

Human brain activity for machine attention

Lukas Muttenthaler[†] Nora Hollenstein[‡] Maria Barrett^{†,⊙}

[†] Dept. of Computer Science, University of Copenhagen

[‡] Dept. of Computer Science, ETH Zurich

[⊙] Dept. of Computer Science, IT University of Copenhagen

mnd926@alumni.ku.dk, noraho@inf.ethz.ch, mjb@di.ku.dk

Abstract

Cognitively inspired NLP leverages human-derived data to teach machines about language processing mechanisms. Recently, neural networks have been augmented with behavioral data to solve a range of NLP tasks spanning syntax and semantics. We are the first to exploit neuroscientific data, namely electroencephalography (EEG), to inform a neural attention model about language processing of the human brain. The challenge in working with EEG data is that features are exceptionally rich and need extensive pre-processing to isolate signals specific to text processing. We devise a method for finding such EEG features to supervise machine attention through combining theoretically motivated cropping with random forest tree splits. After this dimensionality reduction, the pre-processed EEG features are capable of distinguishing two reading tasks retrieved from a publicly available EEG corpus. We apply these features to regularise attention on relation classification, and show that EEG is more informative than strong baselines. This improvement depends on both the cognitive load of the task and the EEG frequency domain. Hence, informing neural attention models with EEG signals is beneficial but requires further investigation to understand which dimensions are the most useful across NLP tasks.

1 Introduction

Cognitively inspired NLP is a research field at the intersection of Cognitive Neuroscience and Natural Language Processing (NLP), which lately has received growing attention. Cognitive Neuroscience aims to investigate cognitive processes that occur in the human brain through high-level explanations, whereas NLP has the overall objective to teach machines to read and understand human language. The recent merge of these fields has led to the overarching goal of introducing human bias (Wilson et al., 2015; Schwartz et al., 2019; Toneva and We-

hbe, 2019) into machines, and hence augment neural networks with cognitive data in solving NLP tasks.

Human readers process words without notable effort. Human text processing can be studied e.g. through eye tracking (ET) records from reading where fixation durations on word level are robustly correlated with cognitive load (Drieghe et al., 2005; Fitzsimmons and Drieghe, 2011; Rayner et al., 2011). Recently, numerous studies have proven that words represented as ET features can help a wide range of NLP tasks including part-of-speech tagging (Barrett et al., 2016), relation detection (Hollenstein et al., 2019a), and sentiment analysis (Mishra et al., 2017, 2016; Barrett et al., 2018).

Fixation durations reliably correlate with cognitive load (Rayner et al., 2011) but prolonged fixation duration will not help differentiate which cognitive process occurs. Therefore, eye movements are considered indirect measures of human text processing whereas electroencephalography (EEG) and fMRI technologies are direct measures of human brain activity. EEG measures electric activity in the brain. When used non-intrusively, a number of electrodes are placed on the scalp to measure brain surface activity. In the field of Cognitive Science, EEG data plays a vital role to explain various phenomena such as cognitive load (Antonenko et al., 2010; Zarjam et al., 2011; Kumar and Kumar, 2016). It is even possible to decode human cortical activity to synthesize audible speech (Anumanchipalli et al., 2019).

In NLP, however, there has just been a single study that investigated how EEG data can enhance machines' ability to perform named entity recognition, sentiment, and relation classification (Hollenstein et al., 2019a). In contrast to eye movements, EEG contains not only signals about cognitive load and mental state, but *all* brain activity within the temporal domain (including facial muscle activity which must be removed from the data after record-

ing). Therefore, EEG requires considerably more de-noising and pre-processing steps before any text processing signals are exploitable for NLP.

Providing a rigorous and clear approach for the latter is the main aim of this study. On top of that, we go one step further and show how our approach improves performance on three sequence classification tasks - even with access to limited amounts of brain data.

1.1 Contributions

- We devise a method for extracting human language processing signals from EEG recordings of a publicly available corpus (Hollenstein et al., 2018). As a sanity check, we demonstrate that these multi-dimensional feature vectors let us distinguish between two different reading tasks, namely Normal Reading (NR) and Annotation Reading (AR).
- We inject this human bias into a multi-task neural model for sequence classification. In doing so, we use the most informative EEG signals on the word-level to regularize the attention mechanism of an RNN.
- We observe that the differences in EEG signals from NR and AR affect NLP downstream performances differently. Moreover, we show that downstream performance further varies as a function of EEG frequency domain.

Together, these insights have decisive implications about which human EEG signals are most suitable to inject into Machine Learning (ML) models for enhancing their language processing performance. We make our code publicly available. ¹

2 Related work

Recently, a number of studies has investigated how external cognitive signals, and thus the injection of human bias, can enhance the capacity of artificial neural networks to understand natural language (Hollenstein et al., 2019a; Strzyz et al., 2019; Schwartz et al., 2019; Toneva and Wehbe, 2019; Gauthier and Levy, 2019), and vice versa, how language processing in neural networks might enhance our understanding of human language processing (Hollenstein et al., 2019b). Others scrutinized whether machine attention deviates from

human attention when disentangling visual scenes (Das et al., 2016).

Most studies, however, focused on the use of ET data, and hence exploited gaze features on the word level. Some utilized gaze features as word embeddings to inform neural networks about which syntactic linguistic features humans deem decisive in their language processing. In so doing, they have successfully enhanced state-of-the-art models for named entity recognition (NER; Hollenstein and Zhang (2019)), part-of-speech tagging (Barrett et al., 2016) and dependency parsing (Strzyz et al., 2019). Others have drawn attention to the enhancement of semantic disentanglement, and improved tasks such as sarcasm detection (Mishra et al., 2017), or sentiment analysis (Mishra et al., 2016) through leveraging human gaze.

One recent study, from which we take inspiration, exploited gaze information to regularize attention in a multi-task-like setting for sequence classification (Barrett et al., 2018). Here, gaze information improved grammatical error detection, hate speech detection, and sentiment classification. The authors enabled neural networks to utilize human attention during training time without the need of accessing this information during test time. Although this study is similar to ours and serves as the foundation for our code, we go one step beyond, and regularize attention with human brain activity instead of indirect proxies such as gaze.

3 EEG feature extraction

In this section, we introduce the EEG corpus used for this study, then explain our feature extraction and dimensionality reduction approach in detail.

3.1 EEG data

In this study, we use the recently created Zurich Cognitive Language Processing Corpus (ZuCo)² (Hollenstein et al., 2018). It is a corpus of simultaneous ET and EEG recordings of 12 English native speakers reading individual sentences. Thus, through the ET and EEG alignment, neural activities are available for single words in English. The corpus contains cognitive data from the following three reading tasks:

Task 1 Normal Reading (NR) of 400 sentences that contain sentiment annotations from the Stanford Sentiment Treebank (Socher et al., 2013). In NR, reading comprehension was the sole task.

¹<https://github.com/LukasMut/NER-with-EEG-and-ET>

²<https://osf.io/q3zws/>

Task 2 NR of 300 sentences that contain relations between named entities extracted from the Wikipedia relation extraction corpus (Culotta et al., 2006).

Task 3 Annotation Reading (AR) of 407 sentences that contain relations between named entities, also from Culotta et al. (2006). Participants were required to annotate sentences while reading by pressing a key to answer whether one specific relation is present in the sentence. This results in a different cognitive load for the human reader.

Each of the 12 participants had to perform all three reading tasks. Participants completed the readings in two sessions on different days. For our experiments, we exclusively used data from NR (Task 2) and AR since these tasks contain text from the same domain.

EEG features are recorded simultaneously with a corresponding ET feature (e.g., First Fixation Duration). As such, eight 105-dimensional EEG vectors were extracted for each ET feature corresponding to the EEG activity during the specific eye-movement event (i.e., ET feature) w.r.t. this word. Each of the eight 105-dimensional vectors corresponds to one of the following EEG frequency domains: θ_1 (4-6 Hz), θ_2 (6.5-8 Hz), α_1 (8.5-10 Hz), α_2 (10.5-13 Hz), β_1 (13.5-18 Hz), β_2 (18.5-30 Hz), γ_1 (30.5-40 Hz) and γ_2 (40-49.5 Hz) (Hollenstein et al., 2018). Each of the 105 dimensions corresponds to a specific electrode on a 128-channel EEG cap used for recording (see Figure 1). 23 electrodes were removed through Automagic (Pedroni et al., 2019) prior to any data analysis, since they did not contain signals relevant for cognitive processing due to their position (Hollenstein et al., 2018).

Given that n = number of recorded ET features, m = number of EEG frequency domains, and, k = number of electrodes, this results in $n \times m \times k = 4200$ EEG features per word per ET feature, where $n = 5$, $m = 8$, and, $k = 105$. Since most of these 4200 features reflect signals that are not relevant to cognitive text processing, we apply various feature extraction techniques to reduce the dimensionality.

3.2 EEG feature reduction

We split the EEG data for NR and AR into a train, development and test split with 70/15/15% of the data. Splits were created w.r.t. sentence-level samples. We calculated feature reduction on the train

set, tuned it on development splits, and evaluated the reduced features on a held-out test set.

We first investigated how to pre-process the EEG features to find differences between AR and NR. This was done to scrutinize potential differences in EEG signals as a function of cognitive load, and evaluate whether the dimensionality reduction is able to capture those. As a validation step, we also tuned the method for classifying Task 1 into its first and second halves respectively. Participants completed Task 2 and the first half of Task 1 in the first session, and Task 3 and the second half of Task 2 in the second session. Such experimental designs may bias the data with session-specific effects. Hence, this validation step serves both as a data quality check of the dimensionality reduction as well as of the noise removal. Ideally, the text processing signal should remain coherent for both halves of Task 1. For both steps, we employed a Random Forest classifier to the train set to find the most informative EEG channels for this task.

Firstly, to extract the k most predictive features per word, per ET feature, and per EEG frequency domain we reduced m and n based on literature.

Reducing m We binned the eight frequency domains (see above) into the four general frequency bands, θ (4-8 Hz), α (8.5-13 Hz), β (13.5-30 Hz) and γ (30.5-49.5 Hz). This strategy was applied to manually decrease the dimensionality prior to exploiting any machine learning techniques, and reduce computational cost at an early stage. To yield binned frequency domains, we calculated the average power spectrum per electrode for each of the four frequency pairs (e.g., $\bar{x}(\alpha_1, \alpha_2)$), thus $m = 4$. Due to the fact that γ frequencies mainly relate to emotionality (Zheng and Lu, 2015; Li and Lu, 2009; Oathes et al., 2008; Luo et al., 2008), and not to attentiveness, we further reduced $m = 4$ to $m = 3$, and computed embeddings in low-dimensional brain space for α , β , and θ only.

Reducing n In initial data analyses, kernel density estimates (KDE) and t -test bootstrapping showed that power spectra vary across EEG frequency domains but are highly correlated among the different ET features. Therefore, we decided to extract EEG features that correspond to total reading time (TRT), as this ET feature covers all activity related to an individual word, and has proven to be the most informative gaze feature in previous studies (Hollenstein and Zhang, 2019; Hollenstein




Figure 1: Differences in EEG power spectra for TRT in specific electrodes between Normal Reading (NR) and Annotation Reading (AR) for the four frequency bands θ , α , β . Best viewed in color. Electrodes colored in green denote higher EEG power spectra for AR, whereas red electrodes indicate higher EEG signals for NR. All differences are statistically significant with $p < 0.01$. Blue colored electrodes refer to no or non-significant differences between the two reading tasks. All electrodes that are not colored are electrodes that have been excluded during pre-processing prior to data analysis. The reference electrode (Cz), however, was not excluded as a pre-processing step, but not considered during statistical analyses since it always has the minimum value of 0.

et al., 2019a). Hence, $n = 1$.

Furthermore, we exploited a Random Forest classifier (Breiman, 2001) with 100 trees. A random forest³ reveals the respective feature indices it requires to solve the classification task. Those can easily be mapped to electrodes on the EEG cap, and hence help to find the brain regions where activity varied across the different reading tasks.

We extracted the k most informative features to distinguish between NR and AR according to our Random Forest implementation. We conducted experiments over three different values for k (i.e., $k \in [5, 15, 30]$), and scrutinized performances of different k -dimensional embeddings as input for an LSTM (Hochreiter and Schmidhuber, 1996) to classify sentences in the respective tasks. This notably reduced the feature space from 4200 dimensions to $n \times m \times k \in [15, 45, 90]$, where $n = 1, m = 3, k = [5, 15, 30]$.

We computed t -test bootstrapping for each electrode in the four general frequency bands between NR and AR to inspect which electrodes show higher power spectra in which task (see Figure 1). All electrodes in the left temporal cortex, which is responsible for both language comprehension and production, that show significantly higher power spectra for AR compared to NR, were included in the k most predictive features to solve the binary

classification task through Random Forest. This means that Random Forest assigned more importance to brain signals that are enhanced for AR. We suspect that stronger EEG activity in these (fronto-)temporal areas of the human brain are due to higher cognitive load for AR compared to NR in language comprehension and production areas. However, further investigation must go into this line of research to draw definite conclusions.

3.3 Classifying task and session using reduced EEG features

As a sanity check, we test the reduced EEG features in a simple, binary classification task to verify that we isolated the cognitive processing of text, and removed noise. We, therefore, use the reduced features to classify NR and AR, as well as whether signals were part of the first or the second half of Task 1. The latter was done to rule out the possibility that differences in brain signals between tasks occurred merely as a function of daytime.

LSTM architecture We use a vanilla LSTM⁴ with 1 hidden layer, 50 hidden units per layer, a layer dropout rate of 0.5, and an Adam optimizer (Kingma and Ba, 2015) with the default learning rate of $\alpha = 0.001$. Since the main goal was to predict the two different reading tasks, we minimized binary cross-entropy loss through mini-batch train-

³Scikit learn (Pedregosa et al., 2011) with default parameters. However, we set the bootstrap parameter to false to use the entire data set to build each tree which led to better classification results in initial experiments.

⁴implemented end-to-end in PyTorch (Paszke et al., 2019).

TASK	VALIDATION			TEST		
	15	45	90	15	45	90
NR-AR	100	100	93.8	100	98.4	85.9
Ses1-Ses2	48.4	57.8	50.5	45.	50.	48.8

Table 1: LSTM binary classification accuracies in % with EEG word embeddings of different dimensions for two different tasks on both development and test set. Embedding dimensions were extracted through Random Forest tree splits for each task individually.

ing with a batch size of 32.

$$\vec{h}_t = \text{LSTM}(\vec{h}_{t-1}, \mathbf{x}_t), t = 1, \dots, |x| \quad (1)$$

LSTM denotes the LSTM function (Hochreiter and Schmidhuber, 1996), h_t is the hidden state at time step t , and x_t represents the word input at current time step, where x was embedded in k -dimensional EEG space ($x_t \in \mathbb{R}^k$).

The results in Table 1 show that NR and AR appear to result in different brain signals and can thus be easily classified into two distinct classes in line with Figure 1. We experimented with different values for k and found lower values of k to be better at task classification. In contrast, session could not be classified into its respective days. It is, therefore, likely that our features capture a great amount of cognitive text processing signals and little noise. We visualize each word that appeared in sentences in NR or AR respectively by the two most useful features (found by the Random Forest) in Figure 2. The plot shows that even in 2D space, EEG features well reflect the differences between NR and AR.




Figure 2: Visualisation of the most informative feature for classifying task. Words embedded in two-dimensional brain space through Random Forest tree splits.

4 Sequence labeling with EEG attention

EEG activity can be used as a supervision, or guidance, signal for the attention function in neural networks. In the following section, we outline our attention-based model for sequence classification, present the task-specific datasets, and explain how we transform EEG to attention scores.

4.1 EEG features to scalar

Attention scores are scalar values that weigh their respective hidden word representations accordingly. To regularize attention, we were thus required to further reduce each k -dimensional word vector x_t represented through EEG activity to a single dimension. We experimented with different approaches such as averaging across all k electrodes or taking the maximum EEG power spectrum per x_t . Initial experiments on the development set revealed that taking the mean over a k dimensional word vector leads to notably worse results than max-pooling. This might be due to information loss of brain activity as a result of averaging.

Domain specific scores The final attention scores were computed through taking the maximum electrode value per k dimensional word embedding for each frequency band. k was one of [15, 45, 90] for concatenated embeddings, and one of [5, 15, 30] for embeddings per frequency domain. To yield values within the range [0, 1), we normalized each EEG attention score by the maximum attention score of the respective sentence. We observed that dividing each normalized attention score by some small constant e leads to better performance. We assume this is due to the fact that EEG attention scores are somewhat peaky prior to dividing by e . Thus, the computation is as follows:

$$\mathbf{a}_t^i = \max(\mathbf{x}_t \in \mathbb{R}^k), t = 1, \dots, |x| \quad (2)$$


$$\mathbf{a}_t^i = \frac{\mathbf{a}_t^i}{\max(\mathbf{a}^i)}, t = 1, \dots, |x| \quad (3)$$

$$\mathbf{a}_t^i = \frac{\mathbf{a}_t^i}{e}, t = 1, \dots, |x| \quad (4)$$


where x_t denotes a word representation embedded in k -dimensional EEG space at time step t for a sentence i . To compute $\mathbf{a}^i$ we did not exploit concatenated EEG embeddings but used isolated EEG embeddings for each of the three frequency domains α, β, θ . Therefore, k was one of [5, 15, 30]. The constant e was set to 2. The computed EEG

TASK	SOURCE	TRAIN		VAL	TEST
		n sents	% positive	n sents	n sents
Relation Detection	SemEval 2010	8,096	19.3	1,361	1,372
Relation Detection	Wikipedia	1,733	10.0	361	354
NE Detection	Ontonotes 5.0	89,389	29.7	11,289	11,318

Table 2: Overview of the data sets.



(a) Normal Reading (NR)



(b) Annotation Reading (AR)

Figure 3: Max-pooled EEG attention scores $\mathbf{a}_i$ (averaged over all 12 participants) for two example sentences read in Normal Reading (NR) and Annotation Reading (AR) respectively.

attention scores served as inputs for our multi-task sequence classification model to supervise attention in the auxiliary task. Hence, final attention weights α^i were yielded through passing the EEG attention scores through the softmax function. The latter computation, however, happened automatically during training and was not done externally.

$$\alpha^i = \text{softmax}(\mathbf{a}^i) \quad (5)$$

Max-pooled EEG attention scores for two example sentences read in NR and AR respectively are depicted in Figure 3. We observe that brain activity transitions between words are smoother in NR compared to AR. We observe that for AR, in particular, there are higher activations for words relating to the relation AWARD.

4.2 Model

The model is an adaptation of Rei and Søgaard (2018) leveraged by Barrett et al. (2018).⁵ It is a

⁵[https://github.com/coastalcph/Sequence_classification_with_human_](https://github.com/coastalcph/Sequence_classification_with_human_attention)

biLSTM architecture that jointly learns the model parameters and the attention function by alternating training (Luong et al., 2016), much related to multi-task learning (Dong et al., 2015; Søgaard and Goldberg, 2016). The inputs are token-level labelled sequences of EEG scalars to learn the attention function (auxiliary task) and a set of sentence-level labelled sequences for training the model parameters (main task).

If the data point is from the main task, we perform normal training and model parameter update through comparing the model’s class prediction ($\hat{y}$) against the true label (y) on the sentence level.

$$L_{sent} = \sum_i (y^i - \hat{y}^i)^2 \quad (6)$$

If the data point is sampled from the EEG corpus, however, we do not update model parameters. We only modify the attention weights by minimizing the squared error between the EEG value and the attention score as described below.

$$L_{tok} = \sum_i \sum_t (\mathbf{a}_t^i - \hat{\mathbf{a}}_t^i)^2 \quad (7)$$

The model has no access to EEG signals during testing.

4.3 Experiments

Recall from Section 3.1 that participants read sentences from the Wikipedia relation extraction corpus (Culotta et al., 2006). We employ three widely used Relation Extraction and NER benchmark data sets for English against baseline models without supervised attention and models whose attention was either supervised through eye-tracking data as in Barrett et al. (2018) or word frequencies computed on the British National Corpus (Kilgarriff, 1995). We perform binary classification and adapt all datasets as described below to obtain sentence-level labels. As such, the main task was a k -vs.-the-rest binary sentence classification task. Overall

attention

statistics about the data sets are displayed in Table 2.

4.3.1 Relation detection

SemEval 2010 We used the SemEval 2010 Task 8 data set that defines the task as a multi-way classification of semantic relations between pairs of nominals (Hendrickx et al., 2010). The data set contains the following nine distinct relations: CAUSE-EFFECT, INSTRUMENT-AGENCY, PRODUCT-PRODUCER, CONTENT-CONTAINER, ENTITY-ORIGIN, ENTITY-DESTINATION, COMPONENT-WHOLE, MEMBER-COLLECTION and MESSAGE-TOPIC. Each sentence that contained the relations ENTITY-ORIGIN or ENTITY-DESTINATION served as a positive example of this relation. We have chosen those relations due to their higher frequency compared to other relations. Positive examples were tested against sentences that consisted of one of the remaining seven relations.

Wikipedia The data set provided by Culotta et al. (2006) contains Wikipedia articles labeled with 53 relation types. Since part of this dataset is included in ZuCo we filtered those sentences. We chose the sentences including the most frequent relation JOB TITLE as positive examples.

4.4 Named entity detection

Ontonotes 5.0 We use the four CoNLL-2003 NER labels PER, LOC, ORG, MISC in the Ontonotes 5.0 data set (Weischedel et al., 2013) as positive examples.

5 Results

SemEval 2010 The results are presented in Table 3 and we observe that EEG attention scores clearly help to solve the task. The best EEG-augmented model, 30k AR (θ), is better than all baselines. The most notable improvements are mainly due to a higher recall. Precision scores appear to be similar across all models although BNC word frequency augmented model show slightly lower precision than the rest.

Wikipedia The $F1$ score of the best EEG-augmented model, 5k NR (θ), outperformed all baselines (see Table 3). The BNC-augmented model provides a strong baseline, which could be explained through both the significantly larger number of data points available for BNC word frequencies and due to the fact that word frequencies

highly correspond to entities which are crucial to link entities.

Ontonotes The results show that all EEG augmented models outperform the baseline by a small margin (see Table 3). The performance improvement for recall is again notable. Precision scores appear to be comparable across all models but the baseline. The best model is 30k NR (α).

6 Discussion

Human brain activity appears to help machine attention in attending toward the most crucial words in a sentence. However, there are differences between the exploited EEG attention scores - measured as the improvement in performance on a particular task. This is dependent on both frequency domain and reading task of the EEG signals, as well as the number of features captured in the embeddings. In general, we observe smaller performance gains for NER compared to Relation Extraction. We suppose this is due to NER results being fairly strong in general for English language data (Lample et al., 2016; Hollenstein and Zhang, 2019). Hence, additional support through cognitive data cannot enhance performance much.

Reading tasks Brain activity extracted from sentences read in NR appear to be more useful to detect named entities compared to EEG signals from AR. This is not surprising since in an NR setting participants read sentences without any additional task to perform (see 3.1), and therefore read each sentence until its end (Hollenstein et al., 2018). On the other hand, brain signals distilled from AR guide the model to detect relations in a given sentence more effectively. This might be due to the fact that the readers were required to search for the respective relation while reading the sentence. Thus, participants drew particular attention to the decisive words that form the relation (Figure 3).

Frequency domains Performance enhancement differed depending on the frequency domain from which EEG attention scores were extracted. Both θ and α frequency bands show more useful signals and lead to better performance compared to β across all tasks. Lower frequency bands such as θ (4-8 Hz) and α (8.5-13 Hz) are linked to cognitive control (Williams et al., 2019) and attentiveness (Klimesch, 2012), respectively. This might explain why brain signals from those domains are particularly useful to guide machine attention. Higher

ATTENTION	SEMEVAL 2010			WIKI			ONTONOTES		
	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1
baseline	80.03	63.21	70.56	54.44	55.00	54.67	88.90	64.46	74.72
BNCFreqInv	78.30	58.00	66.52	61.39	60.00	60.64	91.56	67.38	77.61
MeanFixCont	79.59	60.36	68.62	59.99	58.00	58.75	92.21	66.59	77.33
5k NR (α)	77.90	65.29	70.98	58.76	51.00	54.38	90.41	67.84	77.51
5k NR (θ)	79.70	61.21	69.14	65.47	57.00	60.74	91.02	67.18	77.28
15k NR (α)	79.66	63.93	70.91	54.58	51.00	52.52	91.25	67.36	77.50
15k NR (β)	78.58	64.43	70.79	61.09	57.00	58.82	91.85	67.00	77.47
15k NR (θ)	79.91	60.86	69.04	58.67	55.00	56.65	91.66	67.10	77.47
30k NR (α)	77.88	64.57	70.53	52.83	49.00	50.77	91.04	67.68	77.63
30k NR (θ)	80.32	61.14	69.33	56.88	53.00	54.64	92.25	66.68	77.40
5k AR (α)	79.66	59.64	68.18	60.25	49.00	53.80	91.19	66.84	77.13
5k AR (θ)	79.55	63.00	70.29	56.42	49.00	52.42	91.01	67.01	77.18
15k AR (α)	79.02	64.57	71.04	56.70	60.00	57.75	90.97	67.17	77.23
15k AR (β)	79.15	63.79	70.60	57.16	51.00	53.60	90.96	67.00	77.16
15k AR (θ)	79.14	64.43	71.01	53.78	52.00	52.73	90.57	67.30	77.20
30k AR (α)	79.47	62.00	69.63	60.63	54.00	56.71	91.67	66.92	77.36
30k AR (θ)	79.96	64.50	71.34	59.63	52.00	55.16	91.03	66.87	77.10

Table 3: Relation detection and named entity detection. Results in %. Best scores per metric are displayed in bold face. All scores are averaged over five random seeds.

frequency domains such as β (13.5-30 Hz) and γ (30.5-49.5 Hz), however, are linked to motor activities (Pogossyan et al., 2009) and enhanced emotional responses (Li and Lu, 2009; Oathes et al., 2008), which explains why EEG power spectra in AR increase with the hertz rate (see Figure 1), but are less useful to supervise machine attention than brain signals from lower frequency bands (Table 3).

Dimensions We tested EEG attention scores that were max-pooled over individual frequency domain embeddings of different dimensionality. Overall, the attention scores distilled from 15- and 30-dimensional embeddings carry slightly more informative signals than attention scores extracted from 5-dimensional embeddings (see Table 3). The difference, however, is marginal, and for the Wikipedia relation detection, max-pooled attention scores over 5-dimensional θ embeddings outperform all other supervision signals. Contrary, to classify the word-level EEG signals of sentences into their respective reading tasks, lower-dimensional embeddings lead to better performance than higher-dimensional embeddings (see Table 1). The latter is not surprising since low-dimensional embeddings contain those EEG signals that differ the most between NR and AR. The higher the dimensionality of the embeddings, the more EEG signals are captured that differ less between the two reading tasks.

Improvements despite limited data What is compelling is the fact that we exploited little EEG data to create the attention scores - merely 300 sen-

tences for NR and 407 sentences for AR, averaged over 12 participants. Both ET and BNC frequency attention scores have > 10 times more sentences to train the auxiliary task of the model. Even with such few EEG data samples, useful signals could be extracted that help neural networks to understand language in a similar manner as the human brain does. We assume that more data will lead to even higher performance gains. We plan studies to investigate the latter.

7 Conclusions

We presented the first study that leverages EEG activity to inform machine attention about language processing mechanisms of the human brain. This is compelling for two reasons: First, we successfully isolated the text processing signals from noisy EEG data, which considerably reduced its dimensions. Second, we demonstrated that even a small number of EEG data points from human readers can benefit multi-task neural models for sequence classification. Note that the extracted attention scores may be exploited in various neural architectures that employ any form of attention to solve NLP tasks. Third, we showed that downstream performance varies as a function of both cognitive load and EEG frequency domains. This might have decisive implications about which EEG signals are to inject into neural models. We suspect that more data will provide deeper, and more thorough insights into the latter avenue.

Acknowledgements

We would like to thank Wolfgang Ganglberger, Martin Hebart, Joachim Bingel, and Desmond Elliot for fruitful comments on earlier versions of the paper.

References

- Pavlo Antonenko, Fred Paas, Roland Grabner, and Tamara Van Gog. 2010. Using electroencephalography to measure cognitive load. *Educational Psychology Review*, 22(4):425–438.
- Gopala K Anumanchipalli, Josh Chartier, and Edward F Chang. 2019. Speech synthesis from neural decoding of spoken sentences. *Nature*, 568(7753):493.
- Maria Barrett, Joachim Bingel, Nora Hollenstein, Marek Rei, and Anders Søgaard. 2018. [Sequence classification with human attention](#). In *Proceedings of the 22nd Conference on Computational Natural Language Learning, CoNLL 2018, Brussels, Belgium, October 31 - November 1, 2018*, pages 302–312. Association for Computational Linguistics.
- Maria Barrett, Joachim Bingel, Frank Keller, and Anders Søgaard. 2016. [Weakly supervised part-of-speech tagging using eye-tracking data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 2: Short Papers*. The Association for Computer Linguistics.
- Leo Breiman. 2001. [Random forests](#). *Mach. Learn.*, 45(1):5–32.
- Aron Culotta, Andrew McCallum, and Jonathan Betz. 2006. [Integrating probabilistic extraction models and data mining to discover relations and patterns in text](#). In *Human Language Technology Conference of the North American Chapter of the Association of Computational Linguistics, Proceedings, June 4-9, 2006, New York, New York, USA*. The Association for Computational Linguistics.
- Abhishek Das, Harsh Agrawal, Larry Zitnick, Devi Parikh, and Dhruv Batra. 2016. [Human attention in visual question answering: Do humans and deep networks look at the same regions?](#) In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 932–937. The Association for Computational Linguistics.
- Daxiang Dong, Hua Wu, Wei He, Dianhai Yu, and Haifeng Wang. 2015. [Multi-task learning for multiple language translation](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 1: Long Papers*, pages 1723–1732. The Association for Computer Linguistics.
- Denis Drieghe, Keith Rayner, and Alexander Pollatsek. 2005. Eye movements and word skipping during reading revisited. *Journal of Experimental Psychology: Human Perception and Performance*, 31(5):954.
- Gemma Fitzsimmons and Denis Drieghe. 2011. The influence of number of syllables on word skipping during reading. *Psychonomic Bulletin & Review*, 18(4):736–741.
- Jon Gauthier and Roger Levy. 2019. [Linking artificial and human neural representations of language](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 529–539. Association for Computational Linguistics.
- Iris Hendrickx, Su Nam Kim, Zornitsa Kozareva, Preslav Nakov, Diarmuid Ó Séaghdha, Sebastian Padó, Marco Pennacchiotti, Lorenza Romano, and Stan Szpakowicz. 2010. [Semeval-2010 task 8: Multi-way classification of semantic relations between pairs of nominals](#). In *Proceedings of the 5th International Workshop on Semantic Evaluation, SemEval@ACL 2010, Uppsala University, Uppsala, Sweden, July 15-16, 2010*, pages 33–38. The Association for Computer Linguistics.
- Sepp Hochreiter and Jürgen Schmidhuber. 1996. [LSTM can solve hard long time lag problems](#). In *Advances in Neural Information Processing Systems 9, NIPS, Denver, CO, USA, December 2-5, 1996*, pages 473–479. MIT Press.
- Nora Hollenstein, Maria Barrett, Marius Troendle, Francesco Bigiolli, Nicolas Langer, and Ce Zhang. 2019a. [Advancing NLP with cognitive language processing signals](#). *CoRR*, abs/1904.02682.
- Nora Hollenstein, Jonathan Rotsztejn, Marius Troendle, Andreas Pedroni, Ce Zhang, and Nicolas Langer. 2018. ZuCo, a simultaneous EEG and eye-tracking resource for natural sentence reading. *Scientific data*, 5:180291.
- Nora Hollenstein, Antonio de la Torre, Nicolas Langer, and Ce Zhang. 2019b. [CogniVal: A framework for cognitive word embedding evaluation](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 538–549, Hong Kong, China. Association for Computational Linguistics.
- Nora Hollenstein and Ce Zhang. 2019. [Entity recognition at first sight: Improving NER with eye movement information](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and*

- Short Papers*), pages 1–10. Association for Computational Linguistics.
- Adam Kilgarriff. 1995. BNC database and word frequency lists. <http://www.itri.brighton.ac.uk/~Adam.Kilgarriff/bnc-readme.html>.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Wolfgang Klimesch. 2012. Alpha-band oscillations, attention, and controlled access to stored information. *Trends in cognitive sciences*, 16(12):606–617.
- Naveen Kumar and Jyoti Kumar. 2016. Measurement of cognitive load in HCI systems using EEG power spectrum: an experimental study. *Procedia Computer Science*, 84:70–78.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. [Neural architectures for named entity recognition](#). In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 260–270. The Association for Computational Linguistics.
- Mu Li and Bao-Liang Lu. 2009. Emotion classification based on gamma-band EEG. In *2009 Annual International Conference of the IEEE Engineering in medicine and biology society*, pages 1223–1226. IEEE.
- Qian Luo, Derek Mitchell, Xi Cheng, Krystal Mondillo, Daniel McCaffrey, Tom Holroyd, Frederick Carver, Richard Coppola, and James Blair. 2008. Visual awareness, emotion, and gamma band synchronization. *Cerebral cortex*, 19(8):1896–1904.
- Minh-Thang Luong, Quoc V. Le, Ilya Sutskever, Oriol Vinyals, and Lukasz Kaiser. 2016. [Multi-task sequence to sequence learning](#). In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Abhijit Mishra, Kuntal Dey, and Pushpak Bhattacharyya. 2017. [Learning cognitive features from gaze data for sentiment and sarcasm classification using convolutional neural network](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 377–387. Association for Computational Linguistics.
- Abhijit Mishra, Diptesh Kanojia, Seema Nagar, Kuntal Dey, and Pushpak Bhattacharyya. 2016. [Leveraging cognitive features for sentiment analysis](#). In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning, CoNLL 2016, Berlin, Germany, August 11-12, 2016*, pages 156–166. ACL.
- Desmond J Oathes, William J Ray, Alissa S Yamasaki, Thomas D Borkovec, Louis G Castonguay, Michelle G Newman, and Jack Nitschke. 2008. Worry, generalized anxiety disorder, and emotion: Evidence from the eeg gamma band. *Biological psychology*, 79(2):165–170.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Andreas Pedroni, Amirreza Bahreini, and Nicolas Langer. 2019. [Automagic: Standardized preprocessing of big EEG data](#). *NeuroImage*, 200:460–473.
- Alek Pogosyan, Louise Doyle Gaynor, Alexandre Eusebio, and Peter Brown. 2009. Boosting cortical activity at beta-band frequencies slows movement in humans. *Current biology*, 19(19):1637–1641.
- Keith Rayner, Timothy J Slattery, Denis Drieghe, and Simon P Liversedge. 2011. Eye movements and word skipping during reading: effects of word length and predictability. *Journal of Experimental Psychology: Human Perception and Performance*, 37(2):514.
- Marek Rei and Anders Søgaard. 2018. [Zero-shot sequence labeling: Transferring knowledge from sentences to tokens](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 293–302. Association for Computational Linguistics.
- Dan Schwartz, Mariya Toneva, and Leila Wehbe. 2019. [Inducing brain-relevant bias in natural language processing models](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pages 14100–14110.

- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, EMNLP 2013, 18-21 October 2013, Grand Hyatt Seattle, Seattle, Washington, USA, A meeting of SIGDAT, a Special Interest Group of the ACL*, pages 1631–1642. ACL.
- Anders Søgaard and Yoav Goldberg. 2016. [Deep multi-task learning with low level tasks supervised at lower layers](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 2: Short Papers*. The Association for Computer Linguistics.
- Michalina Strzyz, David Vilares, and Carlos Gómez-Rodríguez. 2019. [Towards making a dependency parser see](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 1500–1506. Association for Computational Linguistics.
- Mariya Toneva and Leila Wehbe. 2019. [Interpreting and improving natural-language processing \(in machines\) with natural language-processing \(in the brain\)](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pages 14928–14938.
- Ralph Weischedel, Martha Palmer, Mitchell Marcus, Eduard Hovy, Sameer Pradhan, Lance Ramshaw, Nianwen Xue, Ann Taylor, Jeff Kaufman, Michelle Franchini, et al. 2013. Ontonotes release 5.0 ldc2013t19. *Linguistic Data Consortium*.
- Chad C. Williams, Mitchel Kappen, Cameron D. Hassall, Bruce Wright, and Olave E. Krigolson. 2019. [Thinking theta and alpha: Mechanisms of intuitive and analytical reasoning](#). *NeuroImage*, 189:574–580.
- Andrew Gordon Wilson, Christoph Dann, Christopher G. Lucas, and Eric P. Xing. 2015. [The human kernel](#). In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 2854–2862.
- Pega Zarjam, Julien Epps, and Fang Chen. 2011. Spectral EEG features for evaluating cognitive load. In *2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 3841–3844. IEEE.
- Wei-Long Zheng and Bao-Liang Lu. 2015. [Investigating critical frequency bands and channels for eeg-based emotion recognition with deep neural networks](#). *IEEE Trans. Auton. Ment. Dev.*, 7(3):162–175.