

Understanding language-elicited EEG data by predicting it from a fine-tuned language model

Dan Schwartz

Carnegie Mellon University
5000 Forbes Ave
Pittsburgh, PA 15213
drscharw@cs.cmu.edu

Tom Mitchell

Carnegie Mellon University
5000 Forbes Ave
Pittsburgh, PA 15213
tom.mitchell@cs.cmu.edu

Abstract

Electroencephalography (EEG) recordings of brain activity taken while participants read or listen to language are widely used within the cognitive neuroscience and psycholinguistics communities as a tool to study language comprehension. Several time-locked stereotyped EEG responses to word-presentations – known collectively as event-related potentials (ERPs) – are thought to be markers for semantic or syntactic processes that take place during comprehension. However, the characterization of each individual ERP in terms of what features of a stream of language trigger the response remains controversial. Improving this characterization would make ERPs a more useful tool for studying language comprehension. We take a step towards better understanding the ERPs by fine-tuning a language model to predict them. This new approach to analysis shows for the first time that all of the ERPs are predictable from embeddings of a stream of language. Prior work has only found two of the ERPs to be predictable. In addition to this analysis, we examine which ERPs benefit from sharing parameters during joint training. We find that two pairs of ERPs previously identified in the literature as being related to each other benefit from joint training, while several other pairs of ERPs that benefit from joint training are suggestive of potential relationships. Extensions of this analysis that further examine what kinds of information in the model embeddings relate to each ERP have the potential to elucidate the processes involved in human language comprehension.

1 Introduction

The cognitive processes involved in human language comprehension are complex and only partially identified. According to the dual-stream model of speech comprehension (Hickok and Poeppel, 2007), sound waves are first converted to




Figure 1: The electrodes from which each event-related potential was recorded in the data from Frank et al. (2015) (after figure 3 in (Frank et al., 2015)). The bottom portion of the figure shows a top-down schematic of the electrode locations with the nose facing towards the top of the page. Each ERP is the mean potential from all of the indicated electrodes during a specific time-window, creating a single scalar value per ERP per word. Overlapping circles indicate multiple ERPs recorded from the same electrode. The ELAN is measured from 125-175ms after stimulus onset, the LAN from 300-400ms, the N400 from 300ms-500ms, the EPNP from 400-600ms, the P600 from 500-700ms, and the PNP from 600-700ms.

phoneme-like features and further processed by a ventral stream that maps those features onto words and semantic structures, and a dorsal stream that (among other things) supports audio-short term memory. The mapping of words onto meaning is thought to be subserved by widely distributed regions of the brain that specialize in particular modalities — for example visual aspects of the

word *banana* reside in the occipital lobe of the brain and are activated when the word *banana* is heard (Kemmerer, 2014) — and the different representation modalities are thought to be integrated into a single coherent latent representation in the anterior temporal lobe (Ralph et al., 2010). While this part of meaning representation in human language comprehension is somewhat understood, much less is known about how the meanings of words are integrated together to form the meaning of sentences and discourses. One tool researchers use to study the integration of meaning across words is electroencephalography (EEG), which measures the electrical activity of large numbers of neurons acting in concert. EEG has the temporal resolution necessary to study the processes involved in meaning integration, and certain stereotyped electrical responses to word presentations, known as event-related potentials (ERPs), have been identified with some of the processes thought to contribute to comprehension.

In this work, we consider six ERP components that have been associated in the cognitive neuroscience and psycholinguistics literature with language processing and which we analyze in the data from Frank et al. (2015) (see Figure 1 for spatial and temporal definitions of these ERP components). Three of these — the N400, EPNP, and PNP responses — are primarily considered markers for semantic processing, while the other three — the P600, ELAN, and LAN responses — are primarily considered markers for syntactic processing. However, the neat division of the ERP responses into either semantic or syntactic categories is controversial. The N400 response has been very well studied (for an overview see (Kutas and Federmeier, 2011)) and it is well established that it is associated with semantic complexity, but the features of language that trigger the other ERP responses we consider here are poorly understood. We propose to use a neural network pretrained as a language model to probe what features of language drive these ERP responses, and in turn to probe what features of language mediate the cognitive processes that underlie human language comprehension, and especially the integration of meaning across words.

2 Background

While a full discussion of each ERP component and the features of language thought to trigger

each are beyond the scope of this document (for reviews see e.g. Frank et al. (2015), Kemmerer (2014), Kutas and Federmeier (2011), Kuperberg et al. (2003), and Van Petten and Luka (2012)), we introduce some basic features of ERP components to help in the discussion later. ERP components are electrical potential responses measured with respect to a baseline that are triggered by an event (in our case the presentation of a new word to a participant in an experiment). The name of each ERP component reflects whether the potential is positive or negative relative to the baseline. The N400 is so-named because it is **N**egative relative to a baseline (the baseline is typically recorded just before a word is presented at an electrode that is not affected by the ERP response) and because it peaks in magnitude at about **400ms** after a word is presented to a participant in an experiment. The P600 is **P**ositive relative to a baseline and peaks around **600ms** after a word is presented to a participant (though its overall duration is much longer and less specific in time than the N400). The post-N400 positivity is so-named because it is part of a biphasic response; it is a positivity that occurs after the negativity associated with the N400. The early post-N400 positivity (EPNP) is also part of a biphasic response, but the positivity has an earlier onset than the standard PNP. Finally, the LAN and ELAN are the left-anterior negativity and early left-anterior negativity respectively. These are named for their timing, spatial distribution on the scalp, and direction of difference from the baseline. It is important to note that ERP components can potentially cancel and mask each other, and that it is difficult to precisely localize the neural activity that causes the changes in electrical potential at the electrodes where those changes are measured.

3 Related Work

This work is most closely related to the paper from which we get the ERP data: Frank et al. (2015). In that work, the authors relate the surprisal of a word, i.e. the (negative log) probability of the word appearing in its context, to each of the ERP signals we consider here. The authors do not directly train a model to predict ERPs. Instead, models of the probability distribution of each word in context are used to compute a surprisal for each word, which is input into a mixed effects regression along with word frequency, word length,

word position in the sentence, and sentence position in the experiment. The effect of the surprisal is assessed using a likelihood-ratio test. In [Hale et al. \(2018\)](#), the authors take an approach similar to [Frank et al. \(2015\)](#). The authors compare the explanatory power of surprisal (as computed by an LSTM or a Recurrent Neural Network Grammar (RNNG) language model) to a measure of syntactic complexity they call “distance” that counts the number of parser actions in the RNNG language model. The authors find that surprisal (as predicted by the RNNG) and distance are both significant factors in a mixed effects regression which predicts the P600, while the surprisal as computed by an LSTM is not. Unlike [Frank et al. \(2015\)](#) and [Hale et al. \(2018\)](#), we do not use a linking function (e.g. surprisal) to relate a language model to ERPs. We thus lose the interpretability provided by the linking function, but we are able to predict a significant proportion of the variance for all of the ERP components, where prior work could not. We interpret our results through characterization of the ERPs in terms of how they relate to each other and to eye-tracking data rather than through a linking function. The authors in [Wehbe et al. \(2014\)](#) also use a recurrent neural network to predict neural activity directly. In that work the authors predict magnetoencephalography (MEG) activity, a close cousin to EEG, recorded while participants read a chapter of *Harry Potter and the Sorcerers Stone* ([Rowling, 1999](#)). Their approach to characterization of processing at each MEG sensor location is to determine whether it is best predicted by the context vector of the recurrent network (prior to the current word being processed), the embedding of the current word, or the probability of the current word given the context. In future work we also intend to add these types of studies to the ERP predictions.

4 Method

Data. We use two sources of data for this analysis. The primary dataset we use is the ERP data collected and computed by [Frank et al. \(2015\)](#), and we also use behavioral data (eye-tracking data and self-paced reading times) from [Frank et al. \(2013\)](#) which were collected on the same set of 205 sentences. In brief, the sentences were selected from sources using British English with a criterion that they be understandable out of context. We use the ERP component values as com-

puted by [Frank et al. \(2015\)](#) which have been high-pass filtered at 0.5 Hz to reduce correlation between ERP components and modulus transformed ([John and Draper, 1980](#)) to make the distribution of component values more normal. We do not use the 100ms pre-trial baseline which is made available by [Frank et al. \(2015\)](#) and which they use as a separate input to the mixed effects regression. For more information about the ERP datasets and data collection procedures we refer the reader to the original papers. For the behavioral data, we use self-paced reading times and four eye-tracking measures. Self-paced reading time is considered a signal of integration difficulty (i.e. as it becomes more difficult to integrate the meaning of the current word into the context, the amount of time a reader spends on the current word increases). The eye-tracking measures are intended to capture both early effects (effects modulated primarily by properties of the word independent of its context, such as word frequency and word length) and late effects (effects modulated by the context in which the word is found, i.e. comprehension difficulty) in word processing ([Rayner and Pollatsek, 2006](#)). In both cases, the eye-tracking measures provide a signal of overt visual attention, which is thought to strongly correlate with covert perceptual attention in normal reading ([Rayner, 2009](#)). We log-transform the self-paced reading time and the eye-tracking measures.

Model. To predict the ERP signals in the data, we start with a 3-layer bidirectional LSTM-based language model encoder using the architecture found in [Merity et al. \(2017\)](#) and pretrained on the WikiText-103 dataset ([Merity et al., 2016](#)) (we use the pretrained model from [Howard and Ruder \(2018\)](#)). The pretraining objective is to minimize the negative log-likelihood of the next word for the forward LSTM and the previous word for the reverse LSTM. The word-embeddings (input embeddings) in the encoder have 400 components, the hidden layer outputs have 1150 components each, and the context-embeddings output from the encoder have 400 components. The forward-encoder and backward-encoder are independently fine-tuned on the baby version of the British National Corpus ([Consortium, 2005](#)) to help with prediction of British English (both the ERP data and eye-tracking data use British English). During task training the two encoders’ output embeddings are concatenated together and fed into a causal-

convolution layer which combines each pair of adjacent timepoints into a single pair-embedding with 10 components. The causal-convolution (i.e. convolution which is left padded) ensures that the pair-embeddings are aligned so that the prediction targets correspond to the later word in the pair. In other words the pair can be thought of as representing the ‘current’ and ‘previous’ words together. A ReLU is applied to the pair-embedding before it, along with the word length and the log probability of the word, is fed into a linear output layer to predict each ERP and behavioral measure (see Figure 2). The convolution and linear layers are initialized using the default PyTorch (Paszke et al., 2017) initialization, i.e. the initialization proposed in He et al. (2015). The encoder portion of the model includes dropout as applied in Merity et al. (2017), but we use different dropout probabilities when we fit the neural and behavioral data (the dropout probability on the input embeddings was 0.05, 0.4 on the input to the LSTM, 0.4 on LSTM hidden layers, 0.5 on the output of the LSTM, and 0.5 on the recurrent weights). We did not find dropout in the decoder to be helpful. We use the Adam optimizer (Kingma and Ba, 2014) with $\beta_1 = 0.95, \beta_2 = 0.999$ for training and we use mean squared error as the loss.

Procedure. We begin our training procedure by fine-tuning the forward- and backward-encoders independently on the baby version of the British National Corpus (Consortium, 2005). This corpus has British English that may help in modeling the University College London corpus, while not overlapping with it.

After the model fine-tuning, we estimate how well the model predicts each of the ERP signals and eye-tracking measures by training the model 100 times with different train/test splits and decoder parameter initializations. We use 10% of the data for testing and the remainder for training. The sentences in the ERP data are split at random. After we split the data, we compute the mean and standard deviation of each ERP signal (and each eye-tracking measure and the self-paced reading time) within participant on the training data. We use these values to standardize the training data within participant, and then average the data from all of the participants together. After we average, we again compute the mean and standard deviation to standardize the average. We follow a similar procedure for the test data, but we use the mean

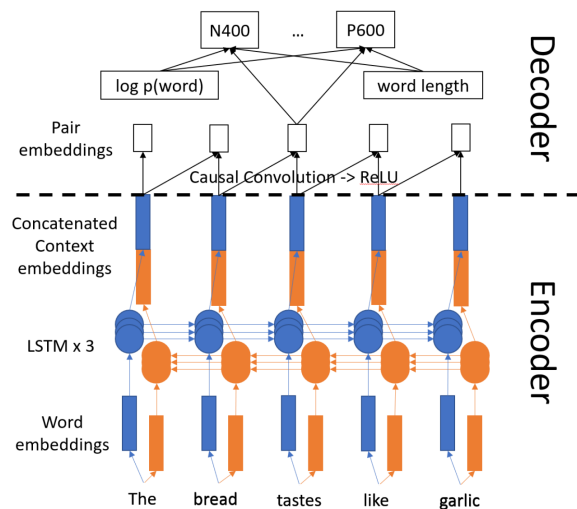


Figure 2: The model uses an encoder based on the architecture and regularization in Merity et al. (2017) and pretrained by Howard and Ruder (2018). Within this architecture 2 independent 3-layer LSTM models encode a sentence. The context-embeddings output from each encoder are then concatenated together to give a single representation to each word in the sentence. These concatenated context-embeddings are fed into a causal-convolution, which learns a function to combine each pair of context-representations into a pair-embedding. A rectified linear unit (ReLU) non-linearity is applied to the pair-embedding, after which independent linear layers map the pair-embedding along with the log-probability of a word and the word-length to a prediction of each ERP or behavioral signal.

and standard deviation from the training data when standardizing. Note that we use the log of the behavior measures, and the log is taken before the data-standardization.

In the loss function (and when we evaluate model performance) we only consider content words. We mark as a content word any word that is an adjective, adverb, auxiliary verb, noun, pronoun, proper noun, or verb (including to-be verbs). All other words are considered function words.

During the first 20 epochs of training, only the parameters of the decoder are modified. Following this, we train the model for an additional 15 epochs during which the parameters of the decoder and the final layer of the encoder (the final LSTM layer in both the forward and backward encoder) can be modified. We also experimented with additional training epochs and allowing all parameters of the model to be modified, but we found that this caused overfitting.

Comparing models trained with different loss functions. To better understand the relationship between ERP signals, and between ERP signals and behavioral data, we train the model with different loss functions that include mean squared error terms corresponding to various combinations of the ERP signals and behavioral data. For example, one of the training variations includes a mean squared error term for the P600 and a mean squared error term for the N400 in the loss, but does not use the other signals during training. In this variation, for a mini-batch of size B , where example b has T_b content tokens and the superscripts p and a denote the predicted and actual values for a measure respectively, the loss function can be written as:

$$\frac{1}{\sum_{b=1}^B T_b} \sum_{b=1}^B \sum_{t=1}^{T_b} (\text{P600}_{b,t}^p - \text{P600}_{b,t}^a)^2 + (\text{N400}_{b,t}^p - \text{N400}_{b,t}^a)^2 \quad (1)$$

For each of the training variations, we repeat the training procedure described above (but fine-tuning the language model on the British National Corpus is done only once). We use a consistent train/test split procedure, such that the split for the i th run of the 100 runs is the same across all training variations, but the split changes between run i and run j . This enables us to use paired statistical testing when we test for significance.

We test for whether the proportion of variance explained (computed as $1 - \frac{\text{MSE}}{\text{variance}}$ on the validation set) on each ERP and behavioral measure is significantly different from 0 using the single sample t-test controlled for false discovery rate using the Benjamini-Hochberg-Yekutieli procedure (Benjamini and Yekutieli, 2001) with a false discovery rate of 0.01.

To test whether the proportion of variance explained is different between different training variations (for example training with just the N400 signal included in the loss vs. training with both the N400 and the LAN included in the loss), we use a paired t-test. We then adjust for the false discovery rate again with a rate of 0.01.

5 Results

All ERP components are predictable. In the original study on this dataset, the investigators found that when surprisal was used as a linking function between the language model and the

mixed effects regression, the only ERP for which the surprisal showed a significant effect in the regression was the N400 (Frank et al., 2015). In contrast, we find that when we directly predict the ERP signals we are able to predict a significant proportion of the variance for all of them (see Table 1).

Joint training benefits ERP component prediction.

To explore the relationship between ERP components, we train $63 = \binom{6}{1} + \binom{6}{2} + \dots + \binom{6}{6}$ different models using all of the possible combinations of which of the six ERP signals are included in the loss function during training. For each of the six ERP components, we look for the best performing models (see Table 1). The N400 is best predicted when the model is trained on that component independently, but every other ERP component prediction can be improved by including a second ERP component in the training. Thus multitask learning has a clear benefit when applied to the ERP data and some information is shared between ERP component predictions via the model parameters. We also note that it is not the case that training with more ERP components is always better, or that the signals which are most correlated benefit each other most (see Appendix A). The relationship between components clearly impacts whether the prediction of one ERP component benefits from the inclusion of others in model training. The results suggest that 8 pairs of ERP signals are related to each other: the LAN is paired with the P600, EPNP, and PNP, the ELAN with the N400, EPNP, PNP, and P600, and the EPNP is paired with the P600. We discuss these relationships in the Discussion section.

In an additional analysis, we modified our training procedure slightly to probe how jointly training on multiple ERP components compares to training individually on each ERP component. In this analysis we compare only training on each ERP component individually to training on all six ERP components together. We also train for a total of 60 epochs (rather than the 35 epochs used elsewhere). During the first 20 epochs we allow only the parameters of the decoder to be modified. During the next 20 epochs, we allow the parameters of the decoder and the final layer of the encoder (i.e. the final recurrent layer) to be modified. During the last 20 epochs, we allow all of the parameters of the model to be modified. The mean squared error for each of the ERP components from this anal-

Target	Additional	POVE	Target	Additional	POVE	Target	Additional	POVE
ELAN		0.20	LAN		0.30	N400		0.26
ELAN	+ EPNP	0.22	LAN	+ EPNP	0.31			
ELAN	+ N400	0.22	LAN	+ PNP	0.32			
ELAN	+ PNP	0.22	LAN	+ P600	0.32			
ELAN	+ P600	0.22	LAN	+ PNP, N400	0.33			
EPNP		0.34	P600		0.27	PNP		0.33
EPNP	+ LAN	0.35	P600	+ EPNP	0.30	PNP	+ LAN	0.36
EPNP	+ GROUP A	0.36	P600	+ LAN	0.30	PNP	+ GROUP B	0.36

Table 1: Proportion of variance explained (POVE) for each of the ERP components (mean of 100 training runs). The second column in each cell shows which ERP components in addition to the target ERP component were included in training. All combinations of training signals were explored. Shown is the best combination for each ERP target as well as every combination which is (i) significantly different from training on the target component alone, (ii) not significantly different from the best training combination, and (iii) uses no more than the number of signals used by the best combination. The N400 is predicted best when only the N400 signal is included in training. All values are significantly different from 0. GROUP A refers to (PNP, ELAN, LAN, P600) and GROUP B refers to (EPNP, ELAN, LAN, P600).

ysis is shown for each epoch in Figure 3. From the loss curves, we make a few observations. First, we see inflection points at epochs 20 and 40, when we allow more parameters of the model to be modified. The first inflection point indicates that allowing the recurrent layer to be modified benefits the prediction, while the second inflection point shows that overfitting becomes more severe if we allow all parameters of the model to be modified. We also see from these curves that part of the benefit of joint training is that it helps reduce overfitting – we see less of a climb in the validation loss after the minimum point in the joint training. Beyond this reduction in overfitting severity, we note that for some of the ERP components (the LAN, EPNP and PNP components) joint training actually gives a better overall minimum in prediction error.

Behavioral data benefits the prediction of ERP components. We are also interested in whether behavioral data can be used to improve ERP prediction since it should signal both the amount of overt attention required at various points in a sentence as well as integration difficulty. To study this question, we again train models using different combinations of training signals that include or do not include the behavioral data predictions in the loss function (see Table 2). We see that self-paced reading time indeed can improve prediction of a target ERP component relative to training on the target ERP component alone by about the same amount as the best combination of ERP

components for all but the N400. Eye-tracking data can also improve the prediction accuracy of the ELAN, P600, and PNP components.

Insensitivity to choice of architecture. One potential concern about our results is the degree to which the relationships we see between ERP components and between ERP components and behavioral data is an artefact of our rather arbitrary choice of network architecture. We partially address this by running the same analysis using (i) only the forward direction of the encoder, and (ii) only the word-embeddings (the input embeddings) and not the context-embeddings (the output embeddings) of the encoder. The proportion of variance explained for each ERP component is lower using these variants of the analysis than using the bidirectional variant (see Appendix A), but qualitatively the relationships are similar. We leave further analysis of the sensitivity of our qualitative results to choice of architecture for future work.

6 Discussion

In this work we find that all six of the ERP components from Frank et al. (2015) can be predicted above chance by a model which has been pre-trained using a language modeling objective and then directly trained to predict the components. This is in contrast to prior work which has successfully linked language models to the N400 (Frank et al., 2015) and P600 (Hale et al., 2018) but not the other ERP components. We also note that con-

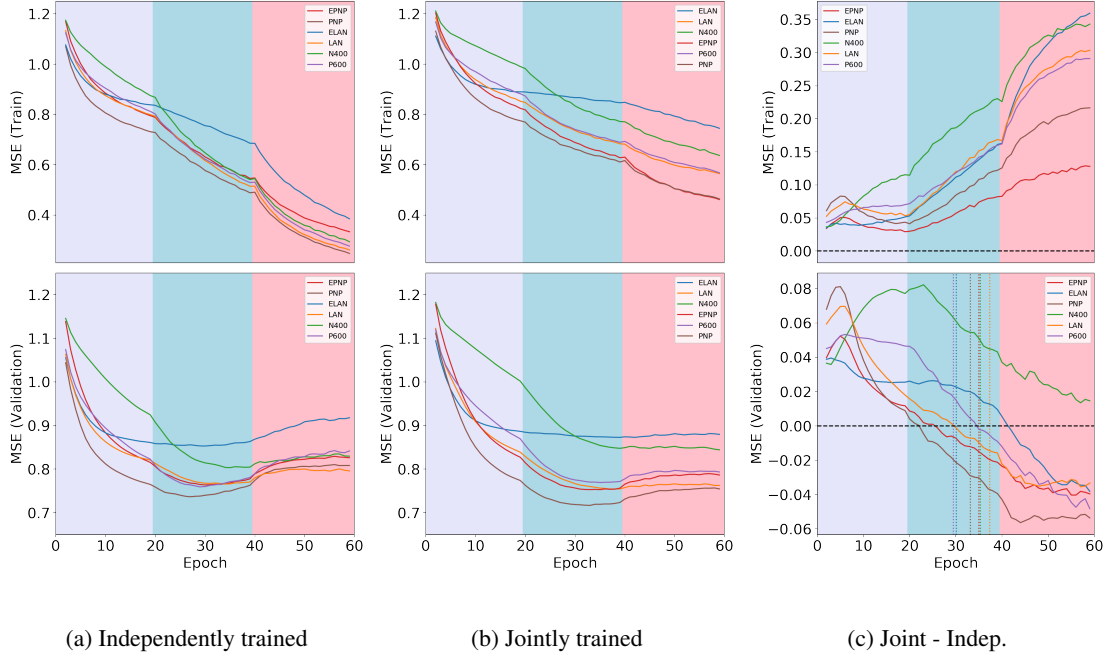


Figure 3: The mean squared error (MSE) for prediction of each of the ERP signals during each epoch of training (mean of 100 training runs). The first 2 epochs have been omitted for clarity. During the first 20 epochs (lavender background), only the decoder parameters are modified. During the next 20 epochs (light blue background), the parameters in the final layer of the encoder are also modified. During the last 20 epochs (pink background), all of the parameters are modified. Note that in this model architecture, information can be shared between ERP signals even when only the decoder is modified. The figure shows the MSE when separate models are trained for each ERP independently (a), the MSE when a single model is trained on all ERPs jointly (b), and the difference between these two scenarios (c). The top row in each column shows the MSE on the training data while the bottom row shows the MSE on the validation data. In the bottom row right, the dotted vertical lines indicate the epoch at which the minimum MSE is reached in the lower of the independent or joint training. The LAN, EPNP, and PNP all show modest benefits from joint training before overfitting sets in (the minimum value occurs in the joint training scenario), while all ERP signals other than the N400 show reduced overfitting in joint training.

Target	Additional	POVE	Target	Additional	POVE	Target	Additional	POVE
ELAN		0.20	LAN		0.30	N400		0.26
ELAN	+ ERP	0.22	LAN	+ ERP	0.33	N400	+ ERP	0.26
ELAN	+ READ	0.22	LAN	+ READ	0.31	N400	+ READ	0.27
ELAN	+ EYE	0.22	LAN	+ EYE	0.30	N400	+ EYE	0.25
EPNP		0.34	P600		0.27	PNP		0.33
EPNP	+ ERP	0.36	P600	+ ERP	0.30	PNP	+ ERP	0.36
EPNP	+ READ	0.35	P600	+ READ	0.29	PNP	+ READ	0.34
EPNP	+ EYE	0.34	P600	+ EYE	0.29	PNP	+ EYE	0.34

Table 2: Proportion of variance explained (POVE) for each of the ERP components (mean of 100 training runs). +ERP indicates the best combination of ERP training signals for the target ERP component, + READ indicates the inclusion of self-paced reading times, +EYE indicates the inclusion of eye-tracking data, and bold font indicates a significant difference from training on the target component alone.

trary to Hale et al. (2018), we find that an LSTM does contain information that can be used to predict EEG data, and in particular that it can predict the P600. We speculate that the analysis used in Hale et al. (2018) did not find reliable effects because the language models were related to the EEG data through functions chosen a priori (the surprisal, and the ‘distance’ metric). These functions, though interpretable, might be interpretable at the cost of losing much of the information in the representations learned by the network.

In addition, we show through our multitask learning analysis that information is shared between ERP components, and between ERP components and behavioral data. Although these relationships must be viewed with caution until they can be verified across multiple datasets and with more variation in neural network architectures, here we consider some potential reasons for our findings. The broad point we wish to make is that by better understanding which ERP components share information with each other and with behavioral data through the type of analysis we present here (multitask learning) or other means, we can better understand what drives each ERP component and in turn the processes involved in human language comprehension.

Relationship between ERPs. Our findings that the LAN and P600 are related, and that the ELAN and P600 are related are expected from both a theoretical perspective and from previous work examining the interactions of ERP components (Gunter et al., 1997; Hagoort et al., 2003a; Hahne and Friederici, 1999; Kutas et al., 2006; Palolahti et al., 2005). Since the ELAN and LAN have been theorized by some to mark word-category (i.e. part-of-speech) or morpho-syntactic (e.g. subject-verb number agreement) violations (Friederici, 2011; Hahne and Friederici, 2002; Hagoort et al., 2003b) and the P600 is considered a marker for syntactic effort (Coulson et al., 1998; Huettig, 2015; Kemmerer, 2014; Kuperberg, 2007; Kuperberg et al., 2003; Van Petten and Luka, 2012), these signals would naturally be related to each other.

The other relationships we find are more surprising. Some researchers have speculated that the LAN and ELAN are markers for working memory demands (King and Kutas, 1995; Kutas et al., 2006), and that indeed these might be part of sustained negativities that are frequently masked by the P600 (Kemmerer, 2014). If we take this view,

then we would expect to find them in the presence of semantic and syntactic complexity, and this might explain why they seem to benefit from joint training with the other ERP component signals (and benefit prediction of other ERP signals with which they are trained). However, it is notable that predictions of the LAN and ELAN do not benefit each other in our analysis, and that the N400 (a marker for semantic complexity) is not benefited by the prediction of any other ERP component. This absence is by no means definitive, but it undermines the argument that all of these relationships can be explained by complexity and working memory demands alone.

The relative isolation of the N400 from other ERP components in our analysis is interesting. If the N400 is a marker for semantic memory retrieval (Kutas and Federmeier, 2011), then it might be expected to be somewhat isolated from the other components, which may involve syntactic processing or later integration effects.

Alternatively, the relationships we find in our analysis might be an artefact of the way the ERPs are operationalized in Frank et al. (2015). Several of the pairings we find overlap spatially and are near to each other in time, so the ERP components might spill over into each other. Further work is required to disambiguate between these possibilities.

Relationship between behavioral data and ERPs. It is reassuring to see that jointly training models to predict behavioral data along with a target ERP component benefits the prediction of the ERP component compared to training on the target ERP component alone. The benefit to prediction in this case cannot be explained as an artefact of how the ERP components are operationalized in the datasets we use for analysis.

Self-paced reading times widely benefit ERP prediction, while eye-tracking data seems to have more limited benefit to just the ELAN, LAN, and PNP ERP components. It’s difficult to know why this might be the case, but perhaps it is not a coincidence that these three ERP components also show up frequently in the pairs of components that benefit from joint training. If indeed the PNP marks semantic role irregularities (Van Petten and Luka, 2012) and the ELAN and LAN mark working memory or look-forward or look-back operations (Kutas et al., 2006), then its possible that eye-movements might be more related to these

types of operations than to general semantic and syntactic complexities marked by other ERP components. Self-paced reading might better capture these generic difficulties. This explanation is highly speculative, and further work is required to determine whether the relationships between the ERP components and behavioral data are consistent across datasets, and if so, what the explanation is for these relationships.

Choice of bidirectional architecture. We emphasize that the neural network architecture we chose for these analyses was motivated primarily by its success on downstream NLP tasks, public availability of pre-trained models and code, and prior work studying how best to fine-tune the model (Howard and Ruder, 2018; Merity et al., 2017). We do not claim that this architecture reflects human processing. We experimented with a forward-only model variant of our analysis, and found that the bidirectional model predicts brain activity better than the forward-only version (see Appendix A). Although the bidirectional model has access to ‘future’ language input, it does not have access to future brain-activity, so the bidirectional model is not ‘cheating’ when it makes predictions. There are at least three possible explanations for why the bidirectional model performs better than the forward-only model. First, it is possible that when a human reads a sentence, he or she predicts the upcoming language input. Under this hypothesis, a model with access to the future language input can do a better job of predicting the current brain activity because the future language is reflected in that brain activity. Second, it is possible that a bidirectional model is simply able to produce better embeddings for each word in the input because it has more context than a forward-only model. For example, the bidirectional model might be (implicitly) better at anaphora resolution given more context. Under this hypothesis, the additional context given to the model partially compensates for its relative deficit of real-world knowledge compared to a human. Where a human can in many cases solve the anaphora resolution problem by using background knowledge and does not need to see the future language input, a model benefits from additional context. Finally, in our setup, the bidirectional model has more parameters than the forward-only model, and the additional degrees of freedom might give the model an advantage in predicting brain activity. Explo-

ration of why the bidirectional model is better than the forward-only model is an interesting question, but it is left to future work. Additionally, as we noted earlier, the qualitative results of our analysis (e.g. how ERP components relate to each other) should be viewed with caution until they are replicated across multiple choices of architecture.

7 Conclusion

We have shown that ERP components can be predicted from neural networks pretrained as language models and fine-tuned to directly predict those components. To the best of our knowledge, prior work has not successfully used statistical models to predict all of these components. Furthermore, we have shown that multitask learning benefits the prediction of ERP components and can suggest how components relate to each other. At present, these joint-training benefit relationships are only suggestive, but if these relationships ultimately lead to insights about what drives each ERP component, then the components become more useful tools for studying human language comprehension. By using multitask learning as a method of characterization, we have found some expected relationships (LAN+P600 and ELAN+P600) and several more surprising relationships. We believe that this is exactly the kind of finding that makes multitask learning an interesting exploratory technique in this area. Additionally, we have shown that information can be shared between heterogeneous types of data (eye-tracking, self-paced reading, and ERP components) in the domain of human language processing prediction, and in particular between behavioral and neural data. Given the small datasets associated with human language processing, using heterogeneous data is a potentially major advantage of a multitask approach. In future work, we will further explore what information is encoded into the model representations when neural and behavioral data are used to train neural networks, and how these representations differ from the representations in a model trained on language alone.

Acknowledgments

We thank our reviewers for their valuable feedback. This work is supported in part by National Institutes of Health grant number U01NS098969.

References

- Yoav Benjamini and Daniel Yekutieli. 2001. The control of the false discovery rate in multiple testing under dependency. *Annals of statistics*, pages 1165–1188.
- British National Consortium. 2005. The BNC Baby. <http://www.natcorp.ox.ac.uk/>.
- Seana Coulson, Jonathan W King, and Marta Kutas. 1998. Expect the unexpected: Event-related brain response to morphosyntactic violations. *Language and cognitive processes*, 13(1):21–58.
- Stefan L Frank, Irene Fernandez Monsalve, Robin L Thompson, and Gabriella Vigliocco. 2013. Reading time data for evaluating broad-coverage models of english sentence processing. *Behavior Research Methods*, 45(4):1182–1190.
- Stefan L Frank, Leun J Otten, Giulia Galli, and Gabriella Vigliocco. 2015. The ERP response to the amount of information conveyed by words in sentences. *Brain and language*, 140:1–11.
- Angela D Friederici. 2011. The brain basis of language processing: from structure to function. *Physiological reviews*, 91(4):1357–1392.
- Thomas C Gunter, Laurie A Stowe, and Gusbertus Mulder. 1997. When syntax meets semantics. *Psychophysiology*, 34(6):660–676.
- Peter Hagoort, Marlies Wassenaar, and Colin Brown. 2003a. Real-time semantic compensation in patients with agrammatic comprehension: Electrophysiological evidence for multiple-route plasticity. *Proceedings of the National Academy of Sciences*, 100(7):4340–4345.
- Peter Hagoort, Marlies Wassenaar, and Colin M Brown. 2003b. Syntax-related erp-effects in dutch. *Cognitive Brain Research*, 16(1):38–50.
- Anja Hahne and Angela D Friederici. 1999. Electrophysiological evidence for two steps in syntactic analysis: Early automatic and late controlled processes. *Journal of cognitive neuroscience*, 11(2):194–205.
- Anja Hahne and Angela D Friederici. 2002. Differential task effects on semantic and syntactic processes as revealed by erps. *Cognitive Brain Research*, 13(3):339–356.
- John Hale, Chris Dyer, Adhiguna Kuncoro, and Jonathan R Brennan. 2018. Finding syntax in human encephalography with beam search. *arXiv preprint arXiv:1806.04127*.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034.
- Gregory Hickok and David Poeppel. 2007. The cortical organization of speech processing. *Nature Reviews Neuroscience*, 8(5):393.
- Jeremy Howard and Sebastian Ruder. 2018. Fine-tuned language models for text classification. *arXiv preprint arXiv:1801.06146*.
- Falk Huettig. 2015. Four central questions about prediction in language processing. *Brain Research*, 1626:118–135.
- JA John and Norman R Draper. 1980. An alternative family of transformations. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 29(2):190–197.
- David Kemmerer. 2014. *Cognitive neuroscience of language*. Psychology Press.
- Jonathan W King and Marta Kutas. 1995. Who did what and when? using word-and clause-level erps to monitor working memory usage in reading. *Journal of cognitive neuroscience*, 7(3):376–395.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Gina R Kuperberg. 2007. Neural mechanisms of language comprehension: Challenges to syntax. *Brain research*, 1146:23–49.
- Gina R Kuperberg, Tatiana Sitnikova, David Caplan, and Phillip J Holcomb. 2003. Electrophysiological distinctions in processing conceptual relationships within simple sentences. *Cognitive Brain Research*, 17(1):117–129.
- Marta Kutas and Kara D Federmeier. 2011. Thirty years and counting: finding meaning in the n400 component of the event-related brain potential (ERP). *Annual review of psychology*, 62:621–647.
- Marta Kutas, Cyma K Van Petten, and Robert Kluender. 2006. Psycholinguistics electrified ii (1994–2005). In *Handbook of Psycholinguistics (Second Edition)*, pages 659–724. Elsevier.
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. 2017. Regularizing and optimizing LSTM language models. *arXiv preprint arXiv:1708.02182*.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2016. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*.
- Maria Palolahti, Sakari Leino, Markus Jokela, Kreetta Kopra, and Petri Paavilainen. 2005. Event-related potentials suggest early interaction between syntax and semantics during on-line sentence comprehension. *Neuroscience Letters*, 384(3):222–227.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in pytorch. In *NIPS-W*.

- Matthew A Lambon Ralph, Karen Sage, Roy W Jones, and Emily J Mayberry. 2010. Coherent concepts are computed in the anterior temporal lobes. *Proceedings of the National Academy of Sciences*, 107(6):2717–2722.
- Keith Rayner. 2009. Eye movements and attention in reading, scene perception, and visual search. *The quarterly journal of experimental psychology*, 62(8):1457–1506.
- Keith Rayner and Alexander Pollatsek. 2006. Eye-movement control in reading. In *Handbook of Psycholinguistics (Second Edition)*, pages 613–657. Elsevier.
- J. K. Rowling. 1999. *Harry Potter and the Sorcerer’s Stone*, 1 edition, volume 1. Scholastic, New York.
- Cyma Van Petten and Barbara J Luka. 2012. Prediction during language comprehension: Benefits, costs, and erp components. *International Journal of Psychophysiology*, 83(2):176–190.
- Leila Wehbe, Ashish Vaswani, Kevin Knight, and Tom Mitchell. 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)*, pages 233–243, Doha, Qatar. Association for Computational Linguistics.

A Appendix

Here we present a visualization (Figure 4) of the results presented in Table 1 of the main paper, and a visualization (Figure 5) of a more complete set of results from which the information in Table 2 of the main paper is drawn. We also show supplemental results for variants of our primary analysis on multitask learning with eye-tracking, self-paced reading time and ERP data. In the variants we modify the input representation to our decoder network to see whether the relationships between the behavioral data and neural activity appear to be consistent with different choices of encoder architectures. Additional (and more varied) choices or architectures are left to future work. The results in Table 3 reflect using only the forward-encoder (rather than the bi-LSTM) in the encoder network, while the results in Table 4 reflect using only the word embeddings (i.e. bypassing the LSTM entirely). While the results are clearly worse for each of these choices of architecture than for using a bi-LSTM encoder, the relationships between the behavioral data and the ERP signals is qualitatively similar. Finally, 5 shows the Pearson correlation coefficient between different measures. We note that the patterns of correlation are different than the patterns of which measures benefit from joint training with each other.




Figure 4: The proportion of variance explained for prediction of each of the ERP signals (mean of 100 training runs). The target ERP is indicated by color; each group of bars shows performance for a different target ERP. The top bar in each group shows the proportion of variance explained when the model is trained using only the target ERP. The bottom bar in each group shows the maximum proportion of variance explained over all combinations of training ERPs (or in the case of the N400, the second best). Also shown in each group are any training combinations that (i) used no more than the number of ERP signals used by the combination that achieved the maximum, and (ii) which were not significantly different from the maximum. Bars are statistically different from each other if a black dot on one bar is connected by a contiguous vertical line to a white dot on the other bar. The bars in the N400 group are not significantly different from each other. The N400 signal is best predicted when the model is trained on just that signal. In every other group, there is at least one ERP that, when combined with the target ERP during training, improves the prediction of the target ERP. The results suggest that these pairs are related: (LAN, P600), (LAN, EPNP), (LAN, PNP), (ELAN, N400), (ELAN, EPNP), (ELAN, PNP), (ELAN, P600), (EPNP, P600).

Target	Additional	POVE	Target	Additional	POVE	Target	Additional	POVE
ELAN		0.20	LAN		0.23	N400		0.20
ELAN	+ERP	0.22	LAN	+ ERP	0.26	N400	+ ERP	0.20
ELAN	+READ	0.22	LAN	+ READ	0.25	N400	+ READ	0.20
ELAN	+EYE	0.21	LAN	+ EYE	0.24	N400	+ EYE	0.18
EPNP		0.28	P600		0.24	PNP		0.28
EPNP	+ ERP	0.28	P600	+ ERP	0.25	PNP	+ ERP	0.31
EPNP	+ READ	0.29	P600	+ READ	0.25	PNP	+ READ	0.30
EPNP	+ EYE	0.29	P600	+ EYE	0.24	PNP	+ EYE	0.29

Table 3: Proportion of variance explained for each of the ERP components when using only the forward direction of the encoder (mean of 100 training runs). +ERP indicates the best combination of ERP training signals for the target ERP component, + READ indicates the inclusion of self-paced reading times, +EYE indicates the inclusion of eye-tracking data, and bold font indicates a significant difference from training on the target component alone.




Figure 5: The proportion of variance explained for prediction of each of the ERP signals (mean of 100 training runs). The target ERP is indicated by color; each group of bars shows performance for a different target ERP. The top bar in each group shows the proportion of variance explained when the model is trained using only the target ERP. Moving down, the next bar in each group, labeled *ERP* shows the proportion of variance explained by the best combination of ERP signals for the target ERP. The other bars in each group moving from top to bottom show training variations that use behavioral data with either just the target ERP, or with the best combination of ERP signals. *READ* denotes self-paced reading data, and *EYE* denotes all four eye-tracking measures (in this analysis we use right-bounded pass time, gaze duration, go-past time, and first-fixation duration). Pairs of bars are significantly different from each other (paired t-test, false discovery rate ≤ 0.01) if a black dot on one bar is connected to a white dot on the other bar by a contiguous vertical line. Self-paced reading time benefits prediction of all target ERP components except the N400. In the case of the ELAN, LAN, and PNP, self-paced reading time also has marginal benefit compared to the best combination of ERP training signals. Eye-tracking data benefits prediction of the ELAN, P600, and PNP components.

Target	Additional	POVE	Target	Additional	POVE	Target	Additional	POVE
ELAN		0.15	LAN		0.17	N400		0.05
ELAN	+ ERP	0.18	LAN	+ ERP	0.19	N400	+ ERP	0.05
ELAN	+ READ	0.18	LAN	+ READ	0.19	N400	+ READ	0.08
ELAN	+ EYE	0.19	LAN	+ EYE	0.19	N400	+ EYE	0.10
EPNP		0.18	P600		0.10	PNP		0.20
EPNP	+ ERP	0.20	P600	+ ERP	0.13	PNP	+ ERP	0.23
EPNP	+ READ	0.20	P600	+ READ	0.13	PNP	+ READ	0.22
EPNP	+ EYE	0.21	P600	+ EYE	0.14	PNP	+ EYE	0.23

Table 4: Proportion of variance explained for each of the ERP components when using only the word embeddings as input to the decoder and bypassing the LSTM entirely (mean of 100 training runs). +ERP indicates the best combination of ERP training signals for the target ERP component, + READ indicates the inclusion of self-paced reading times, +EYE indicates the inclusion of eye-tracking data, and bold font indicates a significant difference from training on the target component alone.

Signal	ELAN	EPNP	LAN	N400	P600	PNP	FIX	PASS	GO	RIGHT	READ
ELAN	1.00	0.27	0.32	0.11	0.10	0.24	0.27	0.26	0.27	0.26	-0.04
EPNP	0.27	1.00	0.66	0.41	0.50	0.83	0.17	0.17	0.19	0.17	0.02
LAN	0.32	0.66	1.00	0.58	0.33	0.47	0.12	0.11	0.13	0.12	0.01
N400	0.11	0.41	0.58	1.00	0.47	0.33	-0.04	-0.04	-0.02	-0.03	0.11
P600	0.10	0.50	0.33	0.47	1.00	0.69	0.14	0.14	0.16	0.14	0.10
PNP	0.24	0.83	0.47	0.33	0.69	1.00	0.25	0.24	0.26	0.25	0.03
FIX	0.27	0.17	0.12	-0.04	0.14	0.25	1.00	1.00	1.00	1.00	0.04
PASS	0.26	0.17	0.11	-0.04	0.14	0.24	1.00	1.00	1.00	1.00	0.04
GO	0.27	0.19	0.13	-0.02	0.16	0.26	1.00	1.00	1.00	1.00	0.05
RIGHT	0.26	0.17	0.12	-0.03	0.14	0.25	1.00	1.00	1.00	1.00	0.04
READ	-0.04	0.02	0.01	0.11	0.10	0.03	0.04	0.04	0.05	0.04	1.00

Table 5: Raw Pearson’s correlation coefficients (computed on content words after the standardization and participant-averaging) between each neural and behavioral measure and each other measure. FIX indicates first-fixation time, PASS indicates first-pass time, GO indicates go-past time, RIGHT indicates right-bounded reading time, and READ indicates self-paced reading. Many of the measures are highly correlated, but the pattern of correlations is different from the pattern of benefits that we find during joint-training. In particular we note that the N400 is correlated with the other ERP signals, and yet we do not see benefit in prediction of the N400 when jointly training a model to predict it and other signals.