



Decoding EEG Brain Activity for Multi-Modal Natural Language Processing

Nora Hollenstein^{1*}, Cedric Renggli², Benjamin Glaus², Maria Barrett³, Marius Troendle⁴, Nicolas Langer⁴ and Ce Zhang²

¹ Department of Nordic Studies and Linguistics, University of Copenhagen, Copenhagen, Denmark, ² Department of Computer Science, Swiss Federal Institute of Technology, ETH Zurich, Zurich, Switzerland, ³ Department of Computer Science, IT University of Copenhagen, Copenhagen, Denmark, ⁴ Department of Psychology, University of Zurich, Zurich, Switzerland

OPEN ACCESS

Edited by:

Lars Kuchinke,
International Psychoanalytic University
Berlin, Germany

Reviewed by:

Stéphane Dufau,
Centre National de la Recherche
Scientifique (CNRS), France
Yohei Oseki,
The University of Tokyo, Japan

*Correspondence:

Nora Hollenstein
nora.hollenstein@hum.ku.dk

Specialty section:

This article was submitted to
Speech and Language,
a section of the journal
Frontiers in Human Neuroscience

Received: 27 January 2021

Accepted: 14 June 2021

Published: 13 July 2021

Citation:

Hollenstein N, Renggli C, Glaus B,
Barrett M, Troendle M, Langer N and
Zhang C (2021) Decoding EEG Brain
Activity for Multi-Modal Natural
Language Processing.
Front. Hum. Neurosci. 15:659410.
doi: 10.3389/fnhum.2021.659410

Until recently, human behavioral data from reading has mainly been of interest to researchers to understand human cognition. However, these human language processing signals can also be beneficial in machine learning-based natural language processing tasks. Using EEG brain activity for this purpose is largely unexplored as of yet. In this paper, we present the first large-scale study of systematically analyzing the potential of EEG brain activity data for improving natural language processing tasks, with a special focus on which features of the signal are most beneficial. We present a multi-modal machine learning architecture that learns jointly from textual input as well as from EEG features. We find that filtering the EEG signals into frequency bands is more beneficial than using the broadband signal. Moreover, for a range of word embedding types, EEG data improves binary and ternary sentiment classification and outperforms multiple baselines. For more complex tasks such as relation detection, only the contextualized BERT embeddings outperform the baselines in our experiments, which raises the need for further research. Finally, EEG data shows to be particularly promising when limited training data is available.

Keywords: EEG, natural language processing, frequency bands, brain activity, machine learning, multi-modal learning, physiological data, neural network

1. INTRODUCTION

Recordings of brain activity play an important role in furthering our understanding of how human language works (Murphy et al., 2018; Ling et al., 2019). The appeal and added value of using brain activity signals in linguistic research are intelligible (Stemmer and Connolly, 2012). Computational language processing models still struggle with basic linguistic phenomena that humans perform effortlessly (Ettinger, 2020). Combining insights from neuroscience and artificial intelligence will take us closer to human-level language understanding (McClelland et al., 2020). Moreover, numerous datasets of cognitive processing signals in naturalistic experiment paradigms with real-world language understanding tasks are becoming available (Alday, 2019; Kandylaki and Bornkessel-Schlesewsky, 2019).

Linzen (2020) advocates for the grounding of NLP models in multi-modal settings to compare the generalization abilities of the models to human language learning. Multi-modal learning in machine learning refers to algorithms learning from multiple input modalities encompassing various aspects of communication. Developing models that learn from such multi-modal inputs

efficiently is crucial to advance the generalization capabilities of state-of-the-art NLP models. Bisk et al. (2020) posit that text-only training seems to be reaching the point of diminishing returns and the next step in the development of NLP is leveraging multi-modal sources of information. Leveraging electroencephalography (EEG) and other physiological and behavioral signals seem especially appealing to model multi-modal human-like learning processes. Although combining different modalities or types of information for improving performance seems intuitively appealing, in practice, it is challenging to combine the varying level of noise and conflicts between modalities (Morency and Baltrušaitis, 2017). Therefore, we investigate if and how we take advantage of electrical brain activity signals to provide a human inductive bias for these natural language processing (NLP) models.

Two popular NLP tasks are sentiment analysis and relation detection. The goal of both tasks is to automatically extract information from text. Sentiment analysis is the task of identifying and categorizing subjective information in text. For example, the sentence “This movie is great fun.” contains a positive sentiment, while the sentence “This movie is terribly boring.” contains a negative sentiment. Relation detection is the task of identifying semantic relationships between entities in the text. In the sentence “Albert Einstein was born in Ulm.”, the relation *Birth Place* holds between the entities “Albert Einstein” and “Ulm”. NLP researchers have made great progress in building computational models for these tasks (Barnes et al., 2017; Rotsztein et al., 2018). However, these machine learning (ML) models still lack core human language understanding skills that humans perform effortlessly (Poria et al., 2020; Barnes et al., 2021). Barnes et al. (2019) find that sentiment models struggle with different linguistic elements such as negations or sentences containing mixed sentiment toward several target aspects.

1.1. Leveraging Physiological Data for Natural Language Processing

The increasing wealth of literature on the cognitive neuroscience of language (see reviews by Friederici, 2000; Poeppel et al., 2012; Poeppel, 2014) enables the use of cognitive signals in applied fields of language processing (e.g., Armeni et al., 2017). In recent years, natural language processing researchers have increasingly leveraged human language processing signals from physiological and neuroimaging recordings for both augmenting and evaluating machine learning-based NLP models (e.g., Hollenstein et al., 2019b; Toneva and Wehbe, 2019; Artemova et al., 2020). The approaches taken in those studies can be categorized as encoding or decoding cognitive processing signals. Encoding and decoding are complementary operations: encoding uses stimuli to predict brain activity, while decoding uses the brain activity to predict information about the stimuli (Naseleris et al., 2011). In the present study, we focus on the decoding process for predicting information about the text input from human brain activity.

Until now, mostly eye tracking and functional magnetic resonance imaging (fMRI) signals have been leveraged for this purpose (e.g., Fyshe et al., 2014). On the one hand, fMRI

recordings provide insights into the brain activity with a high spatial resolution, which furthers the research of localization of language-related cognitive processes. fMRI features are most often extracted over full sentences or longer text spans, since the extraction of word-level signals is highly complex due to the lower temporal resolution and hemodynamic delay. The number of cognitive processes and noise included in brain activity signals make feature engineering challenging. Machine learning studies leveraging brain activity data rely on standard preprocessing steps such as motion correction and spatial smoothing, and then use data-driven approaches to reduce the number of features, e.g., principal component analysis (Beinborn et al., 2019). Schwartz et al. (2019), for instance, fine-tuned a contextualized language model with brain activity data, which yields better predictions of brain activity and does not harm the model’s performance on downstream NLP tasks. On the other hand, eye tracking enables us to objectively and accurately record visual behavior with high temporal resolution at low cost. Eye tracking is widely used in psycholinguistic studies and it is common to extract well-established theory-driven features (Barrett et al., 2016; Hollenstein et al., 2020a; Mathias et al., 2020). These established metrics are derived from a large body of psycholinguistic research.

EEG is a non-invasive method to measure electrical brain surface activity. The synchronized activity of neurons in the brain produces electrical currents. The resulting voltage fluctuations can be recorded with external electrodes on the scalp. Compared to fMRI and other neuroimaging techniques, EEG can be recorded with a very high temporal resolution. This allows for more fine-grained language understanding experiments on the word-level, which is crucial for applications in NLP (Beres, 2017). To isolate certain cognitive functions, EEG signals can be split into frequency bands. For instance, effects related to semantic violations can be found within the gamma frequency range (~ 30 – 100 Hz), with well-formed sentences showing higher gamma levels than sentences containing violations (Penolazzi et al., 2009). Due to the wide extent of cognitive processes and the low signal-to-noise ratio in the EEG data, it is very challenging to isolate specific cognitive processes, so that more and more researchers are relying on machine learning techniques to decode the EEG signals (Affolter et al., 2020; Pfeiffer et al., 2020; Sun et al., 2020). These challenges are the decisive factors why EEG has not yet been used for NLP tasks. Data-driven approaches combined with the possibility of naturalistic reading experiments are now bypassing these challenges.

Reading times of words in a sentence depend on the amount of information the words convey. This correlation can be observed in eye tracking data, but also in EEG data (Frank et al., 2015). Thus, eye tracking and EEG are complementary measures of cognitive load. Compared to eye tracking, EEG may be more cumbersome to record and requires more expertise. Nevertheless, while eye movements indirectly reflect the cognitive load of text processing, EEG contains more direct and comprehensive information about language processing in the human brain. As we show below, this is beneficial for the higher level semantic NLP tasks targeted in this work. For instance, word predictability and semantic similarity show distinct patterns

of brain activity during language comprehension (Frank and Willems, 2017; Ettinger, 2020). The word representations used by neural networks and brain activity observed via the process of subjects reading a story can be aligned (Wehbe et al., 2014). Moreover, EEG effects that reflect syntactical processes can also be found in computational models of grammar (Hale et al., 2018).

The co-registration of EEG and eye-tracking has become an important tool for studying the temporal dynamics of naturalistic reading (Dimigen et al., 2011; Hollenstein et al., 2018; Sato and Mizuhara, 2018). This methodology has been increasingly and successfully used to study EEG correlates in the time domain (i.e., event-related potentials, ERPs) of cognitive processing in free viewing situations such as reading (Degno et al., 2021). In this context, fixation-related potentials (FRPs), which are the evoked electrical responses time-locked to the onset of fixations, have been studied and have received broad interest by naturalistic imaging researchers for free viewing studies. In naturalistic reading paradigms, FRPs allow the study of the neural dynamics of how novel information from currently fixated text affects the ongoing language comprehension process.

As of yet, the related work relying on EEG signals for NLP is very limited. Sassenhagen and Fiebach (2020) find that word embeddings can successfully predict the pattern of neural activation. However, their experiment design does not include natural reading, but reading isolated words. Hollenstein et al. (2019b) similarly find that various embedding type are able to predict aggregate word-level activations from natural reading, where contextualized embeddings perform best. Moreover, Murphy and Poesio (2010) showed that semantic categories can be detected in simultaneous EEG recordings. Muttenthaler et al. (2020) used EEG signals to train an attention mechanism, similar to Barrett et al. (2018a), who used eye tracking signals to induce machine attention with human attention. However, EEG has not yet been leveraged for higher-level semantic tasks such as sentiment analysis or relation detection. Deep learning techniques have been applied to decode EEG signals (Craik et al., 2019), especially for brain-computer interface technologies, e.g., Nurse et al. (2016). However, this avenue has not yet been explored when leveraging EEG signals to enhance NLP models. Through decoding EEG signals occurring during language understanding, more specifically, during English sentence comprehension, we aim to explore their impact on computational language understanding tasks.

1.2. Contributions

More than a practical application of improving real-world NLP tasks, our main goal is to explore to what extent there is additional linguistic processing information in the EEG signal to complement the text input. In this present study, we investigate for the first time the potential of leveraging EEG signals for augmenting NLP models. For the purpose of making language decoding studies from brain activity more interpretable, we follow the recommendations of Gauthier and Ivanova (2018): (1) We commit to a specific mechanism and task, and (2) subdivide the input feature space including theoretically founded preprocessing steps. We investigate the impact of enhancing a neural network architecture for two common NLP tasks with

a range of EEG features. We propose a multi-modal network capable of processing textual features and brain activity features simultaneously. We employ two different well-established types of neural network architectures for decoding the EEG signals throughout the entire study. To analyze the impact of different EEG features, we perform experiments on sentiment analysis as a binary or ternary sentence classification task, and relation detection as a multi-class and multi-label classification task. We investigate the effect of augmenting NLP models with neurophysiological data in an extensive study while accounting for various dimensions:

1. We present a comparison of a purely data-driven approach of feature extraction for machine learning, using full *broadband EEG signals*, to a more theoretically motivated approach, splitting the word-level EEG features into *frequency bands*.
2. We develop two *EEG decoding components* for our multi-modal ML architecture: A recurrent and a convolutional component.
3. We contrast the effects of these EEG features on multiple *word representation types* commonly used in NLP. We compare the improvements of EEG features as a function of various *training data sizes*.
4. We analyze the impact of the EEG features on varying *classification complexity*: from binary classification to multi-class and multi-label tasks.

This comprehensive study is completed by comparing the impact of the decoded EEG signals not only to a *text-only* baseline, but also to baselines augmented with eye tracking data as well as random noise. In the next section, we describe the materials used in this study and the multi-modal machine learning architecture. Thereafter, we present the results of the NLP tasks and discuss the dimensions defined above.

2. MATERIALS AND METHODS

2.1. Data

For the purpose of augmenting natural language processing tasks with brain activity signals, we leverage the Zurich Cognitive Language Processing Corpus (ZuCo; Hollenstein et al., 2018, 2020b). ZuCo is an openly available dataset of simultaneous EEG and eye tracking data from subjects reading naturally occurring English sentences. This corpus consists of two datasets, ZuCo 1.0 and ZuCo 2.0, which contain the same type of recordings. We select the normal reading paradigms from both datasets, in which participants were instructed to read English sentences in their own pace with no specific task beyond reading comprehension. The participants read one sentence at a time, using a control pad to move to the next sentence. This setup facilitated the naturalistic reading paradigm. Descriptive statistics about the datasets used in this work are presented in **Table 1**.

A detailed description of the entire ZuCo dataset, including individual reading speed, lexical performance, average word length, average number of words per sentence, skipping proportion on word level, and effect of word length on skipping proportion, can be found in Hollenstein et al. (2018). In the

TABLE 1 | Details about the ZuCo tasks used in this paper.

	ZuCo 1.0 Task SR	ZuCo 1.0 Task NR	ZuCo 2.0 Task NR
Participants	12	12	18
Sentences	400	300	349
Words	7,079	6,386	6,828
Unique word types	3,080	2,657	2,412
Sentiment analysis	✓	-	-
Relation detection	-	✓	✓

In Task SR participants read sentences from movie reviews, and in Task NR sentences from Wikipedia articles.

TABLE 2 | Example sentences for all three NLP tasks used in this study.

Task	Example sentence	Label(s)
Binary/ternary sentiment analysis	"The film often achieves a mesmerizing poetry."	<i>Positive</i>
Binary/ternary sentiment analysis	"Flaccid drama and exasperatingly slow journey."	<i>Negative</i>
Ternary sentiment analysis	"A portrait of an artist."	<i>Neutral</i>
Relation detection	"He attended Wake Forest University."	<i>Education</i>
Relation detection	"She attended Beverly Hills High School, but left to become an actress."	<i>Education, Job Title</i>

following section, we will describe the methods relevant to the subset of the ZuCo data used in the present study.

2.1.1. Participants

For ZuCo 1.0, data were recorded from 12 healthy adults (between 22 and 54 years old; all right-handed; 5 female subjects). For ZuCo 2.0, data were recorded from 18 healthy adults (between 23 and 52 years old; 2 left-handed; 10 female subjects). The native language of all participants is English, originating from Australia, Canada, UK, USA or South Africa. In addition, all subjects completed the standardized LexTALE test to assess their vocabulary and language proficiency (Lexical Test for Advanced Learners of English; Lemhöfer and Broersma, 2012). All participants gave written consent for their participation and the re-use of the data prior to the start of the experiments. The study was approved by the Ethics Commission of the University of Zurich.

2.1.2. Reading Materials and Experimental Design

The reading materials recorded for the ZuCo corpus contain sentences from movie reviews from the Stanford Sentiment Treebank (Socher et al., 2013) and Wikipedia articles from a dataset provided by Culotta et al. (2006). These resources were chosen since they provide ground truth labels for the machine learning tasks in this work. **Table 2** presents a few examples of the sentences read during the experiments.

For the recording sessions, the sentences were presented one at a time at the same position on the screen. Text was presented in black with font size 20-point Arial on a light gray background resulting in a letter height of 0.8 mm or 0.674° . The lines were triple-spaced, and the words double-spaced. A maximum of 80 letters or 13 words were presented per line in all three tasks. Long sentences spanned multiple lines. A maximum of 7 lines for Task 1, 5 lines for Task 2 and 7 lines for Task 3 were presented simultaneously on the screen.

During the normal reading tasks included in the ZuCo corpus, the participants were instructed to read the sentences naturally, without any specific task other than comprehension. Participants were told to read the sentences normally without any special instructions. Participants were equipped with a control to trigger the onset of the next sentence. The task was explained to the subjects orally, followed by instructions on the screen.

The control condition for this task consisted of single-choice reading comprehension questions about the content of the previous sentence. Twelve percent of randomly selected sentences were followed by a control question on a separate screen. To test the participants' reading comprehension, these questions ask about the content of the previous sentence. The questions are presented with three answer options, out of which only one is correct.

2.1.3. EEG Data

In this section, we present the EEG data extracted from the ZuCo corpus for this work. We describe the acquisition and preprocessing procedures as well as the feature extraction.

2.1.3.1. EEG Acquisition and Preprocessing

High-density EEG data were recorded using a 128-channel EEG Geodesic Hydrocel system (Electrical Geodesics, Eugene, Oregon) with a sampling rate of 500 Hz. The recording reference was at Cz (vertex of the head), and the impedances were kept below 40 k μ . All analyses were performed using MATLAB 2018b (The MathWorks, Inc., Natick, Massachusetts, United States). EEG data was automatically preprocessed using the current version (2.4.3) of *Automagic* (Pedroni et al., 2019). *Automagic* is an open-source MATLAB toolbox that acts as a wrapper to run currently available EEG preprocessing methods and offers objective standardized quality assessment for large studies. The code for the preprocessing can be found online.¹

Our preprocessing pipeline consisted of the following steps. First, 13 of the 128 electrodes in the outermost circumference (chin and neck) were excluded from further processing as they capture little brain activity and mainly record muscular activity. Additionally, 10 EOG electrodes were used for blink and eye movement detection (and subsequent rejection) during ICA. The EOG electrodes were removed from the data after the preprocessing, yielding a total number of 105 EEG electrodes. Subsequently, *bad* channels were detected by the algorithms implemented in the EEGLAB plugin `clean_rawdata`², which removes flatline, low-frequency, and noisy channels. A channel

¹<https://github.com/methlabUZH/automagic>

²http://scn.ucsd.edu/wiki/Plugin_list_process

was defined as a bad electrode when recorded data from that electrode was correlated at less than 0.85 to an estimate based on neighboring channels. Furthermore, a channel was defined as bad if it had more line noise relative to its signal than all other channels (4 standard deviations). Finally, if a channel had a longer flat-line than 5 s, it was considered bad. These bad channels were automatically removed and later interpolated using a spherical spline interpolation (EEGLAB function `eeg_interp.m`). The interpolation was performed as a final step before the automatic quality assessment of the EEG files (see below). Next, data was filtered using a high-pass filter (−6dB cut-off: 0.5 Hz) and a 60 Hz notch filter was applied to remove line noise artifacts. Thereafter, an independent component analysis (ICA) was performed. Components reflecting artifactual activity were classified by the pre-trained classifier MARA (Winkler et al., 2011). MARA is a supervised machine learning algorithm that learns from expert ratings. Therefore, MARA is not limited to a specific type of artifact, and should be able to handle eye artifacts, muscular artifacts and loose electrodes equally well. Each component being classified with a probability rating >0.5 for any class of artifacts was removed from the data. Finally, residual bad channels were excluded if their standard deviation exceeded a threshold of 25 μV . After this, the pipeline automatically assessed the quality of the resulting EEG files based on four criteria: A data file was marked as bad-quality EEG and not included in the analysis if (1) the proportion of high-amplitude data points in the signals ($>30 \mu\text{V}$) was larger than 0.20; (2) more than 20% of time points showed a variance larger than 15 microvolts across channels; (3) 30% of the channels showed high variance ($>15 \mu\text{V}$); (4) the ratio of bad channels was higher than 0.3.

Free viewing in reading is an important characteristic of naturalistic behavior and imposes challenges for the analysis of electrical brain activity data. Free viewing in the context of our study refers to the participant's ability to perform self-paced reading given the experimental requirement to keep the head still during data recording. In the case of EEG recordings during naturalistic reading, the self-paced timing of eye fixations leads to a temporal overlap between successive fixation-related events (Dimigen et al., 2011). In order to isolate the signals of interest and correct for temporal overlap in the continuous EEG, several methods using linear-regression-based deconvolution modeling have been proposed for estimating the overlap-corrected underlying neural responses to events of different types (e.g., Ehinger and Dimigen 2019; Smith and Kutas 2015a,b). Here, we used the *unfold* toolbox for MATLAB (Ehinger and Dimigen, 2019).³ Deconvolution modeling is based on the assumption that in each channel the recorded signal consists of a combination of time-varying and partially overlapping event-related responses and random noise. Thus, the model estimates the latent event-related responses to each type of event based on repeated occurrences of the event over time, in our case eye fixations.

³<https://github.com/unfoldtoolbox/unfold/>

2.1.3.2. EEG Features

The fact that ZuCo provides simultaneous EEG and eye tracking data highly facilitates the extraction of word-level brain activity signals. Dimigen et al. (2011) demonstrated that EEG indices of semantic processing can be obtained in natural reading and compared to eye movement behavior. The eye tracking data provides millisecond-accurate fixation times for each word. Therefore, we were able to obtain the brain activity during all fixations of a word by computing fixation-related potentials aligned to the onsets of the fixation on a given word.

In this work, we select a range of EEG features with a varying degree of theory-driven and data-driven feature extraction. We define the *broadband EEG signal*, i.e., the full EEG signal from 0.5 to 50 Hz as the averaged brain activity over all fixations of a word, i.e., its total reading time. We compare the full EEG features, a data-driven feature extraction approach, to *frequency band features*, a more theoretically motivated approach. Different neurocognitive aspects of language processing during reading are associated with brain oscillations at various frequencies. These frequency ranges are known to be associated with certain cognitive functions. We split the EEG signal into four frequency bands to limit the bandwidth of the EEG signals to be analyzed. The frequency bands are fixed ranges of wave frequencies and amplitudes over a time scale: *theta* (4–8 Hz), *alpha* (8.5–13 Hz), *beta* (13.5–30 Hz), and *gamma* (30.5–49.5 Hz). We elaborate on cognitive and linguistic functions of each of these frequency bands in section 4.1.

We then applied a Hilbert transform to each of these time-series, resulting in a complex time-series. The Hilbert phase and amplitude estimation method yields results equivalent to sliding window FFT and wavelet approaches (Bruns, 2004). We specifically chose the Hilbert transformation to maintain temporal information for the amplitude of the frequency bands to enable the power of the different frequencies for time segments defined through fixations from the eye-tracking recording. Thus, for each eye-tracking feature we computed the corresponding EEG feature in each frequency band. For each EEG eye-tracking feature, all channels were subject to an artifact rejection criterion of $\pm 90 \mu\text{V}$ to exclude trials with transient noise.

In spite of the high inter-subject variability in EEG data, it has been shown in previous research of machine learning applications (Foster et al., 2018; Hollenstein et al., 2019a), that averaging over the EEG features of all subjects yields results almost as good as the single best-performing subjects. Hence, we also average the EEG features over all subjects to obtain more robust features. Finally, for each word in each sentence, the EEG features consist of a vector of 105 dimensions (one value for each EEG channel). For training the ML models, we split all available sentences into sets of 80% for training and 20% for testing to ensure that the test data is unseen during training.

2.1.4. Eye Tracking Data

In the following, we describe the eye tracking data recorded for the Zurich Cognitive Language Processing Corpus. In this study, we focus on decoding EEG data, but we use eye movement data to compute an additional baseline. As mentioned previously, augmenting ML models with eye tracking yields consistent

improvements across a range of NLP tasks, including sentiment analysis and relation extraction (Long et al., 2017; Mishra et al., 2017; Hollenstein et al., 2019a). Since the ZuCo datasets provide simultaneous EEG and eye tracking recordings, we leverage the available eye tracking data to augment all NLP tasks with eye tracking features as an additional multi-modal baseline based on cognitive processing features.

2.1.4.1. Eye Tracking Acquisition and Preprocessing

Eye movements were recorded with an infrared video-based eye tracker (EyeLink 1000 Plus, SR Research) at a sampling rate of 500 Hz. The EyeLink 1000 tracker processes eye position data, identifying saccades, fixations and blinks. Fixations were defined as time periods without saccades during which the visual gaze is fixed on a single location. The data therefore consists of (x,y) gaze location entries for individual fixations mapped to word boundaries. A fixation lasts around 200–250ms (with large variations). Fixations shorter than 100 ms were excluded, since these are unlikely to reflect language processing (Serenó and Rayner, 2003). Fixation duration depends on various linguistic effects, such as word frequency, word familiarity and syntactic category (Clifton et al., 2007).

2.1.4.2. Eye Tracking Features

The following features were extracted from the raw data: (1) *gaze duration* (GD), the sum of all fixations on the current word in the first-pass reading before the eye moves out of the word; (2) *total reading time* (TRT), the sum of all fixation durations on the current word, including regressions; (3) *first fixation duration* (FFD), the duration of the first fixation on the prevailing word; (4) *go-past time* (GPT), the sum of all fixations prior to progressing to the right of the current word, including regressions to previous words that originated from the current word; (5) number of fixations (nFix), the total amount of fixations on the current word.

We use these five features provided in the ZuCo dataset, which cover the extent of the human reading process. To increase the robustness of the signal, analogously to the EEG features, the eye tracking features are averaged over all subjects (Barrett and Hollenstein, 2020). This results in a feature vector of five dimensions for each word in a sentence. Training and test data were split in the same fashion as the EEG data.

2.2. Natural Language Processing Tasks

In this section, we describe the natural language processing tasks we use to evaluate the multi-modal ML models. As usual in supervised machine learning, the goal is to learn a mapping from given input features to an output space to predict the labels as accurately as possible. The tasks we consider in our work do not differ much in the input definition as they consist of three sequence classification tasks for information extraction from text. The goal of a sequence classification task is to assign the correct label(s) to a given sentence. The input for all tasks consists of tokenized sentences, which we augment with additional features, i.e., EEG or eye tracking. The labels to predict vary across the three chosen tasks resulting in varying task difficulty. **Table 2** provides examples for all three tasks.

2.2.1. Task 1 and 2: Sentiment Analysis

The objective of sentiment analysis is to interpret subjective information in text. More specifically, we define sentiment analysis as a sentence-level classification task. We run our experiments on both binary (*positive/negative*) and ternary (+ *neutral*) sentiment classification. For this task, we leverage only the sentences recorded in the first task of ZuCo 1.0, since they are part of the Stanford Sentiment Treebank (Socher et al., 2013), and thus directly provide annotated sentiment labels for training the ML models. For the first task, binary sentiment analysis, we use the 263 positive and negative sentences. For the second task, ternary sentiment analysis, we additionally use the neutral sentences, resulting in a total of 400 sentences.

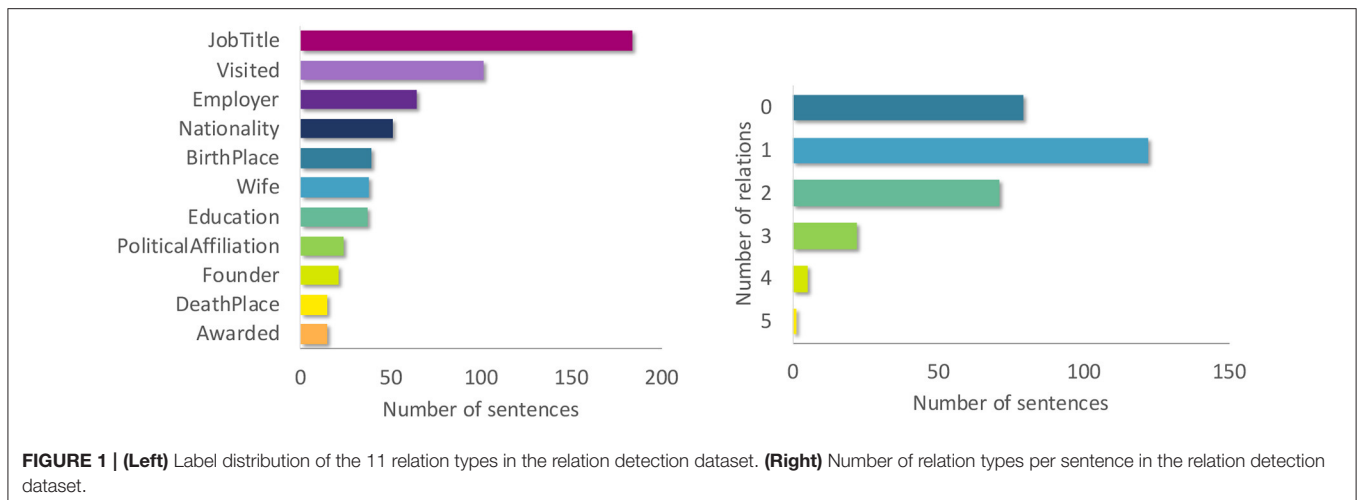
2.2.2. Task 3: Relation Detection

Relation classification is the task of identifying the semantic relation holding between two entities in text. The ZuCo corpus also contains Wikipedia sentences with relation types such as *Job Title*, *Nationality*, and *Political Affiliation*. The sentences in ZuCo 1.0 and ZuCo 2.0, from the normal reading experiment paradigms, include 11 relation types (**Figure 1**). In order to further increase the task complexity, we treat this task differently than the sentiment analysis tasks. Since any sentence can include zero, one or more of the relevant semantic relations (see example in **Table 2**, we treat relation detection as a multi-class *and* multi-label sequence classification task. Concretely, every sample can be assigned to any possible combination out of the 11 classes including none of them. Removing duplicates between ZuCo 1.0 and ZuCo 2.0 resulted in 594 sentences used for training the models. **Figure 1** illustrates the label and relation distribution of the sentences used to train the relation detection task.

2.3. Multi-Modal Machine Learning Architecture

We present a multi-modal neural architecture to augment the NLP sequence classification tasks with any other type of data. Although combining different modalities or types of information for improving performance seems an intuitively appealing task, it is often challenging to combine the varying levels of noise and conflicts between modalities in practice.

Previous works using physiological data for improving NLP tasks mostly implement *early fusion* multi-modal methods, i.e., directly concatenating the textual and cognitive embeddings before inputting them into the network. For example, Hollenstein and Zhang (2019), Barrett et al. (2018b) and Mishra et al. (2017) concatenate textual input features with eye-tracking features to improve NLP tasks such as entity recognition, part-of-speech tagging and sentiment analysis, respectively. Concatenating the input features at the beginning in only one joint decoder component aims at learning a joint decoder across all modalities at risk of implicitly learning different weights for each modality. However, recent multi-modal machine learning work has shown the benefits of *late fusion* mechanisms (Ramachandram and Taylor, 2017). Do et al. (2017) argument in favor of concatenating the hidden layers instead of concatenating the features at input time. Such multi-modal models have been successfully applied in other areas, mostly combining inputs across different domains,



for instance, learning speech reconstruction from silent videos (Ephrat et al., 2017), or for text classification using images (Kiela et al., 2018). Tsai et al. (2019) train a multi-modal sentiment analysis model from natural language, facial gestures, and acoustic behaviors.

Hence, we adopted the late fusion strategy in our work. We present multi-modal models for various NLP tasks, combining the learned representations of all input types (i.e., text and EEG features) in a late fusion mechanism before conducting the final classification. Purposefully, this enables the model to learn independent decoders for each modality before fusing the hidden representations together. In the present study, we investigate the proposed multi-modal machine learning architecture, which learns simultaneously from text and from cognitive data such as eye tracking and EEG signals.

In the following, we first describe the uni-modal and multi-modal baseline models we use to evaluate the results. Thereafter, we present the multi-modal NLP models that jointly learn from text and brain activity data.

2.3.1. Uni-Modal Text Baselines

For each of the tasks presented above, we train uni-modal models on textual features only. To represent the word numerically, we use word embeddings. Word embeddings are vector representations of words, computed so that words with similar meaning have a similar representation. To analyze the interplay between various types of word embeddings and EEG data, we use the following three embedding types typically used in practice: (1) randomly initialized embeddings trained at run time on the sentences provided, (2) GloVe pre-trained embeddings based on word co-occurrence statistics (Pennington et al., 2014)⁴, and (3) BERT pre-trained contextual embeddings (Devlin et al., 2019).⁵

The randomly initialized word representations define word embeddings as n -by- d matrices, where n is the vocabulary size, i.e., the number of unique words in our dataset, and d is the

embedding dimension. Each value in that matrix is randomly initialized and will then be trained together with the neural network parameters. We set $d = 32$. This type of embeddings does not benefit from pre-training on large text collections and hence is known to perform worse than GloVe or BERT embeddings. We include them in our study to better isolate the impact of the EEG features and to limit the learning of the model on the text it is trained on. Non-contextual word embeddings such as GloVe encode each word in a fixed vocabulary as a vector. The purpose of these vectors is to encode semantic information about a word, such that similar words result in similar embedding vectors. We use the GloVe embeddings of $d = 300$ dimensions that are trained on 6 billion words. The contextualized BERT embeddings were pre-trained on multiple layers of transformer models with self-attention (Vaswani et al., 2017). Given a sentence, BERT encodes each word into a feature vector of dimension $d = 768$, which incorporates information from the word's context in the sentence.

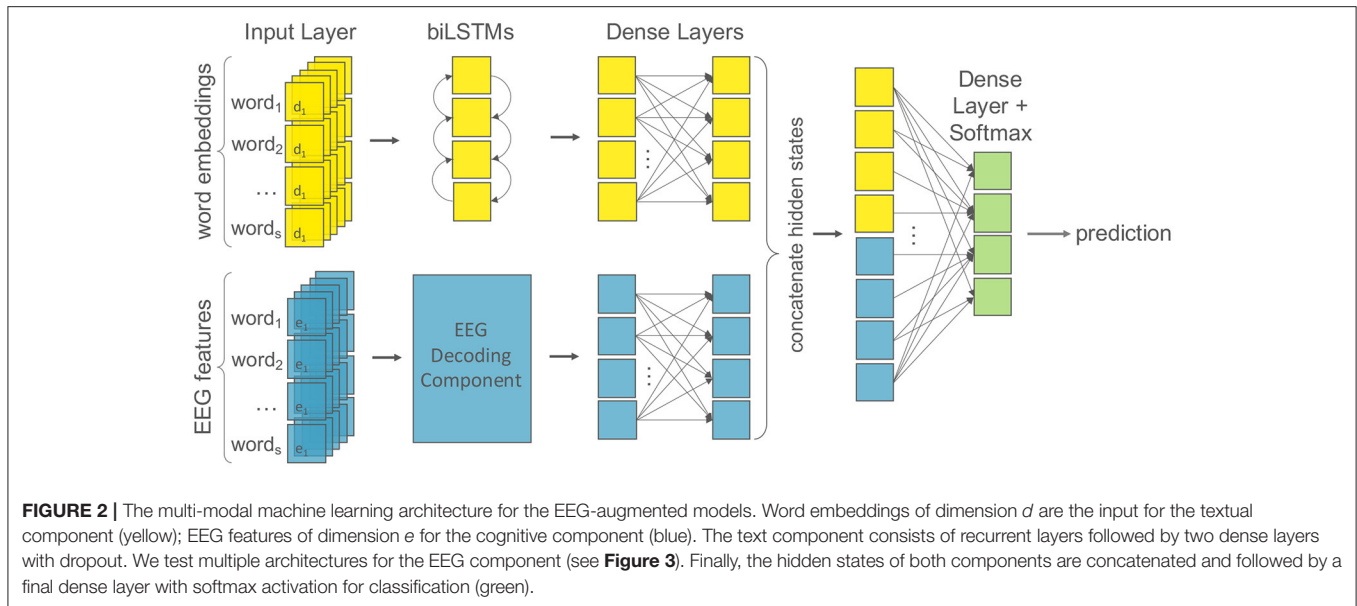
The uni-modal text baseline model consists of a first layer taking the embeddings as an input, followed by a bidirectional Long-Short Term Memory network (LSTM; Hochreiter and Schmidhuber, 1997), then two fully-connected dense layers with dropout between them, and finally a prediction layer using softmax activation. This corresponds to a single component of the multi-modal architecture, i.e., the top component in **Figure 2**. Following best practices (e.g., Sun et al., 2019), we set the weights of BERT to be trainable similarly to the randomly initialized embeddings. This process of adjusting the initialized weights of a pre-trained feature extractor during the training process, in our case BERT, is commonly known as *fine-tuning* in the literature (Howard and Ruder, 2018). In contrast, the parameters of the GloVe embeddings are fixed to the pre-trained weights and thus do not change during training.

2.3.2. Multi-Modal Baselines

To analyze the effectiveness of our multi-modal architecture with EEG signals properly, we not only compare it to uni-modal text baselines, but also to multi-modal baselines using

⁴<https://nlp.stanford.edu/projects/glove/>

⁵<https://huggingface.co/bert-base-uncased>



the same architecture described in the next section for the EEG models, but replacing the features of the second modality with the following alternatives: (1) We implement a gaze-augmented baseline, where the five eye tracking features described in section 2.1.4.2 are combined with the word embeddings by adding them to the multi-modal model in the same manner as the EEG features, as vectors with dimension = 5. The purpose of this baseline is to allow a comparison of multi-modal models learning from two different types of physiological features. Since the benefits of eye tracking data in ML models are well-established (Barrett and Hollenstein, 2020; Mathias et al., 2020), this is a strong baseline. (2) We further implement a random noise-augmented baseline, where we add uniformly sampled vectors of random numbers as the second input data type to the multi-modal model. These random vectors are of the same dimension as the EEG vectors (i.e., $d = 105$). It is well-known that the addition of noise to the input data of a neural network during training can lead to improvements in generalization performance as a form of regularization (Bishop, 1995). Thus, this baseline is relevant because we want to analyze whether the improvements from the EEG signals on the NLP tasks are due to its capability of extracting linguistic information and not merely due to additional noise.

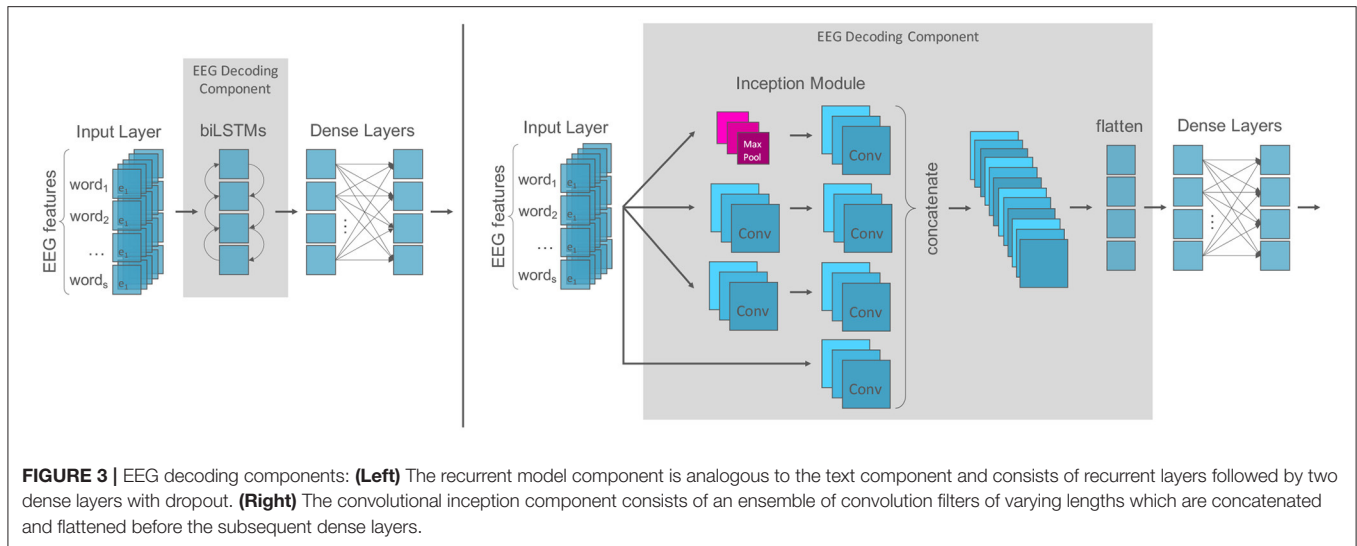
2.3.3. EEG Models

To fully understand the impact of the EEG data on the NLP models, we build a model that is able to deal with multiple inputs and mixed data. We present a multi-modal model with late decision-level fusion to learn joint representations of textual and cognitive input features. We test both a recurrent and a convolutional neural architecture for decoding the EEG signals. **Figure 2** depicts the main structure of our model and we describe the individual components below.

All input sentences are padded to the maximum sentence length to provide fixed-length text inputs to the model. Word

embeddings of dimension d are the input for the textual component, where $d \in \{32, 300, 768\}$ for randomly initialized embeddings, GloVe embeddings and BERT embeddings, respectively. EEG features of dimension e are the input for the cognitive component, where $e = 105$. As described, the text component consists of bidirectional LSTM layers followed by two dense layers with dropout. Text and EEG features are given as independent inputs to their own respective component of the network. The hidden representations of these are then concatenated before being fed to a final dense classification layer. We also experimented with different merging mechanisms to join the text and EEG layers of our two-tower model (concatenation, addition, subtraction, maximum). Concatenation overall achieved the best results, so we report only these. Although the goal of each network is to learn feature transformations for their own modality, the relevant extracted information should be complementary. This is achieved, as commonly done in deep learning, through alternatively running inference and back-propagation of the data through the entire network enabling information to *flow* from the component responsible for one input modality to the other via the fully connected output layers. To learn a non-linear transformation function for each component, we employ the rectified linear units (ReLU) as activation functions after each hidden layer.

For the EEG component, we test a recurrent and a convolutional architecture since both have proven useful in learning features from time series data for language processing (e.g., Lipton et al., 2015; Yin et al., 2017; Fawaz et al., 2020). For the recurrent architecture (**Figure 3**, left), the model component is analogous to the text component: it consists of bidirectional LSTM layers followed by two dense layers with dropout and ReLU activation functions. For the convolutional architecture (**Figure 3**, right), we build a model component based on the *Inception* module first introduced by Szegedy et al. (2015). An *inception* module is an ensemble of convolutions that applies



multiple filters of varying lengths simultaneously to an input time series. This allows the network to automatically extract relevant features from both long and short time series. As suggested by Schirrmeyer et al. (2017) we used exponential linear unit activations (ELUs; Clevert et al., 2015) in the convolutional EEG decoding model component.

For binary and ternary sentiment analysis, the final dense layer has a softmax activation in order to use the maximal output for the classification. For the multi-label classification case of relation detection, we replace the softmax function in the last dense layer of the model with a sigmoid activation to produce independent scores for each class. If the score for any class surpasses a certain threshold, the sentence is labeled to contain that relation type (opposite to simply taking the *max* score as the label of the sentence). The threshold is tuned as an additional hyper-parameter.

This multi-modal model with separate components learned for each input data type has several advantages: It allows for separate pre-processing of each type of data. For instance, it is able to deal with differing tokenization strategies, which is useful in our case since it is challenging to map linguistic tokenization to the word boundaries presented to participants during the recordings of eye tracking and brain activity. Moreover, this approach is scalable to any number of input types. The generalizability of our model enables the integration of multiple data representations, e.g., learning from brain activity, eye movements, and other cognitive modalities simultaneously.

2.3.4. Training Setup

To assess the impact of the EEG signals under fair modeling conditions, the hyper-parameters are tuned individually for all baseline models as well as for all eye tracking and EEG augmented models. The ranges of the hyper-parameters are presented in **Table 3**. All results are reported as means over five independent runs with different random seeds. In each run, five-fold cross-validation is performed on a 80% training and 20% test split. The best parameters were selected according to the model's accuracy on the validation set (10% of the training set) across

TABLE 3 | Tested value ranges included in the hyper-parameter search for our multi-modal machine learning architecture.

Parameter	Range
LSTM layer dimension	64, 128, 256, 512
Number of LSTM layers	1, 2, 3, 4
CNN filters	14, 16, 18
CNN kernel sizes	[1,4,7]
CNN pool sizes	3, 5, 7
Dense layer dimension	8, 16, 32, 64, 128, 256, 512
Dropout	0.1, 0.3, 0.5
Batch size	20, 40, 60
Learning rate	10^{-1} , 10^{-2} , 10^{-3} , 10^{-4} , 10^{-5}
Random seeds	13, 22, 42, 66, 78
Threshold	0.3, 0.5, 0.7

Threshold only applies to relation detection.

all five-folds. We implemented early stopping with a patience of 80 epochs and a minimum difference in validation accuracy of 10^{-7} . The validation set is used for both parameter tuning and early stopping.

3. RESULTS

In this study, we assess the potential of EEG brain activity data to enhance NLP tasks in a multi-modal architecture. We present the results of all augmented models compared to the baseline results. As described above, we select the hyper-parameters based on the best validation accuracy achieved for each setting.

The performance of our models is evaluated based on the comparison between the predicted labels (i.e., positive, neutral or negative sentiment for a sentence; or the relation type(s) in a sentence) and the true labels of the test set resulting in the number of *true positives* (TP), *true negatives* (TN), *false positives* (FP), and *false negatives* (FN) across the classified samples. The

TABLE 4 | Binary sentiment analysis results of the multi-modal model using the recurrent EEG decoding component.

Model	Randomly initialized			GloVe			BERT		
	P	R	F ₁ (std)	P	R	F ₁ (std)	P	R	F ₁ (std)
Baseline	0.572	0.573	0.552 (0.07)	0.751	0.738	0.728 (0.08)	0.900	0.899	0.893 (0.04)
+ noise	0.599	0.574	0.541 (0.08)	0.721	0.715	0.709 (0.09)	0.914	0.916	0.913 (0.03)
+ ET	0.615	0.605	0.586 (0.06)	0.795	0.786	0.781 (0.06)	0.913	0.907	0.904 (0.05)
+ EEG full	0.540	0.538	0.525 (0.06)	0.738	0.729	0.725 (0.07)	0.913	0.909	0.906 (0.04)
+ EEG θ	0.602	0.599	0.584* (0.08)	0.789	0.785	0.783+ (0.05)	0.917	0.916	0.913* (0.04)
+ EEG α	0.610	0.590	0.565 (0.05)	0.763	0.758	0.753 (0.05)	0.912	0.908	0.906 (0.03)
+ EEG β	0.587	0.578	0.555 (0.07)	0.781	0.777	0.774+ (0.06)	0.911	0.911	0.907* (0.04)
+ EEG γ	0.614	0.591	0.553 (0.08)	0.777	0.773	0.769* (0.07)	0.917	0.917	0.915* (0.04)
+ θ + α + β + γ	0.597	0.597	0.569 (0.08)	0.766	0.764	0.760* (0.07)	0.913	0.913	0.911* (0.04)

We report precision (P), recall (R), F₁-score and the standard deviation (std) between five runs. The best results per column are marked in bold. Significance is indicated on the F₁-score with asterisks: * denotes p<0.05 (uncorrected), + denotes p<0.003 (Bonferroni corrected p-value).

TABLE 5 | Ternary sentiment analysis results of the multi-modal model using the recurrent EEG decoding component.

Model	Randomly initialized			GloVe			BERT		
	P	R	F ₁ (std)	P	R	F ₁ (std)	P	R	F ₁ (std)
Baseline	0.408	0.384	0.351 (0.07)	0.510	0.507	0.496 (0.06)	0.722	0.714	0.710 (0.05)
+ noise	0.373	0.399	0.344 (0.10)	0.531	0.519	0.504 (0.04)	0.711	0.706	0.700 (0.06)
+ ET	0.424	0.413	0.388 (0.06)	0.539	0.528	0.513 (0.04)	0.728	0.717	0.714 (0.05)
+ EEG full	0.391	0.387	0.353 (0.07)	0.505	0.505	0.488 (0.07)	0.724	0.715	0.711 (0.06)
+ EEG θ	0.397	0.409	0.360 (0.07)	0.516	0.510	0.498 (0.06)	0.715	0.708	0.704 (0.05)
+ EEG α	0.390	0.390	0.347 (0.08)	0.520	0.516	0.506 (0.05)	0.720	0.712	0.707 (0.05)
+ EEG β	0.350	0.370	0.302 (0.09)	0.523	0.519	0.509 (0.05)	0.732	0.720	0.717 (0.07)
+ EEG γ	0.409	0.397	0.359 (0.07)	0.517	0.513	0.502 (0.04)	0.709	0.705	0.697 (0.06)
+ θ + α + β + γ	0.401	0.400	0.368 (0.06)	0.522	0.516	0.505 (0.05)	0.722	0.717	0.713 (0.05)

We report precision (P), recall (R), F₁-score and the standard deviation (std) between five runs. The best results per column are marked in bold.

terms *positive* and *negative* refer to the classifier's prediction, and the terms *true* and *false* refer to whether that prediction corresponds to the ground truth label. The following decoding performance metrics were computed:

Precision is the fraction of relevant instances among the retrieved instances, and is defined as

$$\text{Precision} = \frac{TP}{TP + FP} \quad (1)$$

Recall is the fraction of the relevant instances that are successfully retrieved:

$$\text{Recall} = \frac{TP}{TP + FN} \quad (2)$$

The F₁-score is the harmonic mean combining precision and recall:

$$F_1 \text{ score} = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}} \quad (3)$$

For analyzing the results, we report macro-averaged precision (P), recall (R), and F₁-score, i.e., the metrics are calculated for each label to counteract the label imbalance in the datasets.

The results for the multi-modal architecture using the recurrent EEG decoding component are presented in **Table 4** for binary sentiment analysis, **Table 5** for ternary sentiment analysis, and **Table 6** for relation detection. The first three rows in each table represent the uni-modal text baseline, the multi-modal noise and eye-tracking baselines. This is followed by the multi-modal models augmented with the full broadband EEG signals and each of the four frequency bands. Finally, in the last row, we also present the results of a multi-modal model with five components, where text and each frequency band extractors are learned separately and concatenated at the end. In both sentiment tasks, the EEG data yields modest but consistent improvements over the text baseline for all word embeddings types. However, in the case of relation detection, the addition of either eye tracking or brain activity data is not helpful for randomly initialized embeddings and only beneficial in some settings using GloVe embeddings. Nevertheless, the

TABLE 6 | Relation detection results of the multi-modal model using the recurrent EEG decoding component.

Model	Randomly initialized			GloVe			BERT		
	P	R	F ₁ (std)	P	R	F ₁ (std)	P	R	F ₁ (std)
Baseline	0.404	0.525	0.452 (0.04)	0.501	0.609	0.539 (0.05)	0.522	0.788	0.623 (0.05)
+ noise	0.420	0.424	0.408 (0.07)	0.577	0.497	0.532 (0.03)	0.675	0.585	0.625 (0.03)
+ ET	0.421	0.404	0.402 (0.06)	0.547	0.476	0.506 (0.04)	0.661	0.631	0.644 (0.03)
+ EEG full	0.345	0.343	0.334 (0.05)	0.511	0.387	0.432 (0.09)	0.652	0.690	0.668* (0.10)
+ EEG θ	0.430	0.421	0.414 (0.07)	0.582	0.508	0.539 (0.07)	0.646	0.736	0.684* (0.08)
+ EEG α	0.368	0.373	0.358 (0.12)	0.582	0.515	0.542 (0.06)	0.652	0.715	0.679* (0.07)
+ EEG β	0.349	0.340	0.329 (0.09)	0.581	0.497	0.532 (0.10)	0.674	0.726	0.696+ (0.06)
+ EEG γ	0.410	0.399	0.397 (0.05)	0.554	0.488	0.514 (0.09)	0.666.	0.715	0.686* (0.07)
+ $\theta + \alpha + \beta + \gamma$	0.370	0.376	0.363 (0.09)	0.554	0.488	0.514 (0.09)	0.675	0.646	0.659 (0.04)

We report precision (P), recall (R), F₁-score and the standard deviation (std) between five runs. The best results per column are marked in bold. Significance is indicated on the F₁-score with asterisks: * denotes $p < 0.05$ (uncorrected), + denotes $p < 0.003$ (Bonferroni corrected p -value).

TABLE 7 | Binary sentiment analysis results of the multi-modal model using the convolutional EEG decoding component.

Model	Randomly initialized			GloVe			BERT		
	P	R	F ₁ (std)	P	R	F ₁ (std)	P	R	F ₁ (std)
Baseline	0.572	0.573	0.552 (0.07)	0.751	0.738	0.728 (0.08)	0.900	0.899	0.893 (0.04)
+ noise	0.558	0.584	0.528 (0.11)	0.780	0.767	0.762 (0.06)	0.895	0.887	0.883 (0.05)
+ ET	0.617	0.623	0.610 (0.07)	0.790	0.790	0.783 (0.06)	0.896	0.887	0.881 (0.05)
+ EEG full	0.588	0.583	0.572 (0.04)	0.778	0.774	0.772+ (0.05)	0.928	0.927	0.926* 0.03
+ EEG θ	0.564	0.569	0.535 (0.08)	0.805	0.792	0.791+ (0.04)	0.922	0.919	0.917* (0.03)
+ EEG α	0.596	0.593	0.563 (0.08)	0.775	0.781	0.772* (0.08)	0.920	0.917	0.916* (0.03)
+ EEG β	0.605	0.597	0.580 (0.08)	0.802	0.797	0.792+ (0.05)	0.920	0.914	0.914* (0.04)
+ EEG γ	0.640	0.625	0.611+ (0.09)	0.787	0.780	0.776+ (0.05)	0.905	0.905	0.901 (0.04)
+ $\theta + \alpha + \beta + \gamma$	0.599	0.579	0.558 (0.07)	0.800	0.794	0.786+ (0.05)	0.909	0.910	0.907 (0.04)

We report precision (P), recall (R), F₁-score and the standard deviation (std) between five runs. The best results per column are marked in bold. Significance is indicated on the F₁-score with asterisks: * denotes $p < 0.05$ (uncorrected), + denotes $p < 0.003$ (Bonferroni corrected p -value).

combination of BERT embeddings and EEG data does improve the relation detection models. Generally, the results show a decreasing maximal performance per task with increasing task complexity measured in terms of the number of classes (see section 4.5 for a detailed analysis).

Furthermore, the results for the multi-modal architecture using the convolutional EEG decoding component are presented in Table 7 for binary sentiment analysis, Table 8 for ternary sentiment analysis, and Table 9 for relation detection. The results of this model architecture yield higher overall results, whereas the trend across tasks is similar to the models using the recurrent EEG decoding component, i.e., considerable improvements for both sentiment analysis tasks, but for relation detection the most notable improvements are achieved with the BERT embeddings. This validates the popular choice of convolutional neural networks for EEG classification tasks (Schirmer et al., 2017; Craik et al., 2019). While recurrent neural networks are often

used in NLP and linguistic modeling (due to the left-to-right processing mechanism), CNNs have shown better performance at learning feature weights from noisy data (e.g., Kvist and Lockvall Rhodin, 2019). Hence, our convolutional EEG decoding component is able to better extract the task-relevant linguistic processing information from the input data.

To assess the results, we perform statistical significance testing with respect to the text baseline in a bootstrap test as described in Dror et al. (2018) over the F₁-scores of the five runs of all tasks. We compare the results of the multi-modal models using text and EEG data to the uni-modal text baseline. In addition, we apply the Bonferroni correction to counteract the problem of multiple comparisons. We choose this conservative correction because of the dependencies between the datasets used (Dror et al., 2017). Under the Bonferroni correction, the global null hypothesis is rejected if $p < \alpha/N$, where N is the number of hypotheses (Bonferroni, 1936). In our setting, $\alpha =$

TABLE 8 | Ternary sentiment analysis results of the multi-modal model using the convolutional EEG decoding component.

Model	Randomly initialized			GloVe			BERT		
	P	R	F ₁ (std)	P	R	F ₁ (std)	P	R	F ₁ (std)
Baseline	0.408	0.384	0.351 (0.07)	0.510	0.507	0.496 (0.06)	0.722	0.714	0.710 (0.05)
+ noise	0.359	0.388	0.334 (0.09)	0.494	0.484	0.476 (0.07)	0.715	0.683	0.684 (0.05)
+ ET	0.417	0.399	0.372 (0.05)	0.509	0.512	0.500 (0.07)	0.721	0.687	0.670 (0.05)
+ EEG full	0.365	0.384	0.333 (0.08)	0.488	0.484	0.476 (0.06)	0.738	0.724	0.723+ (0.04)
+ EEG θ	0.389	0.372	0.330 (0.06)	0.511	0.495	0.477 (0.06)	0.727	0.718	0.716+ (0.05)
+ EEG α	0.357	0.382	0.331 (0.11)	0.534	0.525	0.515+ (0.06)	0.732	0.715	0.713+ (0.04)
+ EEG β	0.425	0.418	0.378 (0.08)	0.534	0.529	0.520+ (0.05)	0.727	0.717	0.715 (0.04)
+ EEG γ	0.404	0.406	0.360 (0.08)	0.539	0.521	0.514 (0.06)	0.733	0.725	0.721+ (0.04)
+ $\theta + \alpha + \beta + \gamma$	0.384	0.402	0.354 (0.10)	0.517	0.504	0.488 (0.05)	0.733	0.717	0.715 (0.06)

We report precision (P), recall (R), F₁-score and the standard deviation (std) between five runs. The best results per column are marked in bold. Significance is indicated on the F₁-score with asterisks: * denotes $p < 0.05$ (uncorrected), + denotes $p < 0.003$ (Bonferroni corrected p-value).

TABLE 9 | Relation detection results of the multi-modal model using the convolutional EEG decoding component.

Model	Randomly initialized			GloVe			BERT		
	P	R	F ₁ (std)	P	R	F ₁ (std)	P	R	F ₁ (std)
Baseline	0.404	0.525	0.452 (0.04)	0.501	0.609	0.539 (0.05)	0.522	0.788	0.623 (0.05)
+ noise	0.424	0.299	0.342 (0.06)	0.547	0.441	0.486 (0.06)	0.532	0.493	0.511 (0.07)
+ ET	0.415	0.307	0.345 (0.08)	0.447	0.413	0.428 (0.07)	0.558	0.665	0.593 (0.13)
+ EEG full	0.225	0.225	0.225 (0.06)	0.548	0.408	0.464 (0.07)	0.647	0.664	0.650 (0.09)
+ EEG θ	0.437	0.380	0.400 (0.05)	0.620	0.493	0.547 (0.05)	0.721	0.698	0.707+ (0.03)
+ EEG α	0.372	0.366	0.352 (0.12)	0.509	0.433	0.461 (0.12)	0.661	0.697	0.675+ (0.08)
+ EEG β	0.394	0.328	0.338 (0.09)	0.627	0.479	0.541 (0.05)	0.643	0.646	0.640 (0.11)
+ EEG γ	0.405	0.363	0.366 (0.09)	0.646	0.490	0.555 (0.04)	0.667	0.699	0.679+ (0.06)
+ $\theta + \alpha + \beta + \gamma$	0.324	0.227	0.257 (0.11)	0.460	0.436	0.437 (0.14)	0.610	0.562	0.584 (0.05)

We report precision (P), recall (R), F₁-score and the standard deviation (std) between five runs. The best results per column are marked in bold. Significance is indicated on the F₁-score with asterisks: * denotes $p < 0.05$ (uncorrected), + denotes $p < 0.003$ (Bonferroni corrected p-value).

0.05 and $N = 18$, accounting for the combination of the 3 embedding types and 6 EEG feature sets, namely broadband EEG; θ , α , β , and γ frequency bands; and all four frequency bands jointly. For instance, in **Table 7** the improvements in 6 configurations out of 18 are also still statistically significant under the Bonferroni correction (i.e., $p < 0.003$), showing that EEG signals bring significant improvements in the sentiment analysis task. In the results tables, we mark significant results under both the uncorrected and the Bonferroni corrected p-value.

4. DISCUSSION

The results show consistent improvements on both sentiment analysis tasks, whereas the benefits of using EEG data are only visible in specific settings for the relation detection task. EEG performs better than, or at least comparable to, eye tracking in many scenarios. This study shows the potential of decoding EEG for NLP and provides a good basis for future studies. Despite

the limited amount of data, these results suggest that augmenting NLP systems with EEG features is a generalizable approach.

In the following sections, we discuss these results from different angles. We contrast the performance of different EEG features, we compare the EEG results to the text baseline and multi-modal baselines (as described in section 2.3.2), and we analyze the effect of different word embedding types. Additionally, we explore the impact of varying training set sizes in a data ablation study. Finally, we investigate the possible reasons for the decrease in performance for the relation detection task, which we associate with the task complexity. We run all analyses with both the recurrent *and* the convolutional EEG components.

4.1. EEG Feature Analysis

We start by investigating the impact of the various EEG features included in our multi-modal models. Different neurocognitive aspects of language processing during reading are associated with

brain oscillations at various frequencies. We first give a short overview of the cognitive functions related to EEG frequency bands that are found in literature before discussing the insights of our results.

Theta activity reflects cognitive control and working memory (Williams et al., 2019), and increases when processing semantic anomalies (Prystauka and Lewis, 2019). Moreover, Bastiaansen et al. (2002) showed a frequency-specific increase in theta power as a sentence unfolds, possibly related to the formation of an episodic memory trace, or to incremental verbal working memory load. High theta power is also prominent during the effective semantic processing of language (Bastiaansen et al., 2005). *Alpha* activity has been related to attentiveness (Klimesch, 2012). Both theta and alpha ranges are sensitive to the lexical-semantic processes involved in language translation (Grabner et al., 2007). *Beta* activity has been involved in higher-order linguistic functions such as the discrimination of word categories and the retrieval of action semantics as well as semantic memory, and syntactic processes, which support meaning construction during sentence processing. There is evidence that suggests that beta frequencies are important for linking past and present inputs and the detection of novelty of stimuli, which are essential processes for language perception as well as production (Weiss and Mueller, 2012). Beta frequencies also affect decisions regarding relevance (Eugster et al., 2014). In reading, a stronger power-decrease in lower beta frequencies has been found for neutral compared to negative words (Scaltritti et al., 2020). Contrarily, emotional processing of pictures enhances *gamma* band power (Müller et al., 1999). Gamma-band activity has been used to detect emotions (Li and Lu, 2009), and increases during syntactic and semantic structure building (Prystauka and Lewis, 2019). In the gamma frequency band, a power increase was observed during the processing of correct sentences in multiple languages, but this effect was absent following semantic violations (Hald et al., 2006; Penolazzi et al., 2009). Frequency band features have often been used in deep learning methods for decoding EEG in other domains, such as mental workload and sleep stage classification (Craik et al., 2019).

The results show that our multi-modal models yield better results with filtered EEG frequency bands than using the broadband EEG signal on almost all tasks and embedding types, as well as on both EEG decoding components. Although all frequency band features show promising results on some embedding types and tasks (e.g., BERT embeddings and gamma features for binary sentiment analysis reported in **Table 4**), the results show no clear sign of a single frequency band outperforming the others (neither across tasks for a fixed embedding type, nor for a fixed task and across all embedding types). For the sentiment analysis tasks, where both EEG decoding components achieve significant improvements, theta and beta features most often achieve the highest results. As described above, brain activity in each frequency band reflects specific cognitive functions. The positive results achieved using theta band EEG features might be explained by the importance of this frequency band for successful semantic processing. Theta power is expected to rise with increasing language processing activity (Kosch et al., 2020). Various studies have shown that theta

oscillations are related to semantic memory retrieval and can be task-specific (e.g., Bastiaansen et al., 2005; Giraud and Poeppel, 2012; Marko et al., 2019). Overall, previous research shows how theta correlates with the cognitive processing involved in encoding and retrieving verbal stimuli (see Kahana, 2006 for a review), which supports our results. The good performance of the beta EEG features might on one hand be explained by the effect of the emotional connotation of words on the beta response (Scaltritti et al., 2020). On the other hand, the role of beta oscillations in syntactic and semantic unification operations during language comprehension (Bastiaansen and Hagoort, 2006; Meyer, 2018) is also supportive of our results.

Based on the complexity and extent of our results, it is unclear at this point whether a single frequency band is more informative for solving NLP tasks. Data-driven methods can help us to tease more information from the recordings by allowing us to test broader theories and task-specific language representations (Murphy et al., 2018), but our results also clearly show that restricting the EEG signal to a given frequency band is beneficial. More research is required in this area to specifically isolate the linguistic processing from the filtered EEG signals.

4.2. Comparison to Multi-Modal Baselines

The multi-modal EEG models often outperform the text baselines (at least for the sentiment analysis tasks). We now analyze how the EEG models compare to the two augmented baselines described in section 2.3.2 (i.e., eye tracking and models augmented with random noise). We find that EEG always performs better than or equal to the multi-modal text + eye tracking models. This shows how promising EEG is as a data source for multi-modal cognitive NLP. Although eye tracking requires less recording efforts, these results corroborate that EEG data contain more information about the cognitive processes occurring in the brain during language understanding.

As expected, the baselines augmented with random noise perform worse than the pure text baselines in all cases except for binary sentiment analysis with BERT embeddings. This model seems to deal exceptionally well with added noise. In the case of relation detection, when no improvement is achieved (e.g., for randomly initialized embeddings), the added noise harms the models similarly to adding EEG signals. It becomes clear for this task that adding the full broadband EEG features is worse than adding random noise (except with BERT embeddings), but some of the frequency band features clearly outperform the augmented noise baseline.

4.3. Comparison of Embedding Types

Our baseline results show that contextual embeddings outperform the non-contextual methods across all tasks. Arora et al. (2020) also compared randomly initialized, GloVe and BERT embeddings and found that with smaller training sets, the difference in performance between these three embedding types is larger. This is in accordance with our results, which show that the type of embedding has a large impact on the baseline performance on all three tasks. The improvements of adding EEG data in all three tasks are especially noteworthy when using BERT embeddings. In combination with the EEG

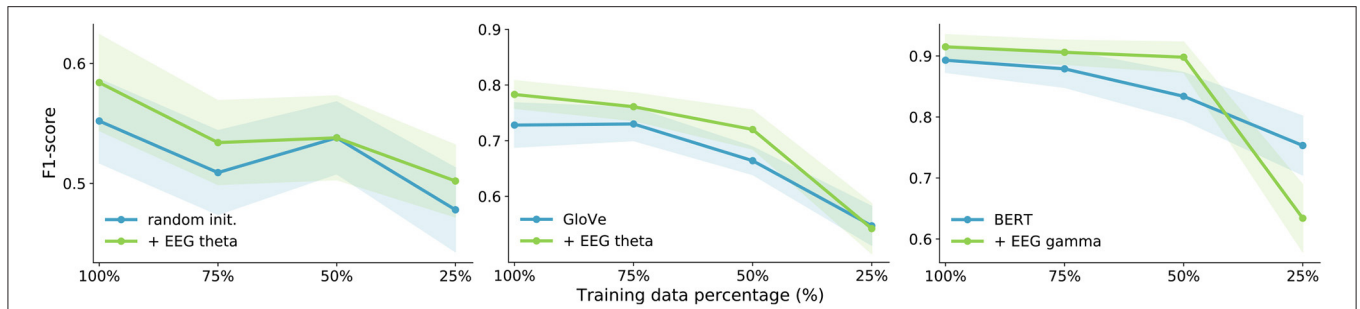


FIGURE 4 | Data ablation for all three word embedding types for the binary sentiment analysis task using the *recurrent EEG decoding component*. The shaded areas represent the standard deviations.

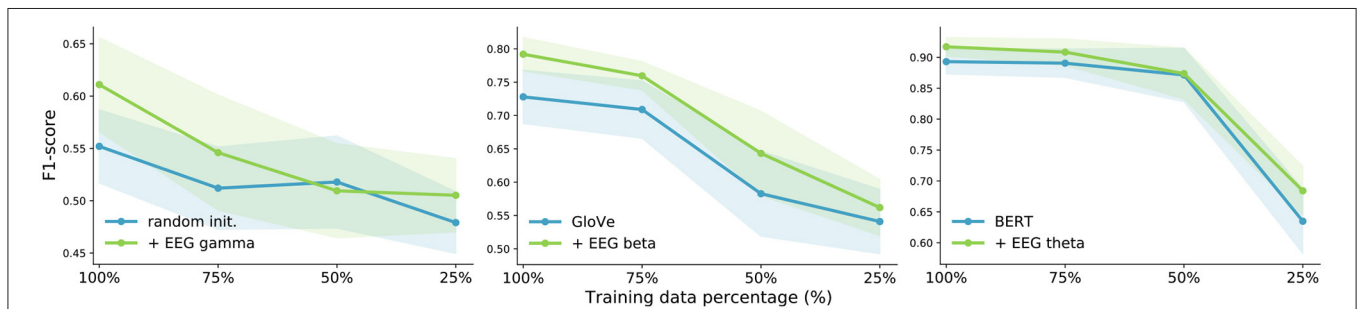


FIGURE 5 | Data ablation for all three word embedding types for the binary sentiment analysis task using the *convolutional EEG decoding component*. The shaded areas represent the standard deviations.

data, these embeddings achieve improvements across all settings, including the full EEG broadband data as well as all individual and combined frequency bands. This shows that state-of-the-art contextualized word representations such as BERT are able to interact positively with human language processing data in a multi-modal learning scenario.

Augmenting our baseline with EEG data on the binary sentiment analysis tasks results in approximately +3% F₁-score across all the different embeddings with the recurrent EEG component. The gain is slightly lower at +1% for all the embeddings in the ternary sentiment classification task. While there is no significant gain for relation detection with random and GloVe embeddings, the improvements with BERT embeddings reach up to +7%. This shows that the improvements gained by adding EEG signals are not only dependent on the task, but also on the embedding type. In foresight, this finding might be useful in the future, when new embeddings will improve the baseline performance even further while possibly also increasing the gain from the EEG signals.

4.4. Data Ablation

One of the challenges of NLP is to learn as much as possible from limited resources. Unlike most machine learning models, one of the most striking aspects of human learning is the ability to learn new words or concepts from limited numbers of examples (Lake et al., 2015). Using cognitive language processing data may allow us to take a step toward meta-learning, the process of discovering the cognitive processes that are used to tackle a task

in the human brain (Griffiths et al., 2019), and in turn be able to improve the generalization abilities of NLP models. Humans can learn from very few examples, while machines, particularly deep learning models, typically need many examples. Perhaps this advantage in humans is due to their multi-modal learning mechanisms (Linzen, 2020).

Therefore, we analyze the impact of adding EEG features to our NLP models with less training data. We performed data ablation experiments for all three tasks. The most conclusive results were achieved on binary sentiment analysis. Randomly initialized embeddings unsurprisingly suffer a lot when reducing training data. The results are shown in **Figures 4, 5**, for both EEG decoding components. We present the results for the best-performing frequency bands only. The largest gain from EEG data is obtained with only 50% of the training data with GloVe and BERT embeddings, which is as little as 105 training sentences. These experiments emphasize the potential of EEG signals for NLP especially when dealing with very small amounts of training data and using popular word embedding types.

4.5. Task Complexity Ablation

From the previously described results, one hypothesis on the reason why augmenting the baseline with EEG data lowers the performance in the relation detection task with randomly initialized and GloVe embeddings lies in the complexity of the task. More concretely, we measure the complexity by counting the number of classes the model needs to learn. Generally, more complex tasks (in terms of number of classes) require more data

TABLE 10 | Binary relation detection results for both EEG decoding components for the relation types *Job Title* and *Visited* using GloVe embeddings.

Job title vs. None	Recurrent EEG decoding			Convolutional EEG decoding		
	Precision	Recall	F ₁ -score	Precision	Recall	F ₁ -score
GloVe	0.789	0.776	0.767 (0.05)	0.789	0.776	0.767 (0.05)
GloVe + EEG full	0.792	0.782	0.773 (0.06)	0.796	0.793	0.789 (0.05)
GloVe + EEG γ	0.780	0.788	0.774 (0.1)	0.817	0.811	0.808 (0.03)
Visited vs. None	Precision	Recall	F ₁ -score	Precision	Recall	F ₁ -score
GloVe	0.762	0.756	0.734 (0.1)	0.762	0.756	0.734 (0.1)
GloVe + EEG full	0.756	0.759	0.745 (0.1)	0.766	0.758	0.750 (0.09)
GloVe + EEG γ	0.773	0.768	0.754 (0.1)	0.819	0.795	0.795 (0.09)

The best result in each column is marked in bold.

to generalize (see for instance Li et al., 2018). Therefore, it is clear that with a fixed amount of data, the impact of augmenting the feature space with additional information (in this case EEG data) is also less visible for the more complex tasks. We see a decrease in performance with increasing complexity over the three evaluated tasks with all embeddings except for BERT. Therefore, we validate this hypothesis by simplifying the relation detection task by reducing the number of classes from 11 to 2. We create binary relation detection tasks for the two most frequent relation types *Job Title* and *Visited* (see **Figure 1**). For example, we classify all the samples containing the relation *Job Title* (184 samples) against all samples with no relation (219 samples).

We train these additional models with GloVe embeddings, since these did not show any significant improvements when augmented with EEG data on the full relation detection task. The results for the full broadband EEG features and the best frequency band from the previous convolutional results (γ) are shown in **Table 10**. It is evident that with the simplification of the relation detection task into binary classification tasks, EEG signals are able to boost the performance of the non-contextualized GloVe embeddings and achieve considerable improvements over the text baseline. The gains are similar as for binary sentiment analysis for both EEG decoding components. This confirms our hypothesis that the EEG features tested yield good results on simple tasks, but more research is needed to achieve improvements on more complex tasks. Note that, as mentioned previously, this is not the case for BERT embeddings, which outperform the baselines on all NLP tasks.

4.6. Conclusion

We presented a large-scale study about leveraging electrical brain activity signals during reading comprehension for augmenting machine learning models of semantic language understanding tasks, namely, sentiment analysis and relation detection. We analyzed the effects of different EEG features and compared the multi-modal models to multiple baselines. Moreover, we compared the improvements gained from the EEG signals on three different types of word embeddings. Not only did we test the effect of varying training set sizes, but also tasks of various difficulty levels (in terms of number of classes).

We achieve consistent improvements with EEG across all three embedding types. The models trained with BERT embeddings yield significant performance increases on all NLP tasks. However, for randomly initialized and GloVe embeddings the improvement magnitude decreases for more difficult tasks. For these two types of embedding, the improvement for the binary and ternary sentiment analysis tasks ranges between 1 and 4% F₁-score. For relation detection, a multi-class and multi-label sequence classification task, it was not possible to achieve any improvements unless the task complexity is substantially reduced. Therefore, our experiments show that state-of-the-art contextualized word embeddings combined with careful EEG feature selection achieve good results in multi-modal learning. Moreover, we find that in the tasks where the multi-modal architecture does achieve considerable improvements, the convolutional EEG decoding component yields even higher results than the recurrent component.

To sum up, we capitalize on the advantages of electroencephalography data to examine if and which EEG features can serve to augment language understanding models. While our results show that there is linguistic information in the EEG signal complementing the text features, more research is needed to isolate language-specific brain activity features. More generally, this work paves the way for more in-depth EEG-based NLP studies.

DATA AVAILABILITY STATEMENT

Publicly available datasets were analyzed in this study. The datasets used for this study can be found in the Open Science Framework: ZuCo 1.0 (<https://osf.io/q3zws/>) and ZuCo 2.0 (<https://osf.io/2urht/>). The code used for this study can be found on GitHub: <https://github.com/DS3Lab/eego>.

ETHICS STATEMENT

The studies involving human participants were reviewed and approved by Ethics Commission, University of Zurich. The participants provided their written informed consent to participate in this study.

AUTHOR CONTRIBUTIONS

NH: conceptualization, methodology, software, formal analysis, writing—original draft, review, and editing, and visualization. CR: conceptualization, methodology, formal analysis, and writing—original draft, review, and editing. BG: methodology, software, and writing—review and editing. MB: conceptualization, methodology, and writing—review and editing. MT: conceptualization and writing—review

and editing. NL and CZ: conceptualization, writing—review and editing, supervision, and funding acquisition. All authors contributed to the article and approved the submitted version.

FUNDING

MB was supported by research grant 34437 from VILLUM FONDEN.

REFERENCES

- Affolter, N., Egressy, B., Pascual, D., and Wattenhofer, R. (2020). Brain2word: decoding brain activity for language generation. *arXiv preprint arXiv:2009.04765*.
- Alday, P. M. (2019). M/EEG analysis of naturalistic stories: a review from speech to language processing. *Lang. Cogn. Neurosci.* 34, 457–473. doi: 10.1080/23273798.2018.1546882
- Armeni, K., Willems, R. M., and Frank, S. (2017). Probabilistic language models in cognitive neuroscience: promises and pitfalls. *Neurosci. Biobehav. Rev.* 83, 579–588. doi: 10.1016/j.neubiorev.2017.09.001
- Arora, S., May, A., Zhang, J., and Ré, C. (2020). Contextual embeddings: when are they worth it? *arXiv preprint arXiv:2005.09117*. doi: 10.18653/v1/2020.acl-main.236
- Artemova, E., Bakarov, A., Artemov, A., Burnaev, E., and Sharaev, M. (2020). Data-driven models and computational tools for neurolinguistics: a language technology perspective. *J. Cogn. Sci.* 21, 15–52.
- Barnes, J., Klinger, R., and im Walde, S. S. (2017). “Assessing state-of-the-art sentiment models on state-of-the-art sentiment datasets,” in *Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis* (Copenhagen), 2–12. doi: 10.18653/v1/W17-5202
- Barnes, J., Øvrelid, L., and Veldal, E. (2019). “Sentiment analysis is not solved! assessing and probing sentiment classification,” in *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP* (Florence: Association for Computational Linguistics), 12–23. doi: 10.18653/v1/W19-4802
- Barnes, J., Veldal, E., and Øvrelid, L. (2021). Improving sentiment analysis with multi-task learning of negation. *Nat. Lang. Eng.* 27, 249–269. doi: 10.1017/S1351324920000510
- Barrett, M., Bingel, J., Hollenstein, N., Rei, M., and Søgaard, A. (2018a). “Sequence classification with human attention,” in *Proceedings of the 22nd Conference on Computational Natural Language Learning* (Brussels) 302–312. doi: 10.18653/v1/K18-1030
- Barrett, M., Bingel, J., Keller, F., and Søgaard, A. (2016). “Weakly supervised part-of-speech tagging using eye-tracking data,” in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, Vol. 2 (Osaka), 579–584. doi: 10.18653/v1/P16-2094
- Barrett, M., González-Garduño, A. V., Frermann, L., and Søgaard, A. (2018b). “Unsupervised induction of linguistic categories with records of reading, speaking, and writing,” in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Vol. 1 (New Orleans, LA), 2028–2038. doi: 10.18653/v1/N18-1184
- Barrett, M., and Hollenstein, N. (2020). Sequence labelling and sequence classification with gaze: novel uses of eye-tracking data for natural language processing. *Lang. Linguist. Compass* 14, 1–16. doi: 10.1111/lnc3.12396
- Bastiaansen, M., and Hagoort, P. (2006). Oscillatory neuronal dynamics during language comprehension. *Prog. Brain Res.* 159, 179–196. doi: 10.1016/S0079-6123(06)59012-0
- Bastiaansen, M. C., Linden, M. v. D., Keurs, M. T., Dijkstra, T., et al. (2005). Theta responses are involved in lexical-semantic retrieval during language processing. *J. Cogn. Neurosci.* 17, 530–541. doi: 10.1162/0898929053279469
- Bastiaansen, M. C., Van Berkum, J. J., and Hagoort, P. (2002). Event-related theta power increases in the human EEG during online sentence processing. *Neurosci. Lett.* 323, 13–16. doi: 10.1016/S0304-3940(01)02535-6
- Beinborn, L., Abnar, S., and Choenni, R. (2019). Robust evaluation of language-brain encoding experiments. *Int. J. Comput. Linguist. Appl.*
- Beres, A. M. (2017). Time is of the essence: a review of electroencephalography (EEG) and event-related brain potentials (ERPs) in language research. *Appl. Psychophysiol. Biofeedback* 42, 247–255. doi: 10.1007/s10484-017-9371-3
- Bishop, C. M. (1995). Training with noise is equivalent to tikhonov regularization. *Neural Comput.* 7, 108–116. doi: 10.1162/neco.1995.7.1.108
- Bisk, Y., Holtzman, A., Thomason, J., Andreas, J., Bengio, Y., Chai, J., et al. (2020). “Experience grounds language,” in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 8718–8735. doi: 10.18653/v1/2020.emnlp-main.703
- Bonferroni, C. (1936). Teoria statistica delle classi e calcolo delle probabilita. *Pubblazioni del R Istituto Superiore di Scienze Economiche e Commerciali di Firenze* 8, 3–62.
- Bruns, A. (2004). Fourier-, Hilbert- and wavelet-based signal analysis: are they really different approaches? *J. Neurosci. Methods* 137, 321–332. doi: 10.1016/j.jneumeth.2004.03.002
- Clevert, D.-A., Unterthiner, T., and Hochreiter, S. (2015). Fast and accurate deep network learning by exponential linear units (ELUs). *arXiv preprint arXiv:1511.07289*.
- Clifton, C., Staub, A., and Rayner, K. (2007). “Eye movements in reading words and sentences,” in *Eye Movements*, eds R. P. G. Van Gompel, M. H. Fischer, W. S. Murray, and R. L. Hill (Amsterdam: Elsevier), 341–371. doi: 10.1016/B978-008044980-7/50017-3
- Craik, A., He, Y., and Contreras-Vidal, J. L. (2019). Deep learning for electroencephalogram (EEG) classification tasks: a review. *J. Neural Eng.* 16:031001. doi: 10.1088/1741-2552/ab0ab5
- Culotta, A., McCallum, A., and Betz, J. (2006). “Integrating probabilistic extraction models and data mining to discover relations and patterns in text,” in *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association of Computational Linguistics* (New York, NY), 296–303. doi: 10.3115/1220835.1220873
- Degno, F., Loberg, O., and Liversedge, S. P. (2021). Co-registration of eye movements and fixation-related potentials in natural reading: practical issues of experimental design and data analysis. *Collabra Psychol.* 7, 1–28. doi: 10.1525/collabra.18032
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). “BERT: pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Vol. 1 (Minneapolis, MN), 4171–4186.
- Dimigen, O., Sommer, W., Hohlfield, A., Jacobs, A. M., and Kliegl, R. (2011). Coregistration of eye movements and EEG in natural reading: analyses and review. *J. Exp. Psychol. Gen.* 140:552. doi: 10.1037/a0023885
- Do, T. H., Nguyen, D. M., Tsiligianni, E., Cornelis, B., and Deligiannis, N. (2017). Multiview deep learning for predicting Twitter users’ location. *arXiv preprint arXiv:1712.08091*.
- Dror, R., Baumer, G., Bogomolov, M., and Reichart, R. (2017). Replicability analysis for natural language processing: testing significance with multiple datasets. *Trans. Assoc. Comput. Linguist.* 5, 471–486. doi: 10.1162/tac1_a_00074

- Dror, R., Baumer, G., Shlomov, S., and Reichart, R. (2018). "The Hitchhiker's guide to testing statistical significance in natural language processing," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics* (Melbourne, VIC), 1383–1392. doi: 10.18653/v1/P18-1128
- Ehringer, B. V., and Dimigen, O. (2019). Unfold: an integrated toolbox for overlap correction, non-linear modeling, and regression-based EEG analysis. *PeerJ* 7:e7838. doi: 10.7717/peerj.7838
- Ephrat, A., Halperin, T., and Peleg, S. (2017). "Improved speech reconstruction from silent video," in *Proceedings of the IEEE International Conference on Computer Vision Workshops* (Venice), 455–462. doi: 10.1109/ICCVW.2017.61
- Ettinger, A. (2020). What BERT is not: lessons from a new suite of psycholinguistic diagnostics for language models. *Trans. Assoc. Comput. Linguist.* 8, 34–48. doi: 10.1162/tacl_a_00298
- Eugster, M. J., Ruotsalo, T., Spapé, M. M., Kosunen, I., Barral, O., Ravaja, N., et al. (2014). "Predicting term-relevance from brain signals," in *Proceedings of the 37th International ACM SIGIR Conference on Research & Development in Information Retrieval* (ACM), 425–434. doi: 10.1145/2600428.2609594
- Fawaz, H. I., Lucas, B., Forestier, G., Pelletier, C., Schmidt, D. F., Weber, J., et al. (2020). Inceptiontime: finding alexnet for time series classification. *Data Mining Knowledge Discov.* 34, 1936–1962. doi: 10.1007/s10618-020-00710-y
- Foster, C., Dharmaretnam, D., Xu, H., Fyshe, A., and Tzanetakis, G. (2018). "Decoding music in the human brain using EEG data," in *2018 IEEE 20th International Workshop on Multimedia Signal Processing (MMSP)* (Vancouver, BC), 1–6. doi: 10.1109/MMSP.2018.8547051
- Frank, S. L., Otten, L. J., Galli, G., and Vigliocco, G. (2015). The ERP response to the amount of information conveyed by words in sentences. *Brain Lang.* 140, 1–11. doi: 10.1016/j.bandl.2014.10.006
- Frank, S. L., and Willems, R. M. (2017). Word predictability and semantic similarity show distinct patterns of brain activity during language comprehension. *Lang. Cogn. Neurosci.* 32, 1192–1203. doi: 10.1080/23273798.2017.1323109
- Friederici, A. D. (2000). The developmental cognitive neuroscience of language: a new research domain. *Brain Lang.* 71, 65–68. doi: 10.1006/brln.1999.2214
- Fyshe, A., Talukdar, P. P., Murphy, B., and Mitchell, T. M. (2014). "Interpretable semantic vectors from a joint model of brain-and text-based meaning," in *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics* (Baltimore, MD), 489–499. doi: 10.3115/v1/P14-1046
- Gauthier, J., and Ivanova, A. (2018). Does the brain represent words? An evaluation of brain decoding studies of language understanding. *arXiv preprint arXiv:1806.00591*. doi: 10.32470/CCN.2018.1237-0
- Giraud, A.-L., and Poeppel, D. (2012). Cortical oscillations and speech processing: emerging computational principles and operations. *Nat. Neurosci.* 15:511. doi: 10.1038/nn.3063
- Grabner, R. H., Brunner, C., Leeb, R., Neuper, C., and Pfurtscheller, G. (2007). Event-related eeg theta and alpha band oscillatory responses during language translation. *Brain Res. Bull.* 72, 57–65. doi: 10.1016/j.brainresbull.2007.01.001
- Griffiths, T. L., Callaway, F., Chang, M. B., Grant, E., Krueger, P. M., and Lieder, F. (2019). Doing more with less: meta-reasoning and meta-learning in humans and machines. *Curr. Opin. Behav. Sci.* 29, 24–30. doi: 10.1016/j.cobeha.2019.01.005
- Hald, L. A., Bastiaansen, M. C., and Hagoort, P. (2006). EEG theta and gamma responses to semantic violations in online sentence processing. *Brain Lang.* 96, 90–105. doi: 10.1016/j.bandl.2005.06.007
- Hale, J., Dyer, C., Kuncoro, A., and Brennan, J. R. (2018). "Finding syntax in human encephalography with beam search," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics* (Melbourne, VIC), 2727–2736. doi: 10.18653/v1/P18-1254
- Hochreiter, S., and Schmidhuber, J. (1997). Long short-term memory. *Neural Comput.* 9, 1735–1780. doi: 10.1162/neco.1997.9.8.1735
- Hollenstein, N., Barrett, M., and Beinborn, L. (2020a). "Towards best practices for leveraging human language processing signals for natural language processing," in *Proceedings of the Second Workshop on Linguistic and Neurocognitive Resources* (Marseille), 15–27.
- Hollenstein, N., Barrett, M., Troendle, M., Bigioli, F., Langer, N., and Zhang, C. (2019a). Advancing NLP with cognitive language processing signals. *arXiv preprint arXiv:1904.02682*.
- Hollenstein, N., de la Torre, A., Langer, N., and Zhang, C. (2019b). "CogniVal: a framework for cognitive word embedding evaluation," in *Proceedings of the 23rd Conference on Computational Natural Language Learning*. doi: 10.18653/v1/K19-1050
- Hollenstein, N., Rotsztein, J., Troendle, M., Pedroni, A., Zhang, C., and Langer, N. (2018). ZuCo, a simultaneous EEG and eye-tracking resource for natural sentence reading. *Sci. Data.* 5, 1–13. doi: 10.1038/sdata.2018.291
- Hollenstein, N., Troendle, M., Zhang, C., and Langer, N. (2020b). "ZuCo 2.0: a dataset of physiological recordings during natural reading and annotation," in *Proceedings of the 12th Language Resources and Evaluation Conference* (Marseille), 138–146.
- Hollenstein, N., and Zhang, C. (2019). "Entity recognition at first sight: Improving NER with eye movement information," in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (Minneapolis, MN).
- Howard, J., and Ruder, S. (2018). "Universal language model fine-tuning for text classification," in *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics* (Melbourne, VIC), 328–339. doi: 10.18653/v1/P18-1031
- Kahana, M. J. (2006). The cognitive correlates of human brain oscillations. *J. Neurosci.* 26, 1669–1672. doi: 10.1523/JNEUROSCI.3737-05c.2006
- Kandylaki, K. D., and Bornkessel-Schlesewsky, I. (2019). From story comprehension to the neurobiology of language. *Lang. Cogn. Neurosci.* 34, 405–410. doi: 10.1080/23273798.2019.1584679
- Kiela, D., Grave, E., Joulin, A., and Mikolov, T. (2018). "Efficient large-scale multi-modal classification," in *Thirty-Second AAAI Conference on Artificial Intelligence* (New Orleans, CA).
- Klimesch, W. (2012). Alpha-band oscillations, attention, and controlled access to stored information. *Trends Cogn. Sci.* 16, 606–617. doi: 10.1016/j.tics.2012.10.007
- Kosch, T., Schmidt, A., Thanheiser, S., and Chuang, L. L. (2020). "One does not simply RSVP: mental workload to select speed reading parameters using electroencephalography," in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI), 1–13. doi: 10.1145/3313831.3376766
- Kvist, E., and Lockvall Rhodin, S. (2019). *A Comparative Study Between MLP and CNN for Noise Reduction on Images: The Impact of Different Input Dataset Sizes and the Impact of Different Types of Noise on Performance*. Stockholm: KTH Royal Institute of Technology School of Electrical Engineering and Computer Science.
- Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science* 350, 1332–1338. doi: 10.1126/science.aab3050
- Lemhöfer, K., and Broersma, M. (2012). Introducing LexTALE: a quick and valid lexical test for advanced learners of english. *Behav. Res. Methods* 44, 325–343. doi: 10.3758/s13428-011-0146-0
- Li, J., Liu, Y., Yin, R., Zhang, H., Ding, L., and Wang, W. (2018). "Multiclass learning: from theory to algorithm," in *Advances in Neural Information Processing Systems*, eds S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Montreal, QC: Curran Associates, Inc.), 1586–1595.
- Li, M., and Lu, B.-L. (2009). "Emotion classification based on gamma-band EEG," in *Engineering in Medicine and Biology Society, 2009. EMBC 2009. Annual International Conference of the IEEE* (Minneapolis, MN: IEEE), 1223–1226.
- Ling, S., Lee, A. C., Armstrong, B. C., and Nestor, A. (2019). How are visual words represented? Insights from EEG-based visual word decoding, feature derivation and image reconstruction. *Hum. Brain Mapp.* 40, 5056–5068. doi: 10.1002/hbm.24757
- Linzen, T. (2020). How can we accelerate progress towards human-like linguistic generalization? *arXiv preprint arXiv:2005.00955*. doi: 10.18653/v1/2020.acl-main.465
- Lipton, Z. C., Berkowitz, J., and Elkan, C. (2015). A critical review of recurrent neural networks for sequence learning. *arXiv preprint arXiv:1506.00019*.
- Long, Y., Lu, Q., Xiang, R., Li, M., and Huang, C.-R. (2017). "A cognition based attention model for sentiment analysis," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (Copenhagen), 462–471. doi: 10.18653/v1/D17-1048

- Marko, M., Cimrová, B., and Riečanský, I. (2019). Neural theta oscillations support semantic memory retrieval. *Sci. Rep.* 9, 1–10. doi: 10.1038/s41598-019-53813-y
- Mathias, S., Kanojia, D., Mishra, A., and Bhattacharyya, P. (2020). A survey on using gaze behaviour for natural language processing. *Proc. IJCAI*. 4907–4913. doi: 10.24963/ijcai.2020/683
- McClelland, J. L., Hill, F., Rudolph, M., Baldrige, J., and Schütze, H. (2020). Placing language in an integrated understanding system: next steps toward human-level performance in neural language models. *Proc. Natl. Acad. Sci. U.S.A.* 117, 25966–25974. doi: 10.1073/pnas.1910416117
- Meyer, L. (2018). The neural oscillations of speech processing and language comprehension: state of the art and emerging mechanisms. *Eur. J. Neurosci.* 48, 2609–2621. doi: 10.1111/ejn.13748
- Mishra, A., Kanojia, D., Nagar, S., Dey, K., and Bhattacharyya, P. (2017). “Leveraging cognitive features for sentiment analysis,” in *Proceedings of The 20th Conference on Computational Natural Language Learning* (Vancouver, BC), 156–166. doi: 10.18653/v1/K16-1016
- Morency, L.-P., and Baltrušaitis, T. (2017). “Multimodal machine learning: integrating language, vision and speech,” in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts* (Vancouver, BC), 3–5. doi: 10.18653/v1/P17-5002
- Müller, M. M., Keil, A., Gruber, T., and Elbert, T. (1999). Processing of affective pictures modulates right-hemispheric gamma band EEG activity. *Clin. Neurophysiol.* 110, 1913–1920. doi: 10.1016/S1388-2457(99)00151-0
- Murphy, B., and Poesio, M. (2010). “Detecting semantic category in simultaneous EEG/MEG recordings,” in *Proceedings of the NAACL HLT 2010 First Workshop on Computational Neurolinguistics* (Los Angeles, CA: Association for Computational Linguistics), 36–44.
- Murphy, B., Wehbe, L., and Fyshe, A. (2018). Decoding language from the brain. *Lang. Cogn. Comput. Models* 53. doi: 10.1017/9781316676974.003
- Muttenthaler, L., Hollenstein, N., and Barrett, M. (2020). Human brain activity for machine attention. *arXiv preprint arXiv:2006.05113*.
- Naselaris, T., Kay, K. N., Nishimoto, S., and Gallant, J. L. (2011). Encoding and decoding in fMRI. *Neuroimage* 56, 400–410. doi: 10.1016/j.neuroimage.2010.07.073
- Nurse, E., Mashford, B. S., Yepes, A. J., Kiral-Kornek, I., Harrer, S., and Freestone, D. R. (2016). “Decoding EEG and LFP signals using deep learning: heading TrueNorth,” in *Proceedings of the ACM International Conference on Computing Frontiers* (Como), 259–266. doi: 10.1145/2903150.2903159
- Pedroni, A., Bahreini, A., and Langer, N. (2019). Automagic: standardized preprocessing of big EEG data. *Neuroimage*. 200, 460–473. doi: 10.1016/j.neuroimage.2019.06.046
- Pennington, J., Socher, R., and Manning, C. D. (2014). “Glove: global vectors for word representation,” in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing* (Doha), 1532–1543. doi: 10.3115/v1/D14-1162
- Penolazzi, B., Angrilli, A., and Job, R. (2009). Gamma EEG activity induced by semantic violation during sentence reading. *Neurosci. Lett.* 465, 74–78. doi: 10.1016/j.neulet.2009.08.065
- Pfeiffer, C., Hollenstein, N., Zhang, C., and Langer, N. (2020). Neural dynamics of sentiment processing during naturalistic sentence reading. *Neuroimage* 218:116934. doi: 10.1016/j.neuroimage.2020.116934
- Poeppel, D. (2014). The neuroanatomic and neurophysiological infrastructure for speech and language. *Current Opinion in Neurobiology* 28, 142–149. doi: 10.1016/j.conb.2014.07.005
- Poeppel, D., Emmorey, K., Hickok, G., and Pyykkänen, L. (2012). Towards a new neurobiology of language. *J. Neurosci.* 32, 14125–14131. doi: 10.1523/JNEUROSCI.3244-12.2012
- Poria, S., Hazarika, D., Majumder, N., and Mihalcea, R. (2020). Beneath the tip of the iceberg: current challenges and new directions in sentiment analysis research. *IEEE Trans. Affect. Comput.* doi: 10.1109/TAFFC.2020.3038167
- Prystauka, Y., and Lewis, A. G. (2019). The power of neural oscillations to inform sentence comprehension: a linguistic perspective. *Lang. Linguist. Compass* 13:e12347. doi: 10.1111/lnc3.12347
- Ramachandram, D., and Taylor, G. W. (2017). Deep multimodal learning: a survey on recent advances and trends. *IEEE Signal Process. Mag.* 34, 96–108. doi: 10.1109/MSP.2017.2738401
- Rotsztein, J., Hollenstein, N., and Zhang, C. (2018). “ETH-DS3Lab at SemEval-2018 Task 7: effectively combining recurrent and convolutional neural networks for relation classification and extraction,” in *Proceedings of The 12th International Workshop on Semantic Evaluation* (New Orleans, CA), 689–696. doi: 10.18653/v1/S18-1112
- Sassenhagen, J., and Fiebach, C. J. (2020). Traces of meaning itself: encoding distributional word vectors in brain activity. *Neurobiol. Lang.* 1, 54–76. doi: 10.1162/nol_a_00003
- Sato, N., and Mizuhara, H. (2018). Successful encoding during natural reading is associated with fixation-related potentials and large-scale network deactivation. *eNeuro* 5, 1–12. doi: 10.1523/ENEURO.0122-18.2018
- Scaltritti, M., Suitner, C., and Peressotti, F. (2020). Language and motor processing in reading and typing: insights from beta-frequency band power modulations. *Brain Lang.* 204:104758. doi: 10.1016/j.bandl.2020.104758
- Schirrmeyer, R. T., Springenberg, J. T., Fiederer, L. D. J., Glasstetter, M., Eggensperger, K., Tangermann, M., et al. (2017). Deep learning with convolutional neural networks for EEG decoding and visualization. *Hum. Brain Mapp.* 38, 5391–5420. doi: 10.1002/hbm.23730
- Schwartz, D., Toneva, M., and Wehbe, L. (2019). “Inducing brain-relevant bias in natural language processing models,” in *Advances in Neural Information Processing Systems*, eds H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (Vancouver, BC), 14100–14110.
- Sereno, S. C., and Rayner, K. (2003). Measuring word recognition in reading: eye movements and event-related potentials. *Trends Cogn. Sci.* 7, 489–493. doi: 10.1016/j.tics.2003.09.010
- Smith, N. J., and Kutas, M. (2015a). Regression-based estimation of erp waveforms: I. The RERP framework. *Psychophysiology* 52, 157–168. doi: 10.1111/psyp.12317
- Smith, N. J., and Kutas, M. (2015b). Regression-based estimation of erp waveforms: II. Nonlinear effects, overlap correction, and practical considerations. *Psychophysiology* 52, 169–181. doi: 10.1111/psyp.12320
- Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C. D., Ng, A., et al. (2013). “Recursive deep models for semantic compositionality over a sentiment treebank,” in *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing* (Seattle, WA), 1631–1642.
- Stemmer, B., and Connolly, J. F. (2012). The EEG/ERP technologies in linguistic research. *Methodological and Analytic Front. Lexical Res.* 47:337. doi: 10.1075/bct.47.16ste
- Sun, C., Qiu, X., Xu, Y., and Huang, X. (2019). “How to fine-tune BERT for text classification?” in *China National Conference on Chinese Computational Linguistics* (Hainan; Cham: Springer), 194–206. doi: 10.1007/978-3-030-32381-3_16
- Sun, P., Anumanchipalli, G. K., and Chang, E. F. (2020). Brain2char: a deep architecture for decoding text from brain recordings. *J. Neural Eng.* doi: 10.1088/1741-2552/abc742
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., et al. (2015). “Going deeper with convolutions,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1–9. doi: 10.1109/CVPR.2015.7298594
- Toneva, M., and Wehbe, L. (2019). “Interpreting and improving natural-language processing (in machines) with natural language-processing (in the brain),” in *Advances in Neural Information Processing Systems*, eds H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (Vancouver, BC: Curran Associates, Inc.), 14928–14938.
- Tsai, Y.-H. H., Bai, S., Liang, P. P., Kolter, J. Z., Morency, L.-P., and Salakhutdinov, R. (2019). “Multimodal transformer for unaligned multimodal language sequences,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 6558–6569. doi: 10.18653/v1/P19-1656
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., et al. (2017). “Attention is all you need,” in *Advances in Neural Information Processing Systems* (Long Beach, CA), 5998–6008.
- Wehbe, L., Vaswani, A., Knight, K., and Mitchell, T. (2014). “Aligning context-based statistical models of language with brain activity during reading,” in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (Doha), 233–243. doi: 10.3115/v1/D14-1030
- Weiss, S., and Mueller, H. M. (2012). “Too many betas do not spoil the broth”: the role of beta brain oscillations in language processing. *Front. Psychol.* 3:201. doi: 10.3389/fpsyg.2012.00201

- Williams, C. C., Kappen, M., Hassall, C. D., Wright, B., and Krigolson, O. E. (2019). Thinking theta and alpha: mechanisms of intuitive and analytical reasoning. *Neuroimage* 189, 574–580. doi: 10.1016/j.neuroimage.2019.01.048
- Winkler, I., Haufe, S., and Tangermann, M. (2011). Automatic classification of artifactual ICA-components for artifact removal in EEG signals. *Behav. Brain Funct.* 7:30. doi: 10.1186/1744-9081-7-30
- Yin, W., Kann, K., Yu, M., and Schütze, H. (2017). Comparative study of CNN and RNN for natural language processing. *arXiv preprint arXiv:1702.01923*.

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2021 Hollenstein, Renggli, Glaus, Barrett, Troendle, Langer and Zhang. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.