and simple monomorphemic words (*almond*). Interestingly, the effect localized to the right fusiform gyrus. To rule out the possibility that the right-lateralized effect was specific to suffixed word forms, Zweig and Pylkkänen (2009) conducted a second experiment using prefixed stimuli. In this follow-up, they again observed an M170 amplitude increase for morphologically complex words, but this time the effect emerged bilaterally. Specifically, the right fusiform gyrus distinguished complex words from simple monomorphemic controls, while the left fusiform gyrus differentiated complex words from orthographic controls. These findings suggest that both hemispheres—and in particular, the fusiform regions—play a role in early morphological decomposition. This bilateral involvement is consistent with prior work identifying the fusiform gyrus, and especially the left-lateralized visual word form area (VWFA), as a key region for mapping visual form to linguistic structure during reading (Tarkiainen et al., 1999; Cohen et al., 2000). Unfortunately, Zweig and Pylkkänen (2009) provide inconclusive evidence regarding the decomposi-

tion of pseudocomplex words (*sweater*): a dipole analysis showed that these stimuli patterned with complex words and elicited an M170 response in the right fusiform gyrus, while sensor-level areal-mean signals and a distributed source analysis showed an opposite effect where responses to complex and pseudocomplex words diverged.

Additional research in MEG using a lexical decision task supports the involvement of VWFA in morphological processing, providing evidence of decomposition for free stems (*taxable*, where *tax* is a free stem that can stand alone as a word) and bound roots (*tolerable*, where *toler* is not a free stem but acts as a common root for multiple morphologically related words), but not for unique root words (like *vulnerable*, which does not have other members in its morphological family) (Solomyak and Marantz, 2010). These M170 effects were sensitive to affix frequency and conditional probability, while earlier (M130) activity was modulated by letter string frequency and transition probability (TP), implying an early visual analysis of morphological form. Gwilliams and Marantz (2018) provided similar results in MEG, linking an M170 response in the left fusiform to the decomposition of complex and pseudocomplex words, but not simplex words.

EEG paired with the masked priming paradigm has demonstrated early, pre-N250 priming effects for complex- and pseudocomplex-based prime-target pairs (Morris et al., 2008). Post-N250 activity suggested that complex words patterned with orthographic controls (*scandal–scan*). Additional EEG masked priming results suggest early or form-based morphological processing that applies to complex and pseudocomplex words 250 ms post-stimulus onset (Morris and Stockall, 2012). This is supported by a recent EEG lexical decision task in Chinese, which found similar evidence

of early decomposition for semantically transparent and opaque compound words, suggesting that early morphological parsing is independent of semantic transparency (Wei et al., 2023). As Chinese compound words do not contain clear visual morpheme boundaries (e.g., the suffix '-er' in English), this provides stronger evidence of form-based parsing rather than overt visual segmentation.

2.3 Decomposition Across Modalities

Reading is unique in that the visual system is presented with all aspects of morphological form at once. Lexical access extends beyond reading, though, and varies significantly when words unfold over time, such as in speech. Accordingly, there is a growing body of research concerning morphological decomposition during auditory processing. An experiment using a cross-modal repetition priming paradigm—where the prime is presented auditorily and the target is presented visually—found that semantically transparent prime–target pairs (*friendly–friend*) showed strong priming effects, even when the primes and targets were not phonologically identical (*serenity–serene*) (Marslen-Wilson et al., 1994). This was not the case for phonologically related, semantically opaque pairs (*tinsel, tin*). Similarly, Longtin et al. (2003) conducted a study in French using auditory-visual cross-modal priming to investigate whether cross-modal priming effects arise from morphological or orthographic relations. They found evidence of morphological decomposition for semantically transparent (*gaufrette–gaufre* "wafer–waffle") but not pseudo-derived (*baguette–bague* "stick–ring") word pairs, suggesting that semantic transparency might be necessary for cross-modal

morphological priming. A body of behavioral research using auditory priming experiments contrasts with these findings, instead suggesting that pseudocomplex forms are decomposed during auditory processing. With an auditory primed lexical decision task, where primes and subsequent targets are played aloud, Creemers et al. (2020) found that both semantically transparent Dutch primes (*aanbieden* "offer") and semantically opaque primes (*verbieden* "forbid") facilitated target recognition (bieden "offer") over semantic and phonological controls. The same auditory priming paradigm was used to test whether processing opaque Dutch words (*herhalen* "repeat") activates the meaning of the pseudostems (*halen* "take"), by testing for priming effects on the recognition of semantically related targets (*brengen* "bring") (Creemers and Embick, 2021). Indeed, the results supported the lexical access of pseudostems for auditorily-presented pseudocomplex words. Similar priming effects for semantically transparent (*treat, treatment*) and semantically opaque (*pig, pigment*) prime–target pairs have been noted with auditory priming in English, though show the strongest effects for transparent cases (Creemers et al., 2023). In a key study, Ettinger et al. (2014) utilized MEG to examine the effects of morphological complexity (*bruiser* vs. *bourbon*) and the predictability of word endings (*bourbon* vs. *burble*). Higher surprisal resulted in increased auditory cortex activity, which was modulated by word complexity, indicating a relationship between morphology and phoneme-level processing.

It is important to consider the role of visual information in lexical access during audiovisual processing. Cues from a speaking face or moving head have been shown to enhance the comprehension of oral speech (Sumby and Pollack, 1954; Munhall

et al., 2004). The fusiform face area (FFA), housed in the right fusiform gyrus, is typically associated with face processing (Kanwisher et al., 1997). Visual inputs can be manipulated to misguide perception, as demonstrated by the classic McGurk effect (McGurk and MacDonald, 1976), where incongruent visual and auditory information leads to a fused interpretation. This is typically interpreted as a low-level perceptual phenomenon, independent of and likely occurring before lexical access. A semantic priming lexical decision experiment from Ostrand et al. (2016) showed that when presented with audio-video conflicts (McGurk stimuli, $bait_{\text{aud}} + date_{\text{vis}} \rightarrow date_{\text{percept}}$), subjects were primed to recognize auditory-only primes related to the true audio stimuli (*bait*), despite consciously perceiving the wrong phoneme. This suggests that AV-integration can occur *after* lexical access, depending on the lexicality of the auditory and visual inputs. Campbell (2008) outlines two neural modes of audiovisual processing: a complementary mode, where vision supplies information absent in sound, and a correlated mode, where visual and auditory speech features mirror each other: a form of robust yet redundant processing. These modes are associated with distinct cortical regions, with pSTS involved in integrating dynamic audiovisual patterns and ventral visual stream regions supporting detailed articulatory analysis. This framework helps clarify how visual speech can influence lexical access through both distinct and overlapping pathways. Supporting the view that visual cues are not just supplemental, Van Wassenhove et al. (2005) show that visual cues speed up auditory processing, suggesting that visual speech actually modulates auditory processing in the superior temporal gyrus. The authors suggest that this finding is a marker of predictive coding, where visual input is used to predict and facilitate

speech processing. Indeed, findings suggest that a mismatch between audio and visual cues results in decreased intelligibility of stimuli when audio leads, but not when video leads (Grant and Greenberg, 2001; Grant et al., 2003).

In cases when auditory information is entirely absent, silent lipreading recruits cortical regions which are active during auditory processing, such as the primary auditory cortex (Calvert et al., 1997; Pekkola et al., 2005). Despite this overlap, evidence suggests that auditory speech comprehension is not a sufficient predictor of lipreading skill (Mohammed et al., 2006). In fact, low variability in speech recognition tasks is contrasted with highly variable performance in lipreading tasks (Summerfield, 1992). Bernstein and Liebenthal (2014) argue that lipreading recruits dedicated visual pathways, distinct from those involved in auditory speech perception. They propose a model in which visual speech is represented in modality-specific regions, such as the temporal visual speech area (TVSA), which are separate from the superior temporal gyrus (STG) regions believed to process acoustic phonemes. Recent work by Bourguignon et al. (2020) used MEG to compare neural responses during video-only and audio-only presentations of a speaker telling a story. In the video-only condition, where no sound was present, auditory cortex activity synchronized with the absent auditory speech envelope rather than with the visible lip movements themselves. This suggests that the brain engages in a bottom-up process that reconstructs coarse auditory features directly from visual input, rather than relying on top-down predictions or prior knowledge. A similar study in EEG suggests that more accurate lipreading facilitates more accurate entrainment to the envelope of the missing speech sound in the left temporal region, and that the envelope can be reconstructed from

the recorded neural signal (Crosse et al., 2015). These findings raise the question of whether higher-level linguistic properties—such as morphological complexity or lexical structure—can also be reconstructed from visual speech alone. In contrast, Baart and Samuel (2015) found early effects (< 400 ms) of both lexicality and lipreading context (facial movements) on auditory Event-Related Potentials (ERPs), but no interaction between the two, suggesting that facial information alone may not directly support lexical access. To our knowledge, this thesis is the first study to investigate whether morphological form can be decoded from the silent visual presentation of a speaking face.

2.4 Probing Lexical Representations with Natural Language Processing Models

Advances in machine learning have offered new methods for exploring how words are represented in the brain. Arguably, these techniques can circumvent the limitations of traditional neurolinguistic approaches by providing quantitatively explicit models that can be directly compared to neural data (Sassenhagen and Fiebach, 2020). In the context of semantic representation, distributional word embedding models like GloVe (Pennington et al., 2014) and Word2Vec (Church, 2017) are designed to represent word meanings, which are learned through co-occurrence relationships within large text corpora. These models represent words in high-dimensional vector space, where semantically related words are "closer together". Sassenhagen and Fiebach (2020) found that embeddings from Word2Vec could significantly predict EEG ac-

tivity collected during visual word reading. Using ridge regression, observed and predicted neural activity were significantly correlated, suggesting a linear relationship between the distances of words in vector space and the neural activations of those same words. Specifically, the embeddings most-strongly predicted N400 activity, an ERP reflective of facilitated lexical access (Lau et al., 2008).

Similar findings have been replicated with fMRI and ECoG data (Schrimpf et al., 2021). Schrimpf et al. (2021) tested the neural predictivity of a wide selection of models, ranging from distributional word vector models (e.g., GloVe) to large language models (LLMs) (e.g., GPT-2). They found that a model's next-word prediction accuracy on the WikiText-2 dataset (Merity et al., 2016) was positively linked to neural predictivity, with GPT-2 achieving the highest performance. In addition to confirming that neural predictivity is a plausible approach for several imaging techniques, this study suggests greater alignment between neural activity and higher-performing NLP models over simpler ones. These results are supported by Caucheteux and King (2022), who report similar convergence between language models and MEG data recorded during language processing. Interestingly, the authors also demonstrate that the representations of a convolutional neural network (CNN) trained on visual character recognition best predicted neural activity in the early visual cortex (V1). This finding emphasizes the promise of language models in extracting linguistic information from non-text inputs. Indeed, SOTA multimodal language models are able to process and integrate linguistic information from distinct modalities, such as text and videos (Zhang et al., 2023; Li et al., 2023; Maaz et al., 2023). Despite these exciting findings, others have critiqued LLM-driven neural encoding analyses.

For example, Hadidi et al. (2025) note flaws in the methodology typically used for these works, and point out that simple confounding variables (e.g., word rate) achieve comparable performance with trained LLMs in neural prediction tasks. These warnings should be carefully considered in future studies that leverage LLMs for neural encoding. In this thesis, we only rely on static word embedding representations as a metric of semantic similarity.

Chapter 3

Experiments

In this chapter, we establish the rationale driving our investigation and list several hypotheses. We provide a description of our participants and stimuli, and detail the procedures of our human-subjects experiments. We include data processing steps and discuss source reconstruction for HD-EEG.

3.1 Rationale and Hypotheses

The experiments described in this chapter were designed to empirically test competing models of morphological decomposition across multiple sensory inputs. Findings on morphological processing in visual word reading fill a mixed bag; MEG paired with the single-word lexical decision task suggests an early sensitivity to morphologically **Complex** words over **Simplex** words (Zweig and Pylkkänen, 2009) in the right fusiform, but yields inconsistent evidence for the parsing of **Pseudocomplex** forms. On the other hand, behavioral (Rastle et al., 2004) and EEG (Morris et al.,

2008; Morris and Stockall, 2012) masked priming experiments show clear evidence of form-based decomposition for **Complex** and **Pseudocomplex** words, though point to a later N250 effect. Whether the temporal discrepancy arises from inconsistent imaging techniques or experimental paradigms, further research in EEG is needed. Furthermore, we question whether and how the post-decomposition processes of the full decomposition model (Stockall and Marantz, 2006) apply following the morpho-orthographic decomposition of **Pseudocomplex** words. If **Pseudocomplex** forms (*whisper*) prime the access of their pseudostems (*whisp*), how do the post-decomposition processes unravel? To answer this, a signifier of semantic access for words and their (pseudo)stems is required. In addition to addressing these issues, we aim to establish source reconstruction as a reliable localization method for HD-EEG (for more details, see Section 3.5.2).

Outside of reading, morphological processing has received comparatively less attention. The serial nature of auditory and visual speech perception distinguishes how morphological information within these modalities is processed temporally, compared to reading. A growing body of work suggests that decomposition applies during auditory processing (Ettinger et al., 2014), and that early effects can be detected following the disambiguating point or critical phoneme onset (Gwilliams and Marantz, 2015). Here, we leverage this idea, constraining our analyses to surround the disambiguating *morpheme boundaries* of our single-word stimuli. As discussed, AV integration is doubly influenced by visual and auditory input channels, which both contribute to lexical access in auditory word recognition (Sumby and Pollack, 1954). Moreover, there are different accounts of whether or not lipreading recruits the same cortical

regions used in auditory processing (Calvert et al., 1997; Bernstein and Liebenthal, 2014), alongside interesting behavioral findings that contrast how well subjects recognize words during listening and lipreading (Mohammed et al., 2006; Summerfield, 1992). Some evidence suggests that coarse-grained information of missing speech can be decoded from neural data recorded during lipreading (Crosse et al., 2015). From these findings, we ask: Does the modality of input (e.g., reading vs. listening vs. lipreading) affect whether and when morphological decomposition occurs? Additionally, can morphological structure be decoded from visual-only speech (i.e., silent lipreading), suggesting that visemic inputs encode morphological complexity?

To address these questions, we adopt a multimodal framework, spanning written, spoken, and visual-only input. We design our experiment accordingly, assessing how the brain processes words of varying morphological complexity. We apply Representational Similarity Analysis (RSA) (see Section 4), which allows us to assess whether neural activity patterns align with models of morphological decomposition. We extend this approach by incorporating distributional semantics from GloVe embeddings, probing spatiotemporal activity associated with semantic processing. This method allows us to assess if, when, and where different words and their (pseudo)stems are accessed during language processing.

Based on our rationale, we establish several hypotheses:

- H1: In visual word reading, morphological decomposition applies to **Complex** and **Pseudocomplex** words, but not to Simplex cases. We expect early decomposition effects, either ∼170 or ∼250 ms post-stimulus onset, emerging in the left or right fusiform areas. Following the work of Stockall and Marantz

(2006), we expect stem activation in the left temporal region ∼350 ms post-stimulus onset, followed by full semantic interpretation.

- H2: In the audio and audio-visual lexical decision tasks, effects of decomposition will emerge following the disambiguating point or average morpheme boundary of the stimuli. These effects might localize to the left temporal region, though audiovisual processing may invoke a range of cortical regions specific to modality-based inputs.

- H3: Silent visual speech supports morphological decomposition. In the absence of auditory input, silent lipreading will evoke neural patterns in auditory regions that reflect morphological parsing, indicating that sufficient linguistic information is encoded in and decoded from visemes. Based on the literature, incongruent findings across behavioral and neural results are expected.

3.2 Materials

Here, we outline the creation of the multimodal stimuli used in the lexical decision experiments. A summary of the conditions is shown in Table 3.1. We prepared:

- 90 grammatical English words ending in '-er' across three key conditions:
  - 30 morphologically **Complex** words ('baker').
  - 30 morphologically **Pseudocomplex** words ('whisper'). Each word in this condition contained a stem that was also an existing English word.

  - 30 **Simplex** words ('monster'). Each word in this condition contained a non-existing English word as a stem.

- 60 (pseudo)stems belonging to the **Complex** and **Pseudocomplex** words ('bake', 'whisp').

- 30 2-syllable monomorphemic controls ('almond').

- 180 1-syllable and 2-syllable nonwords, which were used as fillers ('zort', 'toger'). 30 nonwords were derived from the stems of **Simplex** words ('monst'), and the remaining were generated using a free online fake word generator.

Table 3.1: Summary of stimulus conditions used in the experiments, grouped by lexical status.

Lexicality	**Condition**	**Count**
Words	**Complex** (*baker*)	30
	Pseudocomplex (*whisper*)	30
	Simplex (*monster*)	30
	Complex Stem (*bake*)	30
	Pseudostem (*whisp*)	30
Nonwords	Simplex Stem (*monst*)	30
	Monomorphemic (e.g., *zort*)	30
	Bimorphemic (*toger*)	120

This is equivalent to 180 grammatical English words and 180 nonwords. The full stimuli list is provided in Appendix A. We adapted our stimuli from published materials in prior papers (Zweig and Pylkkänen, 2009; Solomyak and Marantz, 2010; Gwilliams and Marantz, 2018). Key stimuli were controlled for length, whole-word and stem frequency, and phonological and orthographic neighborhood density (all

$p > 0.1$), all of which were calculated using the CELEX2 corpus (Baayen et al., 1996). CELEX includes a total of 17.9 million words. As a novel addition, we also included the (pseudo)stem of each key stimulus.

We recorded an actress's face as she spoke each word into a Canon EOS R5 mirrorless digital camera. Each word was manually parsed into its own clip from a longer recording, with an average clip duration of 1.03 seconds. Using the word list and video data, we prepared four versions of each word corresponding to four modalities: Text, Silent Video, Audio, and Audio-Video (AV). The Silent Video and Audio stimuli were obtained by stripping the audio from each AV clip using the FFmpeg Python package (Tomar, 2006). An example of the stimuli is shown in Figure 3.1.

We conducted text-speech alignment with the Montreal Forced Aligner (McAuliffe et al., 2017), and extracted timestamps of each phoneme from the resulting TextGrid files using PRAAT (Boersma and Van Heuven, 2001). This allowed us to identify the (pseudo)morpheme-boundaries for each word in our stimuli (i.e., when the '-er' occurred). This is critical as it marks the disambiguating point of auditory and visual presentation, which is later used to constrain the search windows of our non-text analyses.

3.3 Participants

Twenty-one self-identified English speakers were recruited by word of mouth from the University of Georgia and greater community. Twenty participants were right-

Figure 3.1: The multimodal stimuli used in the experiments. Audio icons are included for demonstration and were not present during the actual experiment.

handed, as assessed by the Edinburgh Handedness Inventory (Oldfield, 1971). All participants had normal or corrected-to-normal vision, and reported no history of language impairment. Participants provided written informed consent. Participation was voluntary for this IRB-approved study.

3.4 Procedure

3.4.1 HD-EEG Acquisition

EEG signals were recorded using a high-density 128-channel BrainVision actiChamp+ system. Impedance of the EEG sensors was reduced by the application of Super-Visc gel. On-line EEG recording was referenced to FCz according to manufacturer standards, and then re-referenced to average sensors offline.

3.4.2 Electrode Digitization

We used a handheld camera-based CapTrak scanner for 3D electrode localization. This entails the digitization of all electrodes with respect to each participant's head. Additionally, we digitized three anatomical landmarks—the nasion and the preauricular points located in front of each ear. The digitization process is necessary for source reconstruction (see Section 3.5.2).

3.4.3 Lexical Decision Tasks

Each subject participated in four self-paced single-word lexical decision experiments, which varied by the modality in which stimuli were presented: Text, AV, Audio, or Silent video. Each (non)word was presented once in each modality, for a total of 1440 trials. The order of each experiment was counterbalanced across subjects. The lexical decision tasks were programmed with the PsychoPy Python package (Peirce, 2007) and presented on a lab computer monitor.

For each experiment and each given subject, stimuli were pseudo-randomized to ensure counterbalancing of each condition within five presentation blocks. At the start of each experiment, subjects were provided with instructions and notified that they would be presented words in a given modality. For the Audio and AV experiments, subjects were provided with wired headphones, which were kept at a consistent volume between subjects. Subjects were instructed to press the 'f' key on a keyboard when presented with a grammatical English word, and the 'j' key otherwise. Furthermore, subjects were asked to maintain a fixed gaze at the center of the screen for each experiment. After reading the instructions, subjects pressed

the SPACE key to begin each trial. In the Text experiment, each word was displayed on a lab monitor for one second. The Video, AV, and Silent Video experiments presented each word for the duration of their corresponding clip. Upon completing a block (70 trials), subjects were provided with an additional self-paced break.

3.5 HD-EEG Processing

3.5.1 Preprocessing

All EEG preprocessing was conducted in MNE-Python (Gramfort et al., 2013). Raw EEG data was filtered offline between 0.1–40 Hz, using an IIR bandpass filter. We then removed flat or noisy sensors, and interpolated them. Next, we re-referenced the EEG data to an average reference. Afterwards, we extracted epochs from –100ms to 1000ms post-stimulus onset. The 100ms pre-onset period was then used for baseline correction. We then used independent component analysis (ICA) to identify semi-regular endogenous electromagnetic noise sources, including eyeblinks, eye movements, and heartbeats. These components were removed. Following ICA, we automatically rejected all epochs that exceeded a 100μV peak-to-peak threshold. Finally, we visually inspected and removed other problematic epochs.

3.5.2 Source Reconstruction

While EEG offers high temporal resolution, this method is limited spatially, as signals derived in sensor space only allude to electrical activity measured at the scalp. Thus, to estimate the origins of activity in the brain, we utilized *source reconstruc-*

tion. Source reconstruction works by establishing a forward solution, which maps a constructed *source space* on the cortical surface to expected sensor readings. Using this mapping and the observed sensor activity, reconstruction solves the *inverse problem*, estimating activity for each point in the source space. Standardized Low Resolution Brain Electromagnetic Tomography (sLORETA) is one algorithmic solution commonly employed to solve the inverse problem (Pascual-Marqui et al., 2002). This approach has been paired with EEG to localize word-specific neural responses in the visual word form system, converging with the spatial localization of fMRI recordings during language processing (Brem et al., 2009).

The use of source reconstruction for EEG has not gone unchallenged. Dissent for this approach is in part propagated by early comparisons between EEG and MEG, arguing for MEG as a method with superior spatial resolution (Hari and Lounasmaa, 1989). While there is an abundance of research supporting source reconstruction for (HD-)EEG (Cohen and Cuffin, 1991; Klamer et al., 2015; Dattola et al., 2020), others note potential pitfalls of inverse modeling with EEG, such as susceptibility to noise or issues within contemporary inverse algorithms (Whittingstall et al., 2003; Grech et al., 2008). Thus, one motivation of this thesis is to establish that source localization of HD-EEG is a reliable approach for estimating neural activity. For a full review on this topic, see Kaur et al. (2022).

Source reconstruction was performed using MNE-Python (Gramfort et al., 2013). Each participant's digitized sensor positions were coregistered with the FreeSurfer template brain (fsaverage), including distinct layers for outer skull, inner skull, and brain (Fischl, 2012). For each subject, we computed a source space of the corti-

cal surface with approximately 4 mm spacing ('ico-4') and with 2,562 sources per hemisphere. We computed a forward solution using the automated method to produce a boundary element model (BEM). Noise covariance matrices were calculated from baseline periods to model background noise in the data. An inverse operator was created using the forward solution and noise covariance, and source localization was performed using standardized low-resolution electromagnetic tomography (sLORETA) for unbiased and physiologically plausible estimation of neural activity (Pascual-Marqui et al., 2002). The reconstructed source time courses (STCs) were morphed to a common template brain to enable group-level analyses, and results were saved for each trial per subject.

Chapter 4

Representational Similarity Analysis (RSA)

Due to inherent inter-subject variability in neurophysiological data, understanding the organization of linguistic structures across individuals poses a challenge in cognitive neuroscience. Traditional statistical approaches often rely on spatial or temporal aggregation, which is prone to diluting fine-grained information in neural signal.

Representational Similarity Analysis (RSA) provides a powerful framework to circumvent these challenges by shifting the focus from absolute neural activity levels to the relationships between activation patterns (Kriegeskorte et al., 2008). Rather than requiring a strict spatial or temporal alignment across subjects, RSA characterizes neural representations using representational dissimilarity matrices (RDMs), which capture the relative similarity structure of neural responses to different stimuli. These neural RDMs are compared with model RDMs, which can be derived

from a variety of sources, including theoretical predictions or computational models. Model RDMs encode hypothesized dissimilarity structures among stimuli, allowing researchers to assess how well different models account for the patterns observed in neural representations. By correlating neural and model RDMs, RSA provides a powerful method for linking brain activity to abstract representational frameworks, offering insights into how linguistic information is structured in the brain.

Here, we use RSA in two key ways: (1) to compare EEG data with theoretical models of morphological decomposition, which define relationships between whole words and their stems, and (2) to compare EEG data with the semantic relationships between whole words and stems, derived from distributional semantic models like GloVe and GPT-2. In doing so, we establish six key RDMs, introducing the theoretical models in Section 4.1 and the semantic models in Section 4.2.

4.1 Theoretical Model RDMs for Morphological Decomposition

To investigate whether neural responses reflect morphological decomposition, we compared EEG data with three binary models of morphological structure. The use of theoretical binary RDMs is a well-established method for RSA (Wang et al., 2018). Our models define different (pseudo)morphemic relationships and provide a theoretical basis for assessing how whole words relate to their stems. If a word undergoes decomposition, its neural activation should overlap with the activation of its corresponding stem (Rastle and Davis, 2003). This is based on the fundamen-

tal principle of early morphological decomposition: when (pseudo)complex words are processed, they are broken down into their constituent parts, including their (pseudo)stems, yielding a shared activation. Conversely, if a (pseudo)complex word is processed holistically, its neural representation should be distinguishable from its putative stem (i.e., there are separate lexical entries for these words). By computing representational similarity between whole words and their stems, we assess whether neural responses align with each morphological decomposition model.

We evaluate three theoretical models: **Orthographic Stem**, **Lexical Stem**, and **Morphological Stem** (illustrated in Figure 4.1). Each model is represented as a pairwise **distance function** $f(x, y)$, where:

- $f(x, y) = 0$ indicates that two words are **similar**.
- $f(x, y) = 1$ indicates that two words are **dissimilar**.

1. **Orthographic Stem.** This model assumes decomposition is driven purely by surface-level orthographic similarities, without considering lexicality or true morphological structure. Under this model:
    - 'bake' and 'baker' are **similar**: $f(\text{bake}, \text{baker}) = 0$
    - 'whisp' and 'whisper' are **similar**: $f(\text{whisp}, \text{whisper}) = 0$
    - 'monst' and 'monster' are **similar**: $f(\text{monst}, \text{monster}) = 0$
    - All other word pairs are **dissimilar**: $f(x, y) = 1$

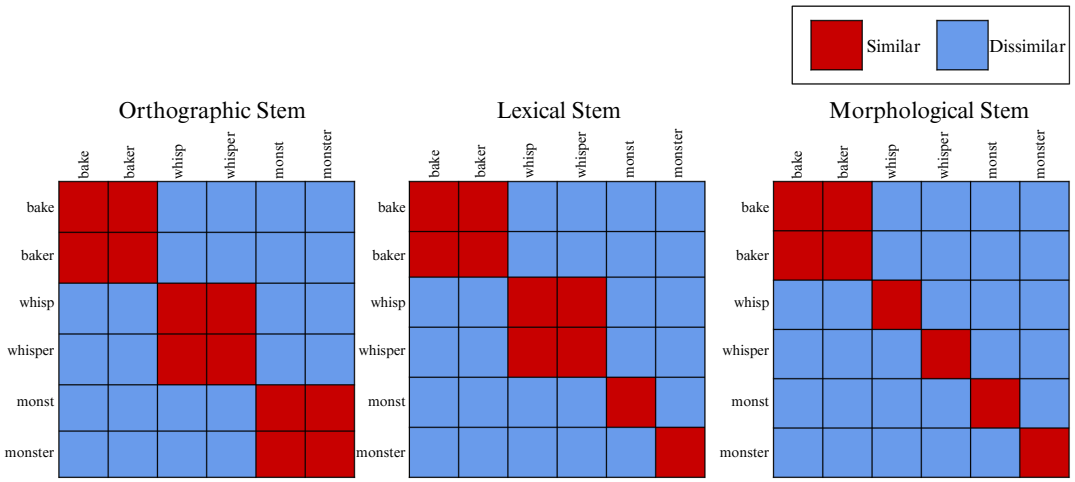

Figure 4.1: Example Binary Models of Morphological Decomposition. When crafting the representational dissimilarity matrices (RDMs), Similar = 0 and Dissimilar = 1.

2. **Lexical Stem.** This model introduces a lexical constraint, acknowledging that not all apparent stems (e.g., 'monst' from 'monster') form valid independent words. Here:

   - 'bake' and 'baker' are **similar**: $f(\text{bake}, \text{baker}) = 0$
   - 'whisp' and 'whisper' are **similar**: $f(\text{whisp}, \text{whisper}) = 0$
   - 'monst' and 'monster' are **dissimilar**: $f(\text{monst}, \text{monster}) = 1$
   - All other word pairs are **dissimilar**: $f(x, y) = 1$

3. **Morphological Stem.** This model recognizes only true morphological relationships, rejecting cases where words share orthographic similarities but lack genuine morphemic structure. Under this model:

- 'bake' and 'baker' are **similar**: $f(\text{bake}, \text{baker}) = 0$
- 'whisp' and 'whisper' are **dissimilar**: $f(\text{whisp}, \text{whisper}) = 1$
- 'monst' and 'monster' are **dissimilar**: $f(\text{monst}, \text{monster}) = 1$
- All other word pairs are **dissimilar**: $f(x, y) = 1$

We created Representational Dissimilarity Matrices (RDMs) for each theoretical model and for our collected HD-EEG epochs using MNE-Python (Gramfort et al., 2013). MNE-Python's searchlight approach extracts local spatiotemporal patches of EEG data, defined by temporal and spatial radii parameters, to compute representational dissimilarity. Each patch is reshaped into a vector by concatenating channel $\times$ time points within that patch, allowing for pairwise similarity comparisons across stimuli. To define local spatiotemporal patches, we used temporal and spatial radii of 50 milliseconds and 3 centimeters, respectively. For the brain model, the default distance metric calculated between epochs is defined as follows:

$$d_{\text{corr}}(u, v) = 1 - \frac{\sum_{i=1}^{n}(u_i - \bar{u})(v_i - \bar{v})}{\sqrt{\sum_{i=1}^{n}(u_i - \bar{u})^2} \cdot \sqrt{\sum_{i=1}^{n}(v_i - \bar{v})^2}}$$

where u and v are EEG response vectors for two different stimuli (words), n is the number of features in each vector (corresponding to selected EEG channels and time points), and $\bar{u}$ and $\bar{v}$ are the means of u and v. The denominator normalizes by the standard deviation of each vector. Thus, the RDMs are simply the inverse of the Pearson correlation coefficient PCC between vectors u and v:

$$d_{\text{corr}}(u, v) = 1 - \text{PCC}(u, v)$$

Thus, a brain RDM is constructed for each searchlight patch for a given subject, and is compared to each model RDM with Spearman's rank correlation coefficient ρ (Spearman, 1961). The resulting coefficient populates a correlation coefficient matrix—which has the same dimensions of the neural data—at the coordinates matching the center of the given searchlight patch.

4.2 Lexical Semantic Association with RSA

Next, we applied RSA to examine shared lexical access between whole words and their putative (pseudo)stems. Whereas the previous application of RSA using binary models of morphological decomposition aimed to detect *any* shared activation between words and their (pseudo)stems, the lexical semantic association RSAs specifically target semantic processing. This approach tests whether semantically related words elicit similar patterns of neural activation. Additionally, we consider whether semantically *unrelated* items elicit similar neural responses, as a consequence of decomposition or pseudostem retrieval. To achieve this, we leverage insights from prior research in natural language processing and machine learning, which has demonstrated convergence during semantic processing between biological and artificial systems. By comparing neural response patterns with computational models of lexical semantics, we aim to identify whether whole words and their stems evoke similar activation during meaning-related processing.

To investigate the extent to which decomposition effects correspond to lexical interpretation of a word's (pseudo)stem, we conducted RSA using **GloVe** (300-

dimensional embeddings). This model provided estimates of semantic relatedness, enabling us to compare word representations derived from distributional semantics. We note that contextual large language models were considered for this task (specifically, GPT-2), but were ultimately not included for two reasons: (1) The stimuli presented to human subjects lacked context; Words were presented in isolation, and thus were not processed as parts of sentences as in many previous studies that have linked LLM representations with neural activation. As such models are trained in a context-dependent fashion, we opted for static representations. Though static GPT embedding models are available, the subtokenization scheme would not parse the morphologically-specific stimuli by any meaningful boundaries. Additionally, (2) layer-level representations of LLMs like GPT-2 have been shown to suffer from anisotropy. As a result, any two random words represented in this manner have near-perfect cosine similarity ($\sim$1) (Ethayarajh, 2019). We provide a demonstration of anisotropy using GPT-2 (small checkpoint) hidden layer representations of our grammatical stimuli (obtained via a forward pass, averaging over the token space) in Figure 4.2. For these reasons, the LSA-RSAs were conducted with the distributional word vectors of GloVe only. To further contextualize this decision, we provide the distribution of cosine similarity values among GloVe representations of our stimuli (Figure 4.3). These values were notably more dispersed than those from GPT-2 (mean = 0.164, range = [–0.187, 0.692]).

Because this approach requires that GloVe can generate representations of its word inputs, it is impossible to extract representations for non-existing words, such as the pseudostems of the **Simplex** stimuli. For this reason, we limit this approach

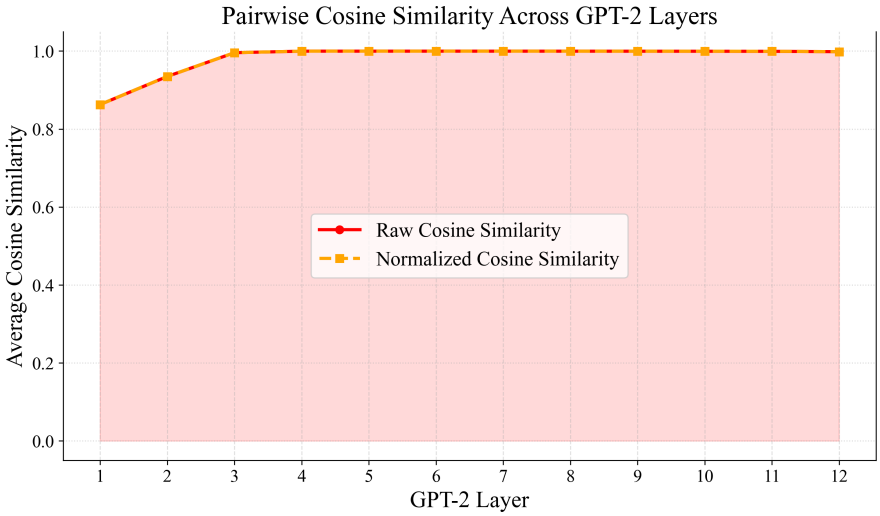

Figure 4.2: Cosine similarity values of GPT-2 representations of our grammatical stimuli. As shown, L2 normalization does not mitigate the issue of anisotropy.

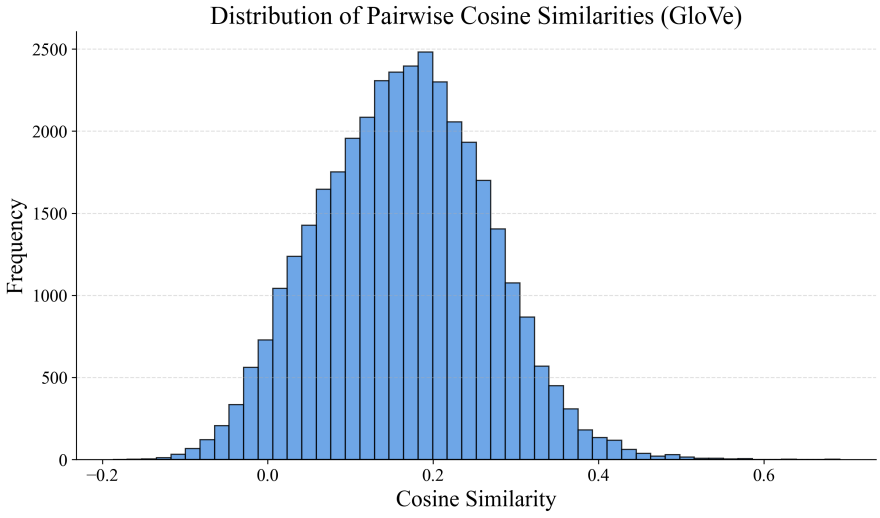

Figure 4.3: Distribution of cosine similarity values among GloVe representations of our real word stimuli.

to **Complex** words, **Pseudocomplex** words, and their putative stems.

For the GloVe-based RSA, we extracted static word embeddings from the pre-trained 300-dimensional GloVe model. After deriving semantic representations, we constructed three model RDMs using the GloVe embeddings. All semantic RDMs were created using **cosine similarity**. However, since MNE-Python expects a dissimilarity matrix for RSA computations, we used the transformed metric:

$$d_{\text{cosine}}(x, y) = 1 - \cos(x, y)$$

where x and y are the vector representations of two words. This transformation ensures consistency with the distance-based RSA framework. The three RDMs were varied as follows:

4. **LSA.** This model does not undergo any alterations. The RDM is populated with the cosine dissimilarities of each word-pair combination.

    - All word combinations: $1 - \cos(\textit{word1}, \textit{word2})$

5. **LSA-Stem.** This model replaces the semantic embeddings of derived **Complex** words with the embeddings of their putative stems. All other word-pair combinations remain the same as in the *LSA* model. This model represents *true* morphological processing, where **Complex** words have shared activations with their stems, and thus, processing **Complex** derived forms triggers the activations of their putative stems.

'baker' $\rightarrow$ **'bake'**

- 'bake' and 'baker': $1 - \cos(\text{bake}, \text{bake}) = 0$
- 'whisp' and 'whisper': $1 - \cos(\text{whisp}, \text{whisper})$
- All additional word combinations: $1 - \cos(\textit{word1}, \textit{word2})$, but replace embeddings of derived **Complex** forms with those of their stems.

6. **LSA-PseudoStem.** This model replaces the semantic embeddings of derived **Complex** and **Pseudocomplex** words with the embeddings of their putative stems. All other word-pair combinations remain the same as in the *LSA* model. This model represents *blind* morphological processing, where **Complex** and **Pseudocomplex** words have shared activations with their stems, and thus, processing **Complex** and **Pseudocomplex** derived forms triggers the activations of their putative stems.

'baker' $\rightarrow$ **'bake'**

'whisper' $\rightarrow$ **'whisp'**

- 'bake' and 'baker': $1 - \cos(\text{bake}, \text{bake}) = 0$
- 'whisp' and 'whisper': $1 - \cos(\text{whisp}, \text{whisp}) = 0$
- All additional word combinations: $1 - \cos(\textit{word1}, \textit{word2})$, but replace embeddings of derived **Complex** and **Pseudocomplex** forms with those of their stems.

The resulting RDMs were subsequently compared with brain RDMs following the procedure in Section 4.1.

4.3 Statistical Analyses on RSA Correlation Coefficients

4.3.1 Spatiotemporal cluster-based permutation t-tests

To assess the statistical significance of the RSA results, we utilized the Eelbrain Python Package (Brodbeck et al., 2023). We performed spatiotemporal cluster-based permutation t-testing using a one-sample t-test across subjects, treating each subject's RSA correlation matrix as an independent observation. The test conducts a one sample t-test at each sensor-time point, to establish whether the mean correlation coefficient value across subjects differs from zero. It subsequently identifies clusters of adjacent sensors where effects are spatially and temporally contiguous. We applied a cluster-forming threshold of $p < 0.05$, and only clusters exceeding a minimum duration of 10 ms and at least four spatially adjacent sensors were considered significant. To simulate the null hypothesis that the true mean correlation was zero, the signs of each subject's correlation values were randomly flipped across 10,000 permutations. This procedure generates a null distribution of cluster-level test statistics, against which the observed clusters are compared. This process defines the null distribution and controls for multiple comparisons, allowing us to identify clusters unlikely to have arisen by chance. Search parameters are provided in Table 4.1. For the Text modality, we restricted the theoretical analyses to the 150–400 ms time window, targeting the N170 and N250 effects. However, because lexical access is expected to occur later than prelexical processing, we use a search window of 300-800 ms for the semantic models in Text. Additional modalities were tested

within two windows divided by the average morpheme boundary or disambiguating point (580 ms) in the key conditions (100-580 ms and 580-880 ms). We again used 10,000 random resamples, a cluster-forming threshold of $p < 0.05$, a minimum cluster duration of 10 ms, and a minimum of four spatially adjacent sensors. Overall, this approach accounts for spatial and temporal dependencies in the EEG data, ensuring robust statistical inference regarding the relationship between brain representations and linguistic models.

Table 4.1: Summary of parameters used for cluster-based permutation tests across different stimulus modalities and model types.

Modality	**Model Type**	**Time Window (ms)**	**Search Space (Source)**
Text	Theoretical	100–400	LTL+LTPJ+LIFG
	Semantic	300–800	LTL+LTPJ+LIFG
Audio, AV, Silent Video	All models	100–580	Whole brain
	All models	580–880	Whole brain*

Note. Cluster-forming threshold: $p < .05$ (one-tailed). Number of permutations: 10,000. Minimum cluster size: 10 ms (temporal) and 4 adjacent sensors (spatial). AV = audiovisual. **After conducting a whole brain analysis on AV correlations, a follow-up test was conducted in the left occipitotemporal region for this modality only.*

This statistical analysis was repeated using the source reconstructed data, which is detailed in Section 3.5.2. To create source space brain RDMs, we replaced sensor-space epochs with the source time courses (STCs) of each subject. For a constrained analysis in Text, we limited the search space using the *aparc* parcellation from FreeSurfer (Fischl, 2012), encompassing the left temporal lobe (LTL), left temporoparietal junction (LTPJ), and the left inferior frontal gyrus (LIFG). Anatomical labels are provided in Appendix B.

Searches for the Audio, AV, and Silent Video modalities were all conducted using the whole brain in source space. However, one follow-up search was conducted for the AV modality encompassing the left middle and inferior temporal regions, the left fusiform gyrus, and the left lateral occipital region. Again, this used the *aparc* parcellation (see Appendix B for anatomical labels).

4.3.2 Repeated Measures One-Way ANOVA

We conducted repeated-measures one-way ANOVAs at each source-time point, to assess whether RSA correlation coefficients differed significantly across model RDMs. Similar to Section 4.3.1, the ANOVAs use spatiotemporal cluster permutation testing to identify clusters of contiguous vertices and timepoints which significantly differ from a null distribution (generated by shuffling model labels within each subject). Specifically, each ANOVA was used to compare model RDMs of the same class (either theoretical or semantic). To reduce the number of comparisons, ANOVAs were only conducted for modality-RDM-class combinations that showed significant or near-significant effects in the preceding spatiotemporal cluster-based permutation t-tests.

Chapter 5

Results

5.1 Behavioral Results

Here, we report the behavioral results of the four lexical decision experiments. Mean accuracies for each key condition (**Complex**, **Pseudocomplex**, and **Simplex**) are presented in Figure 5.1. A one-tailed binomial test confirmed that all subjects performed significantly above chance in the Text, Audio, and AV experiments (all $p < .001$). No subject performed significantly above chance for the Silent Video modality (all $p > .05$). Overall accuracy was highest in the Text ($M = 96.48\%$, $SE = 0.44\%$) modality, followed by Audio ($M = 91.80\%$, $SE = 0.73\%$), AV ($M = 83.89\%$, $SE = 2.46\%$), and Silent Video ($M = 50.15\%$, $SE = 1.40\%$).

To analyze differences in accuracy between word types within each modality, we conducted separate generalized linear models (GLMs) for Text, Audio, AV, and Silent Video modalities, with *Condition* (**Complex**, **Pseudocomplex**, **Simplex**)

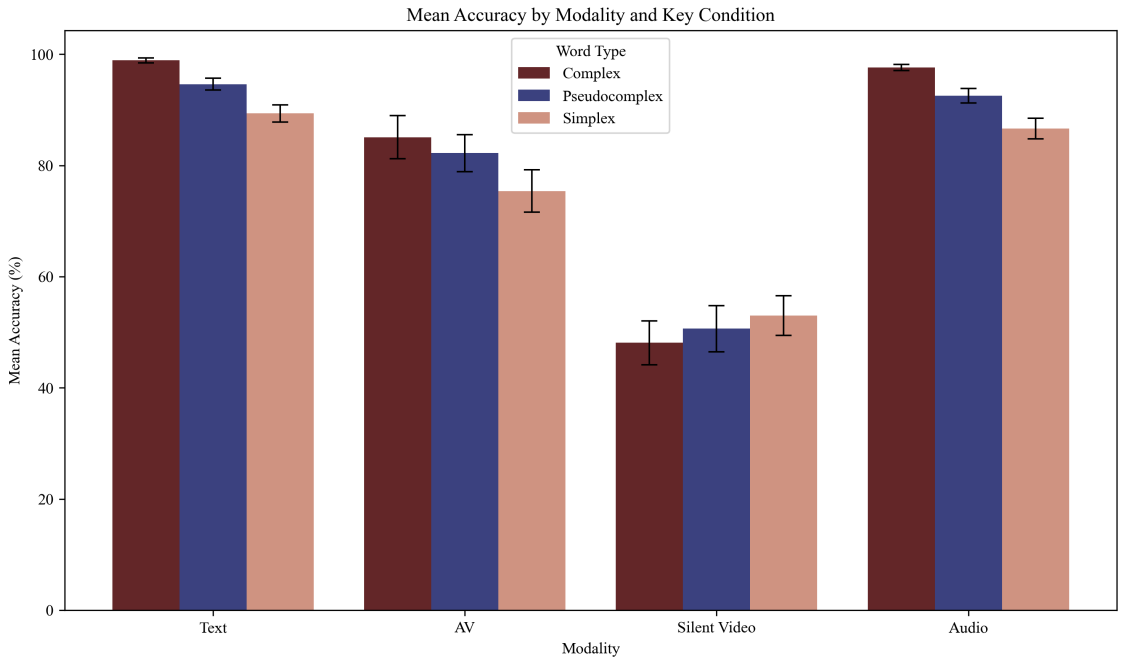

Figure 5.1: Behavioral results for the lexical decision tasks in the Text, AV, Silent Video, and Audio modalities. For each modality, we plot mean classification accuracy for each key condition.

as a fixed effect. In the Text and Audio modalities, accuracy was significantly lower for **Pseudocomplex** words compared to **Complex** words ($p < .001$) and even lower for **Simplex** words ($p < .001$). In the AV modality, accuracy was significantly lower for **Simplex** words compared to **Complex** words ($p < .001$). In the Silent Video modality, no significant differences were found between conditions ($p > .05$).

Prior work comparing the neural correlates of **Complex** and **Simplex** words discarded erroneous trials where subjects made incorrect classifications (Zweig and Pylkkänen, 2009). However, because we report systematic differences in performance across these key conditions, we opt to maintain erroneous trials in our subsequent analyses, as these trials might reflect a difficulty in processing or decomposing **Simplex** cases, rather than inattentive subjects.

5.2 RSA

Here, we report the sensor and source space findings of the representation similarity analyses. We report all significant correlations between each modality-specific brain RDM and all six model RDMs defined in Section 4. To correct for multiple comparisons within each modality, we applied Bonferroni correction for six tests (one for each model RDM). Additionally, we tested near-significant (n.s.) results that did not survive Bonferroni correction with False Discovery Rate (FDR). Primary significant findings are shown in Figure 5.2

5.2.1 Text

For the correlations generated by the Text RSAs in sensor space, spatiotemporal cluster-based permutation testing revealed a left-lateralized spatiotemporal correlation-coefficient cluster for the epoch-based brain RDMs with the LSA RDM (41 sensors, 444-799 ms, $p = .022$). Unfortunately, **this cluster was not significant** after applying Bonferroni correction for multiple (six) tests, or FDR correction (corrected $p = .065$). However, the RSAs identified a **significant**, left-lateralized spatiotemporal correlation-coefficient cluster of 37 sensors between the epoch-based brain RDMs and the LSA-PseudoStem model, from 331-799 ms, which remained significant after Bonferroni correction ($p = .007$).

Permutation t-tests on the Text RSAs in source space revealed no significant relationship with any of the six model RDMs. However, a one-way repeated-measures ANOVA identified a significant difference between the correlation coefficients of the different semantic RDMs ($p < .001$), with LSA-PseudoStem possessing higher correlations than the LSA-Stem and LSA models (left inferior temporal gyrus, 300-604 ms). It is important to note that while this specific cluster exhibits a timecourse that is, on average, positive for LSA-PseudoStem and negative for the other models, the ANOVA only displays *when and where the semantic models differ significantly from one another*. It is possible that there exist space-time courses where two or more model RDMs were positively aligned with source neural RDMs at once, but were not significantly different from zero via spatiotemporal t-tests.

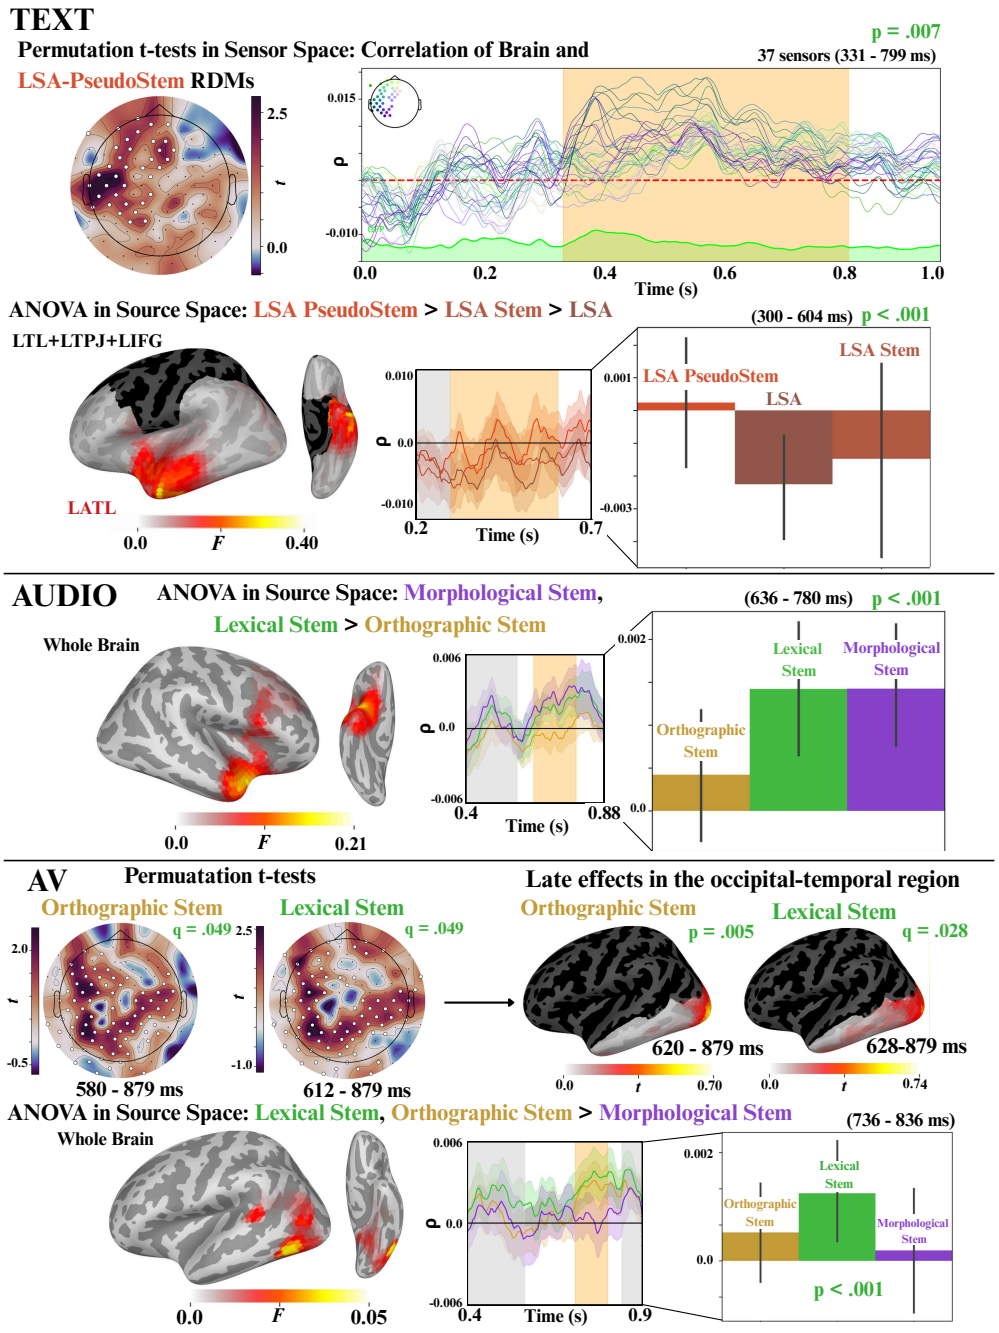

Figure 5.2: Significant statistical results on the RSA correlations for Text, Audio, and AV. Spatiotemporal cluster-based permutation t-tests indicate significant differences from a null distribution, while repeated-measures one-way ANOVAs identify differences between models. FDR-corrected results are reported as q-values.

5.2.2 Audio

The Audio RSA in sensor space revealed a spatiotemporal correlation cluster of 56 sensors between the epoch-based brain RDMs and the binary Morphological Stem RDM in a window **following the average disambiguating point**, from 580-864 ms ($p = .017$). Spatiotemporal t-tests revealed another **n.s.** cluster on the correlation coefficients of the Lexical Stem model, containing 53 sensors from 580-868 ms ($p = .028$). Unfortunately, **neither finding was significant following Bonferroni correction**. No clusters were identified in the earlier search window of 100-580 ms.

Permutation t-tests on the Audio RSAs in source space revealed no significant relationship with any of the six model RDMs. However, a one-way repeated-measures ANOVA identified a significant difference between the correlation coefficients of the different theoretical RDMs ($p < .001$), from 636-780 ms. This trend revealed the Lexical Stem and Morphological Stem correlations patterning together, over the Orthographic Stem correlations.

5.2.3 AV

The AV RSA in sensor space revealed spatiotemporal correlation clusters between the epoch-based brain RDMs and nearly all model RDMs in a window **following the average disambiguating point**. However, **none of these clusters were significant** following Bonferroni correction. To complement Bonferroni correction, which is a conservative approach, we also applied FDR correction to the strongest clusters of all six models. This yielded two statistically significant clusters following

the disambiguating point: Lexical Stem (80 sensors, 612-879 ms, corrected $p = .049$) and Orthographic Stem (79 sensors, 580-879 ms, corrected $p = .049$). No clusters were identified in the earlier search window of 100-580 ms.

Whole-brain spatiotemporal permutation t-tests did not reveal any significant clusters for the theoretical model correlation coefficients. However, a one-way repeated-measures ANOVA identified a significant difference between the correlation coefficients of the different theoretical RDMs ($p < .001$), from 736-836 ms. This trend revealed the Lexical Stem and Orthographic Stem correlations patterning together, over Morphological Stem. As this effect localized to the left occipital-temporal region, additional permutation t-tests were conducted on the theoretical model coefficients in a search space containing the left middle and inferior temporal regions, the left fusiform gyrus, and the left lateral occipital region. This identified significant clusters for the Orthographic Stem (620-879 ms, $p = .005$) and Lexical Stem (628 - 879 ms, corrected $p = .028$) models, displaying late effects in the occipital-temporal region for both models.

5.2.4 Silent Video

We found no significant spatiotemporal correlations between the brain RDMs and any of the predefined model RDMs in the Silent Video modality, for either sensor or source space. With no evidence of significant effects from the permutation tests, we did not conduct ANOVAs between the coefficients of different model RDMs.

Chapter 6

Discussion

Here, we discuss the results of the experiments and their implications. Furthermore, we outline present limitations of the study and suggest directions for future work.

The behavioral results of the lexical decision task characterize a clear trend in response accuracy, which varied significantly across the three key conditions in the Text, Audio, and AV modalities. Namely, **Complex** words (e.g., *baker*) were classified with the highest accuracy, followed by **Pseudocomplex** words (e.g., *hunger*) and then **Simplex** words (e.g., *monster*). These differences were usually statistically significant, with the exception of **Complex** and **Pseudocomplex** words in AV (which still followed the general trend of Text and Audio). The consistency of this trend across modalities provides strong evidence of morphological sensitivity in the lexical decision task, suggesting that morphological transparency and decomposability influence the ease of processing. In other words, subjects more easily recognized and judged morphologically **Complex** words when they align with expected

derivational structure. The Silent Video condition showed no significant differences across word types, likely due to the overall difficulty of the task without auditory input. These results do not indicate that morphological form can be decoded from lip reading. Text stimuli were classified with the highest accuracy ($M = 96.48\%$), suggesting that visual word reading is the easiest modality for classifying stimuli in the lexical decision task. This was followed by Audio ($M = 91.80\%$) and AV ($M = 83.89\%$). While it might be expected that the inclusion of visual information during audio presentation would improve classification accuracy (Sumby and Pollack, 1954; Munhall et al., 2004), that was not the case in our AV experiment. It is important to note that many studies showing the benefit of visual cues in AV presentation involve noisy auditory environments, which was not a quality of our lexical decision task. Subjects' expectations about the study (i.e., the inclusion of putative (pseudo)stems) could have interfered with accuracy. The classification accuracy was lowest for the Silent Video modality ($M = 50.15\%$), reflective of a random guess or chance-level average performance across subjects.

Looking towards the results of the statistical tests on the RSA correlation coefficients, we see evidence of the full decomposition model in visual word reading. This is provided by the significant spatiotemporal correlation between the sensor-based neural RDMs and the LSA-PseudoStem model, showing a positive trend between 331-799 ms in the left hemisphere. The LSA-PseudoStem model normalizes cosine comparisons so that **Complex** forms and their stems are treated as semantically identical, along with **Pseudocomplex** forms and their pseudostems. Thus, alignment between this model RDM and the neural RDMs suggests similar semantic

processing of these conditions and their putative stems. The time course of our observed effect mimics the M350 lexeme lookup effect (Stockall and Marantz, 2006), which is considered the earliest stage of stem activation in the full decomposition model. Unfortunately, the spatiotemporal cluster-based permutation tests in source space did not replicate the same finding. A repeated-measures one-way ANOVA did probe a cluster in the left anterior temporal lobe (LATL) where the coefficients of LSA-Pseudo diverge from the other semantic models. Evidence suggests that LATL acts as a hub for interpreting information into complex semantic memories (i.e., the construction of the meaning of words) (Bonner and Price, 2013). However, it is important to note that this observed effect only identifies *when and where* the alignment of neural activity and the three semantic model RDMs diverge. Thus, it is possible that all semantic models align with the source space STCs simultaneously, but that they do not differ significantly within those clusters.

Hypothesis 1 (H1) proposed that in visual word reading, morphological decomposition would apply to **Complex** and **Pseudocomplex** words, but not to **Simplex** cases, with early effects emerging in the left temporal region, followed by lexeme lookup and semantic interpretation. **The results of the RSA correlation analyses provide support for H1**. While we did not capture early effects of morphological decomposition, results indicate the activation of (pseudo)stems ∼330 ms post-stimulus onset.

In Audio, no spatiotemporal clusters of correlation coefficients significantly aligned with any predefined model RDM in sensor or source space. However, near-significant effects for the Morphological Stem (580-864 ms, $p = .017$) and Lexical Stem (580-

868 ms, $p = .028$) models emerge following the average morpheme boundary of the stimuli in sensor space. A repeated-measures one-way ANOVA over the whole brain in source space supports alignment with these theoretical models, identifying a spatiotemporal cluster where the alignment of these models differ significantly ($p < .001$). In this observed cluster, the Morphological Stem and Lexical Stem coefficients pattern together, displaying higher alignment than the Orthographic Stem model in the right anterior temporal lobe (RATL; 636-780 ms). Although the left anterior temporal lobe is commonly associated with semantic interpretation, there is evidence that RATL is similarly involved with semantic representations (Rice et al., 2018; Lambon Ralph et al., 2010). A meta-analysis from Visser et al. (2010) indicates *both* anterior temporal lobes (ATLs) as semantic hubs, noting an amodal sensitivity that persists for spoken word presentation. An MEG study of semantic dementia presented subjects with words which could not be identified until the final phoneme (*played, plate*). Source results suggested that both ATLs are necessary for word identification, with patients displaying higher RATL activity than a control group following the disambiguating point (final phoneme). Given these findings, it is possible that the Morphological Stem and Lexical Stem models are tapping into semantic processing of whole words and stems following the disambiguating point in audio processing. As nonword entries were included in the Orthographic Stem model (*monst*), it is unsurprising that this RDM does not comparably align within the given spatiotemporal region. However, it is unclear why the theoretical RDMs would probe an effect of semantic processing, rather than one of semantic LSA-models. Thus, there is equal reason to believe that the Audio results signify mor-

phological decomposition, which the Morphological and Lexical Stem models were designed to identify. It is important to reinstate that none of the standalone model RDMs significantly aligned with the neural data recorded during audio processing, and thus, interpretation should remain cautious.

Statistical tests on the RSA correlations for the AV modality provide consistent results in sensor and source space: the Lexical Stem and Orthographic Stem models significantly align with neural trends. Both (1) a whole-brain repeated measures ANOVA over theoretical RDMs and (2) a constrained spatiotemporal analysis indicate that this effect localizes to the occipital-temporal region, following the disambiguating point. The recruitment of the occipital temporal region is not implausible, as early sensory cortices have been shown to be involved in integrating audiovisual information (Ghazanfar and Schroeder, 2006; Luo et al., 2010). While the occipital temporal cortex is commonly associated with reading, there is evidence that this region plays a larger role in integrating visual information with speech and semantics (Price and Devlin, 2011). The alignment of the Lexical Stem and Orthographic Stem models extended beyond that of the Morphological Stem model. Notably, the Lexical Stem and Orthographic Stem RDMs both suggest similar processing of **Pseudocomplex** words and their pseudostems, while the Morphological Stem RDM treats these cases as distinct. Thus, this finding might indicate that the decomposition of **Pseudocomplex** words is a critical component of the decomposition process in audio-visual processing.

Hypothesis 2 (H2) proposed that morphological decomposition would occur in the audio and audio-visual modalities, emerging after the disambiguating morpheme

boundary and localizing to the left temporal region. **H2 was supported for the audio-visual modality**, where robust effects of morphological decomposition were observed following the disambiguating point. While the audio modality did not yield statistically significant effects, the n.s. trends and source localization results suggest that audio processing may support morphological decomposition. Together, these findings provide partial support for H2, warranting further investigation.

The absence of a significant or n.s. effects in the Silent Video modality provides an important counterpoint to the decomposition effects observed in Text, Audio, and AV conditions. The lack of a finding aligns with the corresponding behavioral results of the lexical decision task, where we found no evidence that subjects successfully recognized words during lipreading. These results imply that the visemic sequences present in the Silent Video modality lack the granularity necessary to store morpho-orthographic information, or that such information is present but cannot be decoded by the human brain. While prior works suggest that the auditory cortex may synchronize with the absent speech envelope during lipreading (Bourguignon et al., 2020), the present study does not provide evidence that such entrainment extends to the level of morphological structure. Given this, **we reject Hypothesis 3**, which proposed that silent visual speech would support morphological decomposition, with neural patterns in auditory regions reflecting morphological parsing in the absence of sound.

6.1 Limitations and Future Directions

The experiments and analyses included in this thesis are prone to a wide range of limitations. First, we note that the theoretical binary RDMs did not successfully probe early morpho-orthographic decomposition effects in Text. While these models were useful for establishing trends in the Audio and AV conditions, we were unable to make claims about the earliest stages of decomposition in reading. To that end, the findings in Text provide clarity on the post-decomposition effects only, such as the access of pseudostems, which can be used to infer morphological parsing. It is possible that our sample size ($N = 21$) was insufficient in probing more effects, as neural data is highly variable across subjects, and detecting trends could require a larger sample. To counter this point, the representational similarity analysis is expected to circumvent issues of inter-subject variability by abstracting subject-wise patterns. However, we also note emerging yet insignificant trends in the Audio condition, which could signal the need for more data.

The use of representational similarity analysis is accompanied by its own limitations. While RSA is a powerful tool for abstracting representational structures in neural data, it presents several challenges. First, RSA conducts *correlations over correlations*, and is thus largely susceptible to common issues associated with correlation-based analyses, like outlier influence (Popal et al., 2019). Compared to traditional neuroscience approaches that focus on direct measures of neural signal, RSA imposes a level abstraction and is (comparatively) computationally expensive. In the case of our spatiotemporal cluster-based permutation t-tests on RSA correlation coefficients, it is possible to find multiple significant clusters, even amongst

models that conflict conceptually. Moreover, the use of a two-tailed t-test could hypothetically yield negative clusters, signifying misalignment between how the brain and model RDMs relate word items. Consider the LSA RDM, which treats **Pseudocomplex** words and their stems as *dissimilar* (inherently, these pairs are semantically opaque). If the neural data, however, suggests *similar* processing for **Pseudocomplex** words and their stems in a given spatiotemporal region (as we observe in our findings), the RSA would indicate misalignment. While this possibility is conceptually plausible, it does convolute interpretation. As our primary interest in the current study was to identify when and where the brain and model RDMs aligned, we simply conducted one-tailed tests.

We also note that only GloVe was used for the semantic RSA, though other models were considered (GPT-2). While the use of large language models like GPT-2 are fit for contextually-rich stimuli and regression-based encoding analyses, these models were not appropriate for our single-word stimuli. One limitation is that the tokenization schemes used by LLMs rarely consider morphological structure. However, new advances towards morphology-informed tokenization have emerged, and are worth exploring in future studies (Jabbar, 2023). Regardless, additional static word embedding models could have been applied to the current stimuli.

In the present study, we conducted a limited number of constrained source space analyses in general cortical areas. However, the literature identifies several key regions in the brain that might be involved in morphological decomposition. With the goal of identifying missed effects, future work could conduct spatiotemporal cluster-based permutation t-tests in specific functional regions of interest (fROIs).

Chapter 7

Conclusion

This thesis explored how the brain processes morphological structure across reading, listening, and visual speech recognition using HD-EEG with source localization. Representational similarity analyses support a model of full morphological decomposition during visual word reading, showing that both complex (*baker*) and pseudocomplex (*whisper*) words engage early form-based processing, regardless of semantic transparency. Specifically, we provide evidence that reading 'whisper' activates 'whisp', indexing a post-decomposition effect ∼330 ms post-stimulus onset in the left-hemisphere. This response seemingly aligns with the lexeme-lookup step of the full decomposition model. Our approach demonstrates the capacity of lexical semantic association RSAs to probe signatures of lexical access and semantic processing in the brain. While results in the Audio condition were largely insignificant, emerging trends might indicate decomposition effects following the average morpheme boundary of our stimuli. In comparison, we observe late form-based de-

composition effects for the AV condition in the occipital temporal region, following the disambiguating morpheme boundary. This condition yielded agreement in sensor and source space, indicating that the decomposition of pseudocomplex words is a crucial component of the decomposition process. Silent lipreading, however, did not yield morphological effects detectable via RSA. Taken with the behavioral results of the lexical task for this Condition, we conclude that visual cues alone may not robustly support morphological information.

Overall, this work demonstrates how representational similarity analyses and natural language processing techniques can be leveraged to probe patterns of neural activation associated with linguistic processing. By applying these tools across multiple modalities, we show that the brain engages in modality-dependent strategies for morphological decomposition, with robust effects in reading and audiovisual conditions. These findings delineate the temporal dynamics that underlie the processing of morphologically complex, pseudocomplex, and simplex words. Future work should extend these efforts by incorporating additional representational models and experimental techniques.

Bibliography

Martijn Baart and Arthur G Samuel. Turning a blind eye to the lexicon: Erps show no cross-talk between lip-read and lexical context during speech sound processing. *Journal of Memory and Language*, 85:42–59, 2015.

R Harald Baayen, Richard Piepenbrock, and Leon Gulikers. Celex2. *Linguistic Data Consortium, Philadelphia*, 1996.

LE Bernstein and E Liebenthal. Neural pathways for visual speech perception. front neurosci 8: 386, 2014.

Paul Boersma and Vincent Van Heuven. Speak and unspeak with praat. *Glot International*, 5(9/10):341–347, 2001.

Michael F Bonner and Amy R Price. Where is the anterior temporal lobe and what does it do? *Journal of Neuroscience*, 33(10):4213–4215, 2013.

Mathieu Bourguignon, Martijn Baart, Efthymia C Kapnoula, and Nicola Molinaro. Lip-reading enables the brain to synthesize auditory features of unknown silent speech. *Journal of Neuroscience*, 40(5):1053–1065, 2020.

Silvia Brem, Pascal Halder, Kerstin Bucher, Paul Summers, Ernst Martin, and Daniel Brandeis. Tuning of the visual word processing system: distinct developmental erp and fmri effects. *Human brain mapping*, 30(6):1833–1844, 2009.

Christian Brodbeck, Proloy Das, Marlies Gillis, Joshua P Kulasingham, Shohini Bhattasali, Phoebe Gaston, Philip Resnik, and Jonathan Z Simon. Eelbrain, a python toolkit for time-continuous analysis with temporal response functions. *Elife*, 12:e85012, 2023.

Brian Butterworth. *Lexical representation*. 1983.

Joan Bybee. Regular morphology and the lexicon. *Language and cognitive processes*, 10(5):425–455, 1995.

Joan L. Bybee. *Morphology: A Study of the Relation Between Meaning and Form*. John Benjamins, Philadelphia, PA, 1985.

Joan L Bybee. Morphology as lexical organization. *Theoretical morphology*, 119141, 1988.

Gemma A Calvert, Edward T Bullmore, Michael J Brammer, Ruth Campbell, Steven CR Williams, Philip K McGuire, Peter WR Woodruff, Susan D Iversen, and Anthony S David. Activation of auditory cortex during silent lipreading. *science*, 276(5312):593–596, 1997.

Ruth Campbell. The processing of audio-visual speech: empirical and neural bases. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 363(1493): 1001–1010, 2008.

Alfonso Caramazza, Alessandro Laudanna, and Cristina Romani. Lexical access and inflectional morphology. *Cognition*, 28(3):297–332, 1988.

Charlotte Caucheteux and Jean-Rémi King. Brains and algorithms partially converge in natural language processing. *Communications biology*, 5(1):134, 2022.

Kenneth Ward Church. Word2vec. *Natural Language Engineering*, 23(1):155–162, 2017.

David Cohen and B Neil Cuffin. Eeg versus meg localization accuracy: theory and experiment. *Brain topography*, 4(2):95–103, 1991.

Laurent Cohen, Stanislas Dehaene, Lionel Naccache, Stéphane Lehéricy, Ghislaine Dehaene-Lambertz, Marie-Anne Hénaff, and François Michel. The visual word form area: spatial and temporal characterization of an initial stage of reading in normal subjects and posterior split-brain patients. *Brain*, 123(2):291–307, 2000.

Ava Creemers and David Embick. Retrieving stem meanings in opaque words during auditory lexical processing. *Language, Cognition and Neuroscience*, 36(9):1107–1122, 2021.

Ava Creemers, Amy Goodwin Davies, Robert J Wilder, Meredith Tamminga, and David Embick. Opacity, transparency, and morphological priming: A study of prefixed verbs in dutch. *Journal of Memory and Language*, 110:104055, 2020.

Ava Creemers, Nattanun Chanchaochai, Meredith Tamminga, and David Embick. The activation of embedded (pseudo-) stems in auditory lexical processing: impli-

cations for models of spoken word recognition. *Language, Cognition and Neuroscience*, 38(7):966–982, 2023.

Michael J Crosse, Hesham A ElShafei, John J Foxe, and Edmund C Lalor. Investigating the temporal dynamics of auditory cortical activation to silent lipreading. In *2015 7th International IEEE/EMBS Conference on Neural Engineering (NER)*, pages 308–311. IEEE, 2015.

Serena Dattola, Francesco Carlo Morabito, Nadia Mammone, and Fabio La Foresta. Findings about loreta applied to high-density eeg—a review. *Electronics*, 9(4):660, 2020.

David Embick and Alec Marantz. Cognitive neuroscience and the english past tense: Comments on the paper by ullman et al. *Brain and language*, 93(2):243–247, 2005.

Kawin Ethayarajh. How contextual are contextualized word representations? comparing the geometry of bert, elmo, and gpt-2 embeddings, 2019. URL `https://arxiv.org/abs/1909.00512`.

Allyson Ettinger, Tal Linzen, and Alec Marantz. The role of morphology in phoneme prediction: Evidence from meg. *Brain and language*, 129:14–23, 2014.

Bruce Fischl. Freesurfer. *Neuroimage*, 62(2):774–781, 2012.

Kenneth I Forster and Chris Davis. Repetition priming and frequency attenuation in lexical access. *Journal of experimental psychology: Learning, Memory, and Cognition*, 10(4):680, 1984.

Ram Frost, Kenneth I Forster, and Avital Deutsch. What can we learn from the morphology of hebrew? a masked-priming investigation of morphological representation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23(4):829, 1997.

Asif A Ghazanfar and Charles E Schroeder. Is neocortex essentially multisensory? *Trends in cognitive sciences*, 10(6):278–285, 2006.

Alexandre Gramfort, Martin Luessi, Eric Larson, Denis A Engemann, Daniel Strohmeier, Christian Brodbeck, Roman Goj, Mainak Jas, Teon Brooks, Lauri Parkkonen, et al. Meg and eeg data analysis with mne-python. *Frontiers in Neuroinformatics*, 7:267, 2013.

Ken W Grant and Steven Greenberg. Speech intelligibility derived from asynchronous processing of auditory-visual information. In *AVSP*, volume 200, pages 132–137, 2001.

Ken W Grant, Virginie van Wassenhove, and David Poeppel. Discrimination of auditory-visual synchrony. In *AVSP*, pages 31–35, 2003.

Roberta Grech, Tracey Cassar, Joseph Muscat, Kenneth P Camilleri, Simon G Fabri, Michalis Zervakis, Petros Xanthopoulos, Vangelis Sakkalis, and Bart Vanrumste. Review on solving the inverse problem in eeg source analysis. *Journal of neuroengineering and rehabilitation*, 5:1–33, 2008.

Laura Gwilliams and Alec Marantz. Non-linear processing of a linear speech stream:

The influence of morphological structure on the recognition of spoken arabic words. *Brain and language*, 147:1–13, 2015.

Laura Gwilliams and Alec Marantz. Morphological representations are extrapolated from morpho-syntactic rules. *Neuropsychologia*, 114:77–87, 2018.

Nima Hadidi, Ebrahim Feghhi, Bryan H Song, Idan A Blank, and Jonathan C Kao. Illusions of alignment between large language models and brains emerge from fragile methods and overlooked confounds. *bioRxiv*, pages 2025–03, 2025.

Riitta Hari and Olli V Lounasmaa. Recording and interpretation of cerebral magnetic fields. *Science*, 244(4903):432–436, 1989.

Haris Jabbar. Morphpiece: A linguistic tokenizer for large language models. *arXiv preprint arXiv:2307.07262*, 2023.

Nancy Kanwisher, Josh McDermott, and Marvin M Chun. The fusiform face area: a module in human extrastriate cortex specialized for face perception. *Journal of neuroscience*, 17(11):4302–4311, 1997.

Chamandeep Kaur, Preeti Singh, Amandeep Bisht, Garima Joshi, and Sunil Agrawal. Recent developments in spatio-temporal eeg source reconstruction techniques. *Wireless personal communications*, 122(2):1531–1558, 2022.

Silke Klamer, Adham Elshahabi, Holger Lerche, Christoph Braun, Michael Erb, Klaus Scheffler, and Niels K Focke. Differences between meg and high-density eeg source localizations using a distributed source model in comparison to fmri. *Brain topography*, 28:87–94, 2015.

Alexandra Krauska and Ellen Lau. Moving away from lexicalism in psycho-and neuro-linguistics. *Frontiers in Language Sciences*, 2:1125127, 2023.

Nikolaus Kriegeskorte, Marieke Mur, and Peter A Bandettini. Representational similarity analysis-connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, 2:249, 2008.

Matthew A Lambon Ralph, Lisa Cipolotti, Facundo Manes, and Karalyn Patterson. Taking both sides: do unilateral anterior temporal lobe lesions disrupt semantic memory? *Brain*, 133(11):3243–3255, 2010.

Ellen F Lau, Colin Phillips, and David Poeppel. A cortical network for semantics:(de) constructing the n400. *Nature reviews neuroscience*, 9(12):920–933, 2008.

KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. Videochat: Chat-centric video understanding. *arXiv preprint arXiv:2305.06355*, 2023.

Catherine-Marie Longtin, Juan Segui, and Pierre A Hallé. Morphological priming without morphological relationship. *Language and cognitive processes*, 18(3):313–334, 2003.

Huan Luo, Zuxiang Liu, and David Poeppel. Auditory cortex tracks both auditory and visual stimulus dynamics using low-frequency neuronal phase modulation. *PLoS biology*, 8(8):e1000445, 2010.

Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan.

Video-chatgpt: Towards detailed video understanding via large vision and language models. *arXiv preprint arXiv:2306.05424*, 2023.

Leon Manelis and David A Tharp. The processing of affixed words. *Memory & Cognition*, 5(6):690–695, 1977.

William Marslen-Wilson, Lorraine K Tyler, Rachelle Waksler, and Lianne Older. Morphology and meaning in the english mental lexicon. *Psychological review*, 101 (1):3, 1994.

Michael McAuliffe, Michaela Socolof, Sarah Mihuc, Michael Wagner, and Morgan Sonderegger. Montreal forced aligner: Trainable text-speech alignment using kaldi. In *Interspeech*, volume 2017, pages 498–502, 2017.

Harry McGurk and John MacDonald. Hearing lips and seeing voices. *Nature*, 264 (5588):746–748, 1976.

Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*, 2016.

Ezequiel Mikulan, Simone Russo, Sara Parmigiani, Simone Sarasso, Flavia Maria Zauli, Annalisa Rubino, Pietro Avanzini, Anna Cattani, Alberto Sorrentino, Steve Gibbs, et al. Simultaneous human intracerebral stimulation and hd-eeg, ground-truth for source localization methods. *Scientific data*, 7(1):127, 2020.

Tara Mohammed, Ruth Campbell, Mairéad Macsweeney, Fiona Barry, and Michael Coleman. Speechreading and its association with reading among deaf, hearing and dyslexic individuals. *Clinical linguistics & phonetics*, 20(7-8):621–630, 2006.

Sophie Molholm and John J Foxe. Look 'hear', primary auditory cortex is active during lip-reading. *Neuroreport*, 16(2):123–124, 2005.

Joanna Morris and Linnaea Stockall. Early, equivalent erp masked priming effects for regular and irregular morphology. *Brain and language*, 123(2):81–93, 2012.

Joanna Morris, Jonathan Grainger, and Phillip J Holcomb. An electrophysiological investigation of early effects of masked morphological priming. *Language and Cognitive Processes*, 23(7-8):1021–1056, 2008.

Kevin G Munhall, Jeffery A Jones, Daniel E Callan, Takaaki Kuratate, and Eric Vatikiotis-Bateson. Visual prosody and speech intelligibility: Head movement improves auditory speech perception. *Psychological science*, 15(2):133–137, 2004.

Graham A Murrell and John Morton. Word recognition and morphemic structure. *Journal of Experimental Psychology*, 102(6):963, 1974.

RC Oldfield. Edinburgh handedness inventory. *Journal of Abnormal Psychology*, 1971.

Rachel Ostrand, Sheila E Blumstein, Victor S Ferreira, and James L Morgan. What you see isn't always what you get: Auditory word signals trump consciously perceived words in lexical access. *Cognition*, 151:96–107, 2016.

Roberto Domingo Pascual-Marqui et al. Standardized low-resolution brain electromagnetic tomography (sloreta): technical details. *Methods find exp clin pharmacol*, 24(Suppl D):5–12, 2002.

Jonathan W Peirce. Psychopy—psychophysics software in python. *Journal of neuroscience methods*, 162(1-2):8–13, 2007.

Johanna Pekkola, Ville Ojanen, Taina Autti, Iiro P Jääskeläinen, Riikka Möttönen, Antti Tarkiainen, and Mikko Sams. Primary auditory cortex activation by visual speech: an fmri study at 3 t. *Neuroreport*, 16(2):125–128, 2005.

Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543, 2014.

Steven Pinker. Words and rules. *Lingua*, 106(1-4):219–242, 1998.

Haroon Popal, Yin Wang, and Ingrid R Olson. A guide to representational similarity analysis for social neuroscience. *Social cognitive and affective neuroscience*, 14(11): 1243–1253, 2019.

Cathy J Price and Joseph T Devlin. The interactive account of ventral occipitotemporal contributions to reading. *Trends in cognitive sciences*, 15(6):246–253, 2011.

Kathleen Rastle and Matthew H Davis. Reading morphologically complex words. *Masked priming: The state of the art*, pages 279–305, 2003.

Kathleen Rastle, Matt H Davis, William D Marslen-Wilson, and Lorraine K Tyler. Morphological and semantic effects in visual word recognition: A time-course study. *Language and cognitive processes*, 15(4-5):507–537, 2000.

Kathleen Rastle, Matthew H Davis, and Boris New. The broth in my brother's brothel: Morpho-orthographic segmentation in visual word recognition. *Psychonomic bulletin & review*, 11:1090–1098, 2004.

Grace E Rice, Helen Caswell, Perry Moore, Paul Hoffman, and Matthew A Lambon Ralph. The roles of left versus right anterior temporal lobes in semantic memory: a neuropsychological comparison of postsurgical temporal lobe epilepsy patients. *Cerebral Cortex*, 28(4):1487–1501, 2018.

Jona Sassenhagen and Christian J Fiebach. Traces of meaning itself: Encoding distributional word vectors in brain activity. *Neurobiology of Language*, 1(1):54–76, 2020.

Robert Schreuder and R Harald Baayen. How complex simplex words can be. *Journal of memory and language*, 37(1):118–139, 1997.

Martin Schrimpf, Idan Asher Blank, Greta Tuckute, Carina Kauf, Eghbal A Hosseini, Nancy Kanwisher, Joshua B Tenenbaum, and Evelina Fedorenko. The neural architecture of language: Integrative modeling converges on predictive processing. *Proceedings of the National Academy of Sciences*, 118(45):e2105646118, 2021.

Olla Solomyak and Alec Marantz. Evidence for early morphological decomposition in visual word recognition. *Journal of Cognitive Neuroscience*, 22(9):2042–2057, 2010.

Charles Spearman. The proof and measurement of association between two things. 1961.

Linnaea Stockall and Alec Marantz. A single route, full decomposition model of morphological complexity: Meg evidence. *The mental lexicon*, 1(1):85–123, 2006.

Linnaea Stockall, Christina Manouilidou, Laura Gwilliams, Kyriaki Neophytou, and Alec Marantz. Prefix stripping re-re-revisited: Meg investigations of morphological decomposition and recomposition. *Frontiers in Psychology*, 10:1964, 2019.

William H Sumby and Irwin Pollack. Visual contribution to speech intelligibility in noise. *The journal of the acoustical society of america*, 26(2):212–215, 1954.

Quentin Summerfield. Lipreading and audio-visual speech perception. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 335 (1273):71–78, 1992.

Marcus Taft and Kenneth I Forster. Lexical storage and retrieval of prefixed words. *Journal of verbal learning and verbal behavior*, 14(6):638–647, 1975.

Antti Tarkiainen, Päivi Helenius, Peter C Hansen, Piers L Cornelissen, and Riitta Salmelin. Dynamics of letter string perception in the human occipitotemporal cortex. *Brain*, 122(11):2119–2132, 1999.

Suramya Tomar. Converting video formats with ffmpeg. *Linux journal*, 2006(146): 10, 2006.

Virginie Van Wassenhove, Ken W Grant, and David Poeppel. Visual speech speeds up the neural processing of auditory speech. *Proceedings of the National Academy of Sciences*, 102(4):1181–1186, 2005.

Maya Visser, Elizabeth Jefferies, and MA Lambon Ralph. Semantic processing in the anterior temporal lobes: a meta-analysis of the functional neuroimaging literature. *Journal of cognitive neuroscience*, 22(6):1083–1094, 2010.

Xiaosha Wang, Yangwen Xu, Yuwei Wang, Yi Zeng, Jiacai Zhang, Zhenhua Ling, and Yanchao Bi. Representational similarity analysis reveals task-dependent semantic influence of the visual word form area. *Scientific reports*, 8(1):3047, 2018.

Yanjun Wei, Ying Niu, Marcus Taft, and Manuel Carreiras. Morphological decomposition in chinese compound word recognition: Electrophysiological evidence. *Brain and Language*, 241:105267, 2023.

Kevin Whittingstall, Gerhard Stroink, Larry Gates, JF Connolly, and Allen Finley. Effects of dipole position, orientation and noise on the accuracy of eeg source localization. *Biomedical engineering online*, 2:1–5, 2003.

Hang Zhang, Xin Li, and Lidong Bing. Video-llama: An instruction-tuned audio-visual language model for video understanding. *arXiv preprint arXiv:2306.02858*, 2023.

Eytan Zweig and Liina Pylkkänen. A visual m170 effect of morphological complexity. *Language and Cognitive Processes*, 24(3):412–439, 2009.

Appendix A

Word List by Condition

Morphologically Complex

drinker	farmer	sleeper	winner	loser
killer	printer	teacher	weaver	boiler
baker	sinner	caller	catcher	charmer
bidder	freezer	digger	hunter	diver
painter	dancer	heater	faker	singer
rider	preacher	founder	boxer	driver

Morphologically Pseudocomplex

locker	archer	bumper	scooter	breather
folder	buster	porter	cower	patter
hunger	gutter	joiner	plunger	poker
prayer	sweater	sneaker	rubber	drawer
header	merger	bouncer	steamer	teller

ranger	poster	counter	charter	slipper

Simplex

anger	filter	clatter	beaver	finger
cancer	platter	roster	fever	ginger
silver	danger	monster	blunder	banter
lumber	feather	gender	canker	leper
roger	garter	plaster	tiger	holster
leather	fodder	cluster	bluster	panther

Monomorphemic Controls

scramble	assay	middle	notion	convent
tangle	balance	donate	rumble	barrow
beckon	cannon	compare	curtain	salute
socket	adapt	drizzle	pheasant	applaud
rocket	margin	button	bubble	witness
bargain	almond	billiard	trickle	puddle

Stems of Morphologically Complex Words

charm	call	preach	sleep	bid
print	ride	dance	freeze	sin
heat	teach	sing	paint	found
bake	fake	kill	dig	weave
dive	hunt	farm	win	lose
drive	box	catch	boil	drink

Stems of Morphologically Pseudocomplex Words

pray	plunge	join	lock	sneak
merge	range	tell	post	draw
arch	hung	scoot	bounce	chart
gut	breathe	sweat	pat	bump
count	port	bust	cow	rub
head	slip	poke	fold	steam

Stems of Simplex Words

lep	dange	silv	panth	ang
feath	ginge	gend	blund	cank
clust	monst	fodd	roge	feve
plast	blust	platt	gart	tige
fing	leath	rost	holst	filt
cance	beav	bant	clatt	lumb

Bimorphemic '-er' Nonwords

fanper	poder	ferder	ranper	doxer
borter	veaber	kerger	loiner	glimper
jaikler	soger	supner	jinner	seaper
rolber	lanker	gifner	jutter	jomper
joper	lanper	foiner	poger	dincer
scronker	lupner	veader	doobler	ferger
blicer	stuber	folber	pleeper	faikler

soder	lodder	luker	daker	saker
panger	jupner	diser	tuster	paller
dorter	kuster	panker	junger	janfer
dainter	sanfer	sopner	dupner	marfer
jocker	doker	rimper	sepler	jerter
tepder	toger	jorter	bupler	soper
lirter	pocker	vilser	jander	hepder
loxer	paker	blimper	pilser	hutter
busler	vapter	loker	jopner	topner
mider	herger	tuger	daler	faidler
fipper	kaber	dounter	plimper	vanker

Additional Bimorphemic Nonwords

cindam	notin	harkin	clustid	roppit
gendal	parlin	filtan	pubble	pangle
bolate	siddle	dipple	subim	rupit
zopin	fonape	dalute	labbit	lonate
putton	furtain	silvon	hiddle	fipul
pabbit	lopin	sutton	kalute	lobin

Monomorphemic Nonwords

triv	drint	crile	plint	crant
drobb	blor	snope	glim	blunk
zorn	gront	zate	grint	snarp

smick	crint	zwint	plort	fross
brint	crive	sniv	clib	plim
plang	treel	snop	zort	smorn

Appendix B

Anatomical Labels used in Constrained Source Space Analyses

The following anatomical labels were included in the constrained search space analysis using the *aparc* parcellation:

In Text: left temporal lobe (LTL), left temporoparietal junction (LTPJ), and the left inferior frontal gyrus (LIFG):

- *lateraloccipital-lh*
- *cuneus-lh*
- *lingual-lh*
- *pericalcarine-lh*
- *fusiform-lh*
- *middletemporal-lh*
- *inferiortemporal-lh*
- *entorhinal-lh*
- *bankssts-lh*
- *inferiorparietal-lh*
- *insula-lh*
- *superiortemporal-lh*

- *supramarginal-lh*
- *temporalpole-lh*
- *transversetemporal-lh*
- *parsopercularis-lh*
- *parsorbitalis-lh*
- *parstriangularis-lh*
- *caudalmiddlefrontal-lh*
- *rostralmiddlefrontal-lh*
- *lateralorbitofrontal-lh*
- *medialorbitofrontal-lh*

In AV: left middle and inferior temporal regions, the left fusiform gyrus, and the left lateral occipital region:

- *middletemporal-lh*
- *inferiortemporal-lh*
- *ateraloccipital-lh*
- *fusiform-lh*