

Data-driven models and computational tools for neurolinguistics: a language technology perspective

Ekaterina Artemova¹, Amir Bakarov¹, Aleksey Artemov², Evgeny Burnaev², Maxim Sharaev²

¹ National Research University Higher School of Economics, Moscow, Russia

² Skolkovo Institute of Science and Technology, Moscow, Moskovskaya oblast', Russia

Abstract

In this paper, our focus is the connection and influence of language technologies on the research in neurolinguistics. We present a review of brain imaging-based neurolinguistics studies with a focus on the natural language representations, such as word embeddings and pre-trained language model. Mutual enrichment of neurolinguistics and language technologies leads to development of brain-aware natural language representations. The importance of the research area is emphasized by medical applications

1. Introduction

Contemporary scientific research is unimaginable without powerful data processing tools offered by the computational sciences, in particular, by data science and machine learning. Data-driven approaches, aiming to automatically establish dependencies among measurable quantities using algorithmic and optimization techniques, substantially extend the potential scope of modelled variables (ultimately targeting explanation of *all* experimental observables). Modern neuroscience, in particular neuroimaging studies, has greatly benefitted from incorporating data-driven approaches (Sejnowski et al., 2014; Vu et al., 2018).

Neurolinguistics is a focus for this survey: major goals of this paper are to analyse research projects that address brain mapping, neural decoding, characterization of neural processes, their achievements, and connections to natural language processing.

Decades ago it was thought that a single specific brain region is responsible for performing particular linguistic functions. Now the common idea is that brain connectivity is more important than a distinct brain area activation.

Collecting neuroimaging data (e.g. EEG, fMRI, fNIRS, eye-tracking) during medical neurolinguistics studies helps to shed light on the causes of impairments suffered by language-related brain structures.

Modern findings from both brain mapping and linguistics begin to make predictions about the breakdown in the brain function and their spatial location, responsible for these impairments (Oota et al., 2018). Neurological and psychiatric disorders studied with brain mapping and linguistic approaches range from are Mild Cognitive Impairment (MCI), aphasia (McKinney-Bock

and Bedrick, 2019; Wise et al., 1991), Alzheimer's disease (Pakhomov et al., 2016), focal brain lesions (Troyer et al., 1998), as well as schizophrenia (Robert et al., 1998) and others. In less than two decades of research with imaging and data-driven methods, neurolinguistics has made leaps from early attempts on relating separate words with imaging data (Friederici et al., 2000; Fiebach & Friederici, 2004), to whole-brain mapping from natural stimuli (Huth et al., 2016), to decoding of meaning from brain activation data (Pereira et al., 2018).

Commonly research in NL is conducted with the use of volunteering or paid participants that perform certain tasks; both information about the tasks and the internal state of the participants is recorded; then a predictive model is built, relating the features of the task and the features of the subjects' internal state, and then it is interpreted to draw conclusions about the organization or the functionality of the brain structures

A central topic in neurolinguistics is the study of the relations between features of the presented stimulus and the corresponding brain responses. This should be done for the possibly broader set of stimuli, ultimately arriving at a comprehensive structural and functional brain description that could be put to practical use.

Core problem in neurolinguistics is the construction of appropriate semantic space to represent the stimulus. NLP offers particularly attractive tools for representing spoken and written language.

It is common in natural language processing (NLP) to use distributional semantics model (DSM) as representations of word meaning. The success of DSMs in the variety of NLP tasks inspired the idea of utilizing DSMs for neuroscience studies, including brain mapping. These studies address such questions as how humans process and understand spoken and written language. Participants of such experimental studies are presented with vocal or textual language stimuli while their brain activity is recorded by neuroimaging devices. DSMs models allow to capture patterns in the written language stimuli, and (help to interpret brain activity) to correlate the discovered patterns with the brain activity.

Thus, as we conclude, NLP-based methods are increasingly used in most neuroimaging studies. No overview, however, is available that would cover the use of language technology machinery involved in the construction of models for neurolinguistics. This overview is challenging to conduct in a comprehensive or in a systematic way due to the scope of research in neurolinguistics. In what follows, we perform an attempt to perform this overview.

In this survey, we study the influence of the NLP-based feature representation in the brain mapping task. Our goal is to trace the development of data-driven methods for studying brain's capabilities for representing lexical information, correlating it with the development of natural language processing technologies. We do not focus on neuroscience component, but rather concentrate on the applications.

This paper is structured as follows. Section 2 reviews common experimental setups for brain imaging-based neurolinguistics studies and includes discussion of data acquisition issues, addressed tasks and their respective challenges, and pre-built research environments.

Section 3 introduces the major concepts of distributional semantics as well as three generations of DSMs and Section 4 presents brain mapping methods with DS representations of different generations. Section 5 introduces a recent trend of brain aware DSMs. Section 6 lists a few

medical applications of brain mapping and neurolinguistics. Section 7 concludes with a future outlook.

2. Experimental setup of neurolinguistics experiments

A common neuroimaging modality utilised with the semantic mapping studies is functional magnetic resonance imaging (functional MRI, fMRI), which aims to record whole-brain blood-oxygen-level-dependent (BOLD) responses in subjects while they are exposed to various types of stimuli. BOLD responses are commonly represented as dense 4-D arrays of measured data, where time series of blood flow-related activity in tens of thousands of *voxels* (small areas of size approximately $2 \times 2 \times 2 \text{ mm}^3$) are measured across the brain. These excitations are hypothesized to be elicited primarily by the presented stimulus, with minor background contamination due to respiration, heartbeat, or movement.

Brain mapping aims to build interpretable models for recorded BOLD data using tools common to statistics, machine learning, and other areas of computational data science. When performing such modeling, commonly adopted premises are that

- 1) Particular areas selectively respond to particular content encountered in stimulus, and the stronger the influence, the higher the magnitude of the response. This suggests using simple linear models (e.g., generalized or regularized linear regression) to determine which features of the stimulus lead to a high magnitude signal in each area.
- 2) Each area, represented by a voxel, responds largely independently of the other areas, thus a separate model is needed to fit responses in each cortical voxel.

Thus, one common path would be voxel-wise modelling, aiming to estimate the influence of content-related features found in stimuli on the fMRI measurements captured by the scanner, using linear models. Studies aiming to build brain mapping models differ by (1) the type of stimulus involved, (2) the feature representation chosen to describe the stimulus.

Interestingly, studies involving modern mathematical and computational tools have been recently involving a feedback loop, where an improvement to the computational models employed is sought. Such models must be predictive, i.e. one should be able to approximate unseen brain activations from newly presented linguistic stimuli. Another feature is that such models are hypothesized to be good decoders of meaning from brain data, i.e. one may use them to extract a representation of meaning (and possibly decode it into text), based on the current brain signal.

Performing linguistic experiments with brain imaging is tricky due to many physiological processes, which take place in the brain during speech generation and comprehension. These processes highly depend on the experimental settings, so the majority of studies try to concentrate on a particular effect with strong control for others (Wehbe et al., 2014a). Brain activity is monitored during many kinds of audio and visual stimuli, ranging from more simple visual stimuli as written words, in contrast to a non-word letter string (Salmelin, 2007), or a sentence with expected versus unexpected meaning (Kuperberg, 2007), to more complex audio

and audio-visual stimuli like natural speech comprehension and movie watching (Hanke et al., 2014; Huth et al., 2016).

One of the main research questions is the coordinated work, or integration of so-called brain language network including Broca's and Wernicke's brain as well as complicated information flow between them (Wehbe et al., 2014a). Another complication is the high dimensionality of measured brain activity, the correlational nature of neuroimaging techniques, as well as coarse measurements that neuroimaging methods could provide (Abnar et al., 2019). For example, the dataset by (Wehbe et al., 2014a) consists of the fMRI scans of participants reading a particular text. The story was presented to the participants word by word on a screen, each word was displayed for 0.5 seconds and fMRI scan was taken every 2 seconds, which is a usual temporal resolution for fMRI technique (Abnar et al., 2019).

It is also important to consider physiological delay between the stimulus presentation and brain hemodynamic response (Buckner, 1998) which could be around 4 to 6 seconds. Moreover, as raw fMRI signal is very noisy (Murphy et al., 2013; Sharaev et al., 2018a) and noise could be highly correlated with stimuli presentation, the former should be removed before the analysis (Sharaev et al., 2018b). In a well-known 7T fMRI dataset, see (Hanke et al., 2014) participants listened to a German audio-description of the movie. During fMRI scanning, the MR machine as well as radiofrequency coil emits high and low frequency noise, which could be harmful for speech comprehension. To avoid this effect, the audio signal was processed by a series of filters to remove the frequencies that would have caused acoustic distortions in the headphones (Hanke et al., 2014).

While fMRI has good spatial (~1-3mm) resolution, its big disadvantage is low temporal resolution (around 1-2 seconds) which prevents registration of rapid neuronal activity during speech comprehension and generation. When the stimulus is read, it usually takes around 100ms for the visual input to reach the visual cortex, 50ms to be processed as letter strings in a specialized brain region (Salmelin, 2007) and between 200-500ms, to have its semantic properties been processed (Wehbe et al., 2014b). Magnetoencephalography (MEG) measures small magnetic fields outside the skull evoked by neuronal electrical activity, which is similar to electroencephalography (EEG). Thus, MEG recordings, same as EEG, are directly related to neural activity and have no physiological delay. MEG data has very good temporal resolution (sampling frequency around 1kHz) which is ideal for tracking the fast dynamics of language processing, that's why this technique was used in some neurolinguistics studies (Wehbe et al., 2014b). The drawback of MEG is poor spatial resolution, which prevents precise activation localization in the brain, i.e. brain mapping. The possible way out could be the combination of fMRI and EEG imaging to take advantage of the strengths of the methods (Huster et al., 2012).

Different types of stimulus experiences are involved in the studies: read a sentence, look at separate words, hear a story, natural stimulus eg a movie.

3. Three generations of distributional semantics

Distributional semantics model (DSM) stands for a group of methods which are used to map words from a large vocabulary to vectors. These vectors should consist of real numbers, have few zeros and be of relatively small dimensionality: it is common to construct 300-dimensional vectors. Usually such word vectors are addressed as word embeddings. These vectors are treated as mathematical (algebraic) objects: not only similarity (or distance) between them can be computed, but also they can be added together or subtracted. At the core of numerous methods for word embedding construction, is the distributional hypothesis: words that occur in the same contexts tend to have similar meanings (Harris 1954). Word embedding models are trained on large text corpora. They aim at finding words that share contexts and represent them with such vectors that would be close, according to a mathematical similarity measure. For example¹, the embeddings of such words as *coffee* and *tea* should have a high similarity degree, since they are used in a similar way, along with the words *to drink*, *cup*, *to pour* etc. What is more, advanced word embedding models allow to conduct arithmetical operations: *coffee to morning = tea to evening* ; *Madrid to Spain = Berlin to Germany*. Of course, these associations are corpus-specific and may not be present in other models.

From a historical point of view, there are three generations of DSMs: count-based DSMs, distributed or prediction-based DSMs and deep contextualized DSMs. Count-based DSMs arose in the early 90's, followed in the 2010's by prediction-based models. The recent breakthrough of deep contextualized DSMs leverages transfer learning techniques for downstream tasks.

Count-based DSMs build word representations by counting usage of words in different contexts. The basic idea of such algorithms is that given a sufficiently large number of texts, one could create a vocabulary of all unique words and count contexts of each word. By counting these contexts one can represent each word through a real valued vector.

To capture co-occurrence the so-called word-matrix is introduced. In this matrix, each unique word in a corpus is associated with an array of values. Each cell in this array is associated with another word that could be encountered in the context of this word, so basically the number of values is equal to the number of unique words in the corpus. Then, the value in each cell reports degree of closeness of word *a* to word *b* from the context perspective. In the most simple case, it could be a binary value that says whether this word was encountered in the context of the considered word, or not. It also could be a number of encounters of this word, or more complicated measures that should more precisely report "co-occurrence score" (for example, Positive Pointwise Mutual Information). As a result, these actually count the word contexts, so they are referred to as count-based DSMs.

¹ The examples are provided by RusVectors (<https://rusvectors.org>).

The second generation of DSMs are prediction-based models. These models leverage language modelling approach by either predicting the current word from a context, or by predicting the context of the given word. Word2vec (Mikolov et. al, 2013) is a family of neural network architectures, namely CBOW and skip-gram. See Fig. 1 for word2vec configurations.

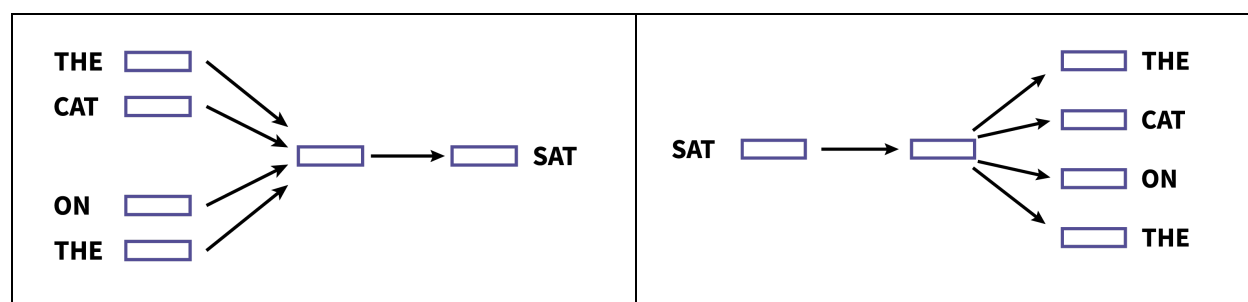


Fig 1. Word2vec configurations. (on the left): continuous bag of words predicts a word based on its context, (on the right): skip-gram predicts the context of the given words.

A word2vec model should be trained on a large text corpus. First, this corpus is preprocessed, so that each text is split in overlapping n-grams by a sliding window. Next, a neural network is trained using one of the word2vec objectives, CBOW or skip-gram. The weights of the hidden layer are the actual learned word embeddings. GloVe (Pennington et al. 2014) is a log-bilinear regression model, trained on aggregated global word-word co-occurrence statistics from a corpus.

Word embeddings can be manipulated as any mathematical object: not only it is possible to calculate a similarity between them, but also to sum them up or to subtract them. Word embeddings can be seen as shallow representation of the grammar and semantics. When trained on a large general corpus, such as Wikipedia, word embeddings capture semantic and grammar patterns, that can be revealed when measuring word similarity. The closer the words are by lexical meaning, the closer the corresponding word embeddings should be. A similarity function of choice is usually a cosine function.

The linguistic interpretation of this similarity measure is obscure, but is always affected by the parameters of the model. One of the most significant parameters is the training corpus, and the impact of training corpus on interpretation of vectors similarity is still an open question. In most cases this similarity reports association (or co-hyponymy² of two words), in some cases it reports relatedness³. The embedding model can be fine-tuned to express a certain type of

² Such word pairs as “blue” and “red” or “tea” and “coffee” are co-hyponyms

³ Such word pairs as “car” and “gasoline” or “cat” and “striped” are related

similarity, such as meronymy⁴ or hyponymy⁵. This approach to fine-tuning is referred as retrofitting.

Word embedding models often fail when faced with such complex language phenomena as polysemy or hyponymy. Although word embeddings are exceptionally powerful for finding words that share a similar meaning, they often mistake for words that have opposite meanings, as they occur in similar contexts. Word embedding models suffer from polysemy and homonymy. Such words as “*bank*” or “*apple*” get a single vector, despite having multiple sense. A few models, such as AdaGram (Bartunov et al. 2016) or SenseGram (Pelevina et al. 2016), try to overcome this issue by simultaneous word sense disambiguation, and word embedding training.

A natural extension of the idea of word embeddings are sentence embeddings. Sentence embeddings are massively used in those applications, which require modeling of sentence similarity, for example, paraphrase identification, question answering and others. Paragraph2vec (frequently addressed as doc2vec, (Le and Mikolov, 2014)) follows the architecture of word2vec to build a sentence embedding. Skip-thoughts (Kiros et al., 2015) model is similar to skip-gram, too. It consists of one encoder, that encodes a sentence, and two decoders, such that one is used to generate the previous sentence and the other is used to generate the next sentence. A skip-thought like task combined with a conversational input-response and multiple classification tasks are used to train Universal Sentence Encoder (Cer et al., 2018) by Google. Unlike skip-thought, which is build from LSTMs, USE exploits deep averaging networks and transformer-like encoders. The majority of the mentioned sentence embedding models suffer from several disadvantages: they are slow to train, give poor results or are difficult to use in real-life settings.

However current pre-trained language models (LMs) are a much more efficient solution to all these issues, as they search for context-dependent word embeddings, which are core to transfer learning paradigm. The transfer learning paradigm is based on the idea of pre-training of a large language model on large corpora and further fine-tuning of these models for downstream tasks. In NLP one can think that during the pre-training stage the model learns the language. During the fine-tuning stage the model is learned to conduct specific tasks. A new generation of DSMs comes with large pre-trained LMs, and they start to make a difference in every natural language processing task. Instead of using every time pre-trained word embeddings and training the whole model from scratch, now a pre-trained model is fine tuned for downstream tasks. Word embeddings are an imperfect way to store language representation, which suffered from language ambiguity. Pre-trained models are less prone to polysemy and antonymy and are able to handle multilinguality at the same time.

⁴ “Tire” is a meronym of “car”. Meronymy stands for “part of” relation.

⁵ “Grey” is a hyponym of “color”. Hy,onymy stands for “is-a” relation. The inverse relation is hypernymy.

Inside the vast majority of transfer learning models are transformer layers, that are more advanced from a technical point of view when compared to other layers. The architecture of transfer learning models is sophisticated, enumerates millions of parameters and take weeks to be pre-trained. Even though pre-training of a large model is expensive and time-consuming, new models appear almost every month as of late 2019. Among others, ELMo (Peters et al., 2018) and BERT by Google (Devlin et al. 2019) are the most popular models. In contrast to other models, ELMo is based on recurrent neural networks, while BERT and other models are transformer-derived. BERT's successors, ALBERT (Lan et al., 2019), RoBERTa (Liu et al., 2019), and other transformer-based models, Transformer-XL (Dai et al., 2019), XLNet (Yang et al., 2019), T5 (Raffel et al., 2019), released by Facebook, Microsoft and other technological companies, are larger and outperform BERT by a significant margin. ELMo, BERT and other models are trained in an unsupervised way by using language modelling objective or similar objectives. ELMo exploits language modelling straightforwardly, while BERT's objective is combined from masked language modelling and next sentence prediction. The latter is massively criticized and omitted in succeeding models, such as RoBERTa. Transformer-XL is specifically designed to process long sequences. Transfer learning models can be exploited as a standalone sentence embedding tool by pooling contextual word embeddings, derived from the model.

4: Brain mapping with DSMs

4.1 Count-bases DSMs

In the pioneering work by Mitchell et al. 2008 for the first time the task to relate language stimuli with the recorded brain activity was introduced. In contrast to primarily descriptive theories this work presents a first computational model to predict spatial patterns of functional magnetic resonance imaging (fMRI) neural activations associated with previously unseen words. The language stimuli consists of 60 word-picture pairs, the words are concrete nouns and belong to twelve semantic categories. Nine healthy participants took part in the study, To predict

activation at voxel y_v a linear model is used: $y_v = \sum_{v_i} c_{v_i} f(w_i)$, where c_{v_i} are trainable weights

and f are word vectors. The word vectors are build from word-context co-occurrence statistics. Two types of contexts are defined: either frequent words according to a large corpus, or the list of manually selected action and sensory-motor verbs. The model is trained in a leave-one-out fashion: all but two words are used for learning model weights, two words are used to test the model precision. The held out words are chosen either from unrelated categories or from the same category. The verb-based context resulted in a model that significantly outperforms models, based on other contexts. Distinguishing between words within the same category appears to be a more challenging task than predicting words from unrelated categories. The results of the paper show a direct, predictive relationship between the statistics of word

co-occurrence and the neural activation associated with thinking about word meaning, and establish a whole new line of research.

4.2 Word embeddings

Murphy et al. (2012) seems to be one of the first papers that used DSMs based on word embeddings for brain mapping. They used the dataset from Mitchel et al. (2008), containing fMRI data from 9 participants with stimuli being line-drawings accompanied by their text labels. The authors compared semantic representations that provided state-of-the-art (SOTA) performance on the brain mapping tasks with new distributional models of semantics that can be derived from arbitrary corpora. The SOTA models used were hand-crafted, based on manual annotation, use of domain-appropriate corpus, etc. As competing corpus-derived models the authors considered word-region co-occurrences, word-collocate features including raw tokens, POS tags, etc. To preprocess data and extract features they used a common pipeline including a co-occurrence frequency cutoff, a frequency normalization weighting, and dimensionality reduction with SVD. As a result the authors claim that they were able to achieve brain activity prediction accuracies comparable to SOTA or slightly higher using unsupervised data-driven models of semantics based on a large random sample of web-text.

In a similar way Sogaard (2016) correlates word embeddings against fMRI data and eye-tracking data. He studied the correlation between distances in word embedding space and fMRI/eye-tracking space, as well as evaluated the squared error of a regression model trained to associate the input word embeddings with fMRI/eye-tracking data. To get word embeddings he used a discriminating and non-probabilistic approaches from Collobert et al. (2011), where the authors to construct embeddings in each training update read an n-gram from a corpus, corrupt it with a random word from a vocabulary and learns a scoring function to score the original phrase lower than the corrupted one. Experimental results confirmed significant statistical dependency and proved that correlating with human word processing data is the valid way for evaluation of word embeddings.

Ettinger et al. (2016) modeled the N400, which is a component of time-locked EEG signals known as event-related potentials (ERP). It is a negative-going deflection that peaks around 400 milliseconds post-stimulus onset, and is a part of the normal brain response to meaningful stimuli, in particular, to words. Ettinger et al. (2016) used vector space word2vec model as explicit models of word relations to predict N400 amplitude in an ERP study. The authors investigated N400 effects in contexts that predict a particular word, and are then completed by words with varying levels of similarity to that predicted word. The already known effect is that unpredicted (zero-cloze) words elicit larger N400s, and when the unpredicted item is similar to the predicted word in strongly predictive contexts, the N400 amplitude is reduced. Ettinger et al. (2016) found that by representing words in a vector space, averaging vectors of informative context words, and taking cosine similarity measures between the averaged context vector and each of its possible completions, they were able to simulate key observations about N400 amplitude behaviour.

Abnar et al. (2018) modeled brain activity using fMRI data from Mitchell et al. (2008) and different types of word embeddings: experiential, distributional and dependency-based. The main aim was to estimate their usefulness for predicting the neural activation patterns associated with concrete nouns. Here in case of experiential word embeddings 65 features are defined and crowdsourcing is used to rate the relatedness of each feature for each word; as distributional word embeddings they used word2vec, fasttext, dependency-based word2vec, GloVe and LexVec approaches based on ideas similar to ones considered in Section 3 above; dependency-based word embeddings contains 25 verb features and non-distributional word vector representations, where words are represented either by the co-occurrence counts of nouns with these selected set of verbs or by a binary vector, indicating whether the represented word has or does not have a specific feature. It turned out that for predicting neural activation patterns in human brains dependency-based word2vec, GloVe and 25 features model are the top-ranked models for at least one of the subjects, and the non-distributional word representation model has the lowest performance. At the same time when predicting the word representation given a brain activation they found that for the 25 features model, the accuracy on the reversed task is much lower. On the other hand, it seems that it is very easy to construct GloVe word vectors from brain activation patterns. Hence, obtaining a high accuracy in the word prediction task using GloVe, supports the fact that the context of the words have a major role in the way we learn the meanings of the words. Besides that the authors also found out that different models makes different kinds of mistakes, and identified the most predictable voxels in the brain for each word embedding model.

Rodrigues et al. (2018) went further in their research and instead of using feature-based models or semantic spaces (aka word embeddings) for words representation they utilized a semantic network to encode lexical semantics. Here the semantic network was used as the base repository of lexical semantic knowledge, namely WordNet. The authors then generated semantic space embeddings from semantic networks, and use it to obtain WordNet embeddings: the intuition is that the larger the number of paths and the shorter paths connecting any two nodes (words) in a network the stronger is their semantic association. The proposed model turned out to demonstrate a competitive performance as its accuracy is comparable to the results obtained with SOTA models based on corpus-based word embeddings reported in previous papers. Although for one third of the 9 subjects this model surpasses Mitchell et al. (2008), on average it did not outperform their model based on hand-selected features.

In their seminal paper Huth et al. (2016) consider another type of stimuli, namely, they model brain responses from naturally spoken narrative stories that contain many different semantic domains and map semantic selectivity across the cortex using voxel-wise modelling of functional MRI (fMRI) data. They used a word embedding space to identify semantic features of each word in the stories constructed by computing the normalized co-occurrence between each word and a set of 985 common English words across a large corpus of English text. Constructed fit models allow to determine which specific semantic domains are represented in each voxel. Due to high dimensionality examining each voxel separately is unfeasible, and so the authors reduced the space of the models by principal component analysis and finally identified four shared dimensions that capture shared aspects of semantic tuning across the subjects and provide a

way to summarize succinctly the semantic selectivity of a voxel. To visualize the semantic space, the authors projected the 10,470 words in the stories from the word embedding space onto each dimension and used k-means clustering to identify 12 distinct categories. Thanks to this the authors visualized the pattern of semantic-domain selectivity across the entire cortex by projecting voxel-wise models onto the shared semantic dimensions. In this way for each voxel it is possible to say to which categories in the semantic space greater observed BOLD responses correspond to. Using specially designed probabilistic model and the constructed semantic labelling of voxels in response to stimuli, the authors created a single atlas that describes the distribution of semantically selective functional areas in human cerebral cortex.

As we discussed above the previous approaches mapped word-level semantic representations using embedding vectors and neglected the effect of context by assuming that the response to each word is independent. On the same dataset as Huth et al. (2016) unlike to these approaches Jain et al. (2018) used representations discovered by a long-short term memory network (LSTM) to incorporate context into encoding language models that predict fMRI responses to natural, narrative speech. These contextual models perform significantly better at predicting brain responses than previously published word embedding models. Moreover, Jain et al. (2018) examined which brain areas can be best modelled by which context length and layers of the LSTM encoding model. It turned out that voxels in low-level language areas (AC) prefer short context, while voxels in higher-level language areas prefer long context. As for layers the authors showed that there is little difference in performance from different layers, although layers 1 and 3 are better suited to model AC, layer 2 better models higher semantic regions. The results illustrates a temporal hierarchy in how the human language is processed. The studies discussed above are limited in the sense that they use relatively small and/or constrained sets of stimuli. Moreover, often the analysis is based on semantic features limited to concrete nouns that does not easily scale to a typical vocabulary of tens of thousands of words. As a result an important challenge arises whether the models would generalise to meanings beyond the limited scope of the training sample. Pereira et al. (2018) proposed an answer to this challenging question. They developed a system that would work on imaging data collected while a subject reads naturalistic linguistic stimuli on potentially any topic. To represent the semantic space they used 300-dimensional GloVe vectors and spectral clustering to group words into 200 regions almost all of which were intuitively interpretable. The authors hand-selected representative words from each of the regions and used them in creating the stimuli. Based on collected data Pereira et al. (2018) constructed a universal brain decoder that can infer the meanings of words, phrases, or sentences from patterns of brain activation data. In the first set of experiments they evaluated a capability of single concept decoding. To that the subject is thinking about the relevant meaning of each words, first, the target word was presented in the context of a sentence that made the relevant meaning salient, second, the target word was presented with a picture that depicted some aspect(s) of the relevant meaning, and third, the target word was presented in a “cloud”, surrounded by five representative words from the semantic cluster. To evaluate decoding performance Pereira et al. (2018) considered a pairwise classification approach and a rank accuracy classification task. In the former they computed and compared the similarity between the decoded vectors and the “true” (text-derived) semantic vectors. In the latter they compared each decoded vector to all 180

text-derived semantic vectors, and ranked them by their similarity. In both cases classification accuracy turned out to be around 0.7 or more. In other experiments on sentence decoding the decoder was applied to the brain images for the sentences from yielding a semantic vector for each sentence. A text-derived semantic vector for each sentence was created by averaging the corresponding word embeddings vectors. To evaluate accuracy Pereira et al. (2018) considered a pairwise classification task to recognise different topics, different passages within the same topic and different sentences from the same passage for all possible pairs in each task. It turned out that it is possible to distinguish sentences at all levels of granularity. As a result this work clearly demonstrates that it is possible to extract a linguistic meaning from imaging data. As we see prediction of semantic meaning of words based on fMRI data prevails. So it is not surprising that Fyshe et al. (2019) recently concentrated on MEG data and analysed semantic representations of adjective-noun phrases. For that they presented phrases consisting of an adjective followed by a noun. Besides that some oddballs were demonstrated, i.e. adjective-adjective pairs. The word vector model used was built using a large corpus of text from web pages. The probability of seeing two words in a particular dependency relationship in a sentence was calculated. To get word embeddings the authors compressed the matrix with the probabilities using singular value decomposition. The first 100 dimensions of the matrix were used for this study. Thanks to the experimental setup, there is a strong correlation between adjectives and nouns in the collected data. Two tasks of predicting the stimulus from the MEG recording were considered: 1) predict adjective semantics: predict the identity of the first word in the phrase (any of the six adjectives); 2) predict noun semantics: predict the identity of the noun (one of eight). Their analysis revealed two novel findings: (a) a neural representation of the adjective is present during noun presentation, but this representation is different from that observed during adjective presentation and (b) the neural representation of adjective semantics observed during adjective reading is reactivated after phrase reading, with remarkable consistency.

4.3 Pretrained language models

Schwartz and Mitchell, 2019 utilize ULMFiT (Howard and Ruder, 2018), a bidirectional multi layer LSTM, to predict several time-locked stereotyped EEG responses, known as event-related potentials, ERPs to word presentations, which are supposed to be markers for semantic or syntactic processes that take place during comprehension. In this study the ERP data comes from the project of (Frank et al., 2019). The context-aware word embeddings from each layer of the architecture are concatenated to achieve a single representation to each word in the sentence, fed into a convolution layer and a non-linear layer. The resulting word embeddings are combined with the log-probability of a word and the word-length to a prediction of each ERP. A single loss might be used to train the whole model to predict six ERP signals (or any subset of them) simultaneously. Such multitask learning allows to find relationships between ERP components: the prediction of one ERP component may benefit from inclusion of another ERP component or multiple components. Shwartz and Mirchel, 2019 claim that their 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. The ablation study shows that the model benefits significantly from the usage of bidirectional architecture. A possible explanation for this may be that the context-aware embeddings benefit from seeing the future language input. Clearly, the model has no access to future brain-activity, so there is no overfitting to the training data. What is more, the bidirectional model has advantageously more parameters, which helps to gain more knowledge from the training data.

Hollestein et al., 2019 present a new framework, Cognival, and conduct an extensive comparison of how different types of DSMs, such as word embeddings and pre-trained language models, fit to various datasets of eye-tracking, EEG and fMRI. They define a single model, namely, neural architecture regression network with one hidden layer, which is used multiple times to predict brain activity from DSMs. The word embeddings under consideration are GloVe, word2vec, fasttext, as well as a less known modification of word2vec, wordnet2vec (Bartusiak et al., 2019), which utilizes WordNet instead of raw textual data. These are compared to two pre-pretrained LMs, namely, BERT and ELMo. The results show that BERT, ELMo and FastText embeddings achieve the best prediction results for predicting all types of brain activity. Pre-trained LMs, BERT and ELMo, significantly outperform word embeddings. This supports the general trend in NLP. There is an evident correlation between predictions of eye-tracking, EEG and fMRI, which supports the idea that the DSMs predict actually brain activity signals and not overfit to preprocessing artifacts of each modality.

In contrast to other research projects, Toneva and Wehbe, 2019 do not predict the brain activity from DSMs, but rather use the brain activity recordings as a proxy for interpreting neural DSMs. The fMRI and MEG data published by Wehbe et al. (2014b) is used as a source for information for interpreting four pre-trained LMs, ELMo, USE, BERT and Transformer-XL. The authors train a linear model with a ridge penalty to predict brain activity from a sentence embedding and show that the middle layers of all DSMs outperform lower and upper layers at the same time. The performance of ELMo, BERT and Transformer-XL is almost the same, with USE having lower performance. Transformer-XL appears to benefit from larger size of contexts, as it is the only DSM specially designed for processing long sequences.

As we can see, there exists various setups of experiments on brain mapping, utilizing different data modalities and types of unsupervised word/sentence embeddings. Obtained results clearly demonstrate the possibility of constructing a brain decoder that is capable to decode semantics from brain activation data.

5. Brain-aware embeddings

Although the multimodal word and sentence embeddings are demanded in NLP, little progress has been achieved in the development of brain-aware embeddings. At the core of the concept of brain aware embeddings is the following hypothesis: if brain activation data encodes semantics, inclusion of brain data in a model of semantics results in a model more consistent

with semantic ground truth (Fyshe et al., 2014). A first attempt to build a brain-aware DSM was made by (Rustandi et al., 2009). Canonical Correlation Analysis (CCA) was applied to correlate corpus-based and brain-based statistics. However, as CCA requires that the data sources to be paired, a vast majority of the corpus-based statistics was lost, thus the benefits of CCA were very limited.

Another of the first models of brain-aware embeddings is introduced by (Fyshe et al., 2014). The proposed model, Joint Non-Negative Sparse Embedding (JNNSE), utilizes matrix factorization techniques to build a hybrid count-based DSM. JNNSE solves the following

objective: $\operatorname{argmin}_{A, D^c, D^b} \sum \|X - A D^c\| + \sum \|Y - A D^b\| + \lambda \|A\|_1$, where X, Y are corpus-based

and brain-based word matrices. X represents standard word-context co-occurrence, while Y is constructed from feature vectors derived from brain activity recording. The JNNSE task is solved with respect matrices A, D^c, D^b being non-negative. The results show that even a small available amount of brain data has a positive impact on the embedding, as they are better in predicting fMRI activations, than text-based embeddings.

Schwartz et al., 2019 have recently introduced the first model specifically designed to capture the way the brain represents language meaning. They fine-tuned the BERT model (Devlin et al., 2018) to predict recordings of brain activity of people, while reading a text. The resulting representations are encoding more brain-activity-relevant information and thus improve the quality of the brain activity prediction. To make sure that the fine-tuned BERT is not overfitted for the task, the authors of the paper conduct a series of transfer learning experiments and show the following. First, the changes to language representation does not harm the performance on downstream NLP tasks. Second, the model is capable of generalizing to new participants. Third, the changes introduced to the model are not an artifact of imaging modality, as the fine-tuned representations learned from both MEG and fMRI are better for predicting fMRI than predicting from fMRI-based representations alone.

To encode the brain activity information, the common procedure for fine-tuning BERT is used: a simple linear layer is added on top of the pre-trained model. This linear layer is trained to map the output embeddings from the base architecture to the fMRI prediction task. The added linear layer can be either trained alone or the whole model can be fine-tuned in the end-to-end fashion. Two datasets of MEG and fMRI experiments (Wehbe et al. 2014a, Wehbe et al. 2014b) were used to fine-tune the pre-trained BERT model by (Devlin et al., 2018).

As fMRI and MEG recordings have different time resolution, they are used differently to supervise the fine-tuning procedure. To predict fMRI, the pooled embedding of the sentence is used, while the MEG recordings are used in a sequence labelling fashion, so that each word in the input sequence is mapped to a certain MEG recording.

To demonstrate, whether the model can be biased towards brain-relevant language information, several models are trained. A vanilla model, which serves as a baseline, is trained for each experiment participant in a such way, that only an added linear layer is updated and the core layers of the model remain unchanged. In contrast to vanilla model, the fully joint model is

trained to simultaneously predict all participants' data from both fMRI and MEG experiments. The results of the paper show that the joint model performs better than vanilla model. This means that the model poses the ability to generalize across participants and imaging modalities. Although the performance of the joint model on such NLP tasks as GLUE benchmark does not differ much from the initial BERT model. It remains unclear how exactly the language representations were changed during the fine-tuning procedure. Finally, this study demonstrated a first step towards brain activity aware language models.

5. Medical applications

Many approaches taken from neurolinguistics, DSM in particular, found applications in medical studies of neurological and psychiatric disorders, like Mild Cognitive Impairment (MCI), aphasia and others. Even without sophisticated neuroimaging techniques, word embeddings models could help in clinical diagnostics and prognosis tasks. Many studies confirm that Semantic Verbal Fluency (SVF, a task where patients should recall as many words from a particular semantic category as they can in a short period of time) is helpful for diagnostics of such brain pathologies as Alzheimer's disease (Pakhomov et al., 2016), MCI, focal brain lesions (Troyer et al., 1998), as well as schizophrenia (Robert et al., 1998). Many studies explored a machine learning approach to automatically detect cognitive impairment based on text/speech samples from ill subjects and healthy controls. In a paper (Linz et al., 2017) the authors compared the classification performance of traditional taxonomic models (Troyer et al., 1997) with word2vec model for MCI detection on a sample of 100 subjects and their results were clearly in favour of the DSM. Another work (Fraser et al., 2019) also suggests that the fully automated cluster models are better able to capture patterns in the data which distinguish the MCI patients from healthy controls. In this study authors extract information units using topic models trained on word embeddings in monolingual and multilingual spaces, and find that the multilingual approach leads to significantly better classification accuracies than training on the target language alone. This might be due to the fact that multilingual data enriches the vocabulary in the relevant areas of the semantic space and topic modelling approach helps to better identify this space (Fraser et al., 2019). Another example is clinical assessment of aphasia using a confrontation naming task, where a stimulus is shown to a subject and she must name it (McKinney-Bock and Bedrick, 2019). This procedure is very time demanding and need automation in assessing semantic similarity. The authors in (McKinney-bock and Bedrick, 2019) as well as (Fergadiotis et al., 2016) showed that word embeddings can be successfully applied to classification of aphasias when assessed by the naming task.

When a person performs either speech generation or comprehension, it is always connected networks of different brain regions working together to perform such a task. For example, the fronto-temporal language system has been shown to respond very specifically in fMRI to linguistic stimuli, but not to other perception or cognitive tasks (Ivanova et al., 2019). One question is whether the language system is engaged during the processing of words meaning

and is necessary for it. In the paper (Ivanova et al., 2019) fMRI study on patients with brain damage was combined with neurolinguistics approach - distributed semantic representations created through extensive training on massive language corpora. The authors found that language brain networks in healthy controls are active during the semantic plausibility task when subjects were shown both sentences and pictures. In patients with global aphasia (those having language areas severely damaged), performance on the picture plausibility task was not decreased, suggesting that the language network is not required for understanding and accessing semantic information in the brain about visually presented events. Much more information about possible mechanisms of brain disorders could be captured when analyzing the same semantic models with different brain imaging modalities. A fresh example is in the paper (Hollenstein et al., 2019), where the authors present the first multimodal framework for evaluating English word representations based on cognitive lexical semantics. They investigated six types of word embeddings by fitting them to 15 datasets of eye-tracking, EEG and fMRI data collected during language processing tasks. A remarkable result here is that the authors found correlations in the embedding performance between different cognitive datasets (i.e. different tasks and different neuroimaging modalities) and their performance on extrinsic NLP tasks. Moreover, the authors developed a framework to automate these comparisons. Thus, neurolinguistic experiments with neuroimaging data collection could give insights to applied NLP tasks and new word embedding creation and evaluation.

Linguistic models might also give insights on aphasia treatment. The authors in (Roelofs, 2019) explore the method called phonological cueing to increase the word retrieval (or proper word finding) performance in patients with aphasia. Picture naming performance of the WEAVER computer model (Roelofs, 1997) was assessed in simulations of immediate and treatment effects of phonological cueing in post-stroke aphasia. The authors claim that the linguistic model successfully simulated the observed effects in naming performance and neural measures. The model could expand our understanding of word finding, associated difficulties, and their improvement by therapy (Roelofs, 2019).

The field of neurolinguistics and semantic models can also help the area, closely related to brain mapping Direct-Speech Brain-Computer-Interface (DS-BCI) research on methods for targeting the imagined speech content for synthesis (Cooney et al., 2018). DS-BCI is aimed at acquiring neural signals (mostly EEG and EcoG due to high temporal resolution) corresponding to imagined speech in order to decode these signals and produce a linguistic output in the form of phonemes, words, or sentences (Cooney et al., 2018) to help disabled people and thus could be considered as an inverse problem of brain mapping. Recent studies have shown the potential of neurolinguistics to improve the ability of DS-BCI to decode imaginary speech from acquired signals with the inclusion of semantics in experimental procedures.

6. Conclusion

The rapid growth of interest towards neurolinguistics research across scientific teams led to an increase in the number of private multimodal datasets with different imaging modalities, different linguistic tasks and types of stimuli. This wide variety of datasets makes it difficult to compare results from different teams or re-use these data in new brain mapping studies. The possible way out for prospective studies could be to standardize (to some extent) the imaging procedure and at least most common linguistic tasks, as it was successfully done previously in the field of machine learning, for instance, in well-known image classification databases, or more general brain mapping experiments, see Human Connectome Project (Van Essen et al. 2013) and similar. In order to make comparisons across studies and perform meta-analysis, we should elaborate on some common performance (or accuracy) measures for brain mapping experiments.

As already mentioned, up to now there is no unified approach or standardization how to carry out neurolinguistics experiments and acquire brain imaging data, so lots of heterogeneous multimodal datasets already exist, which could not be processed and analyzed in the same way. Moreover, usually these datasets are quite small in terms of number of subjects, see for example (Hanke et al. 2014; Huth et al. 2016; Wehbe et al. 2014) where there are less than 10 subjects in the dataset. In order to be able to use the data from already existing different datasets for models training or testing, it is necessary to develop new transfer learning approaches which take into account intersubject variability as well as scanning on different brain imaging equipment (with different settings) and using different experimental paradigms.

References

- Abnar, S., Ahmed, R., Mijnheer, M., and Zuidema, W. (2018). Experiential, distributional and dependency-based word embeddings have complementary roles in decoding brain activity. In *Proceedings of the 8th Workshop on Cognitive Modeling and Computational Linguistics (CMCL 2018)*, pages 57-66.
- Abnar, S., Beinborn, L., Choenni, R., Zuidema, W., 2019. Blackbox Meets Blackbox: Representational Similarity & Stability Analysis of Neural Language Models and Brains. <https://doi.org/10.18653/v1/w19-4820>
- Bartunov, Sergey, Dmitry Kondrashkin, Anton Osokin, and Dmitry Vetrov. "Breaking sticks and ambiguities with adaptive skip-gram." In *Artificial Intelligence and Statistics*, pp. 130-138. 2016
- Bartusiak, Roman, Łukasz Augustyniak, Tomasz Kajdanowicz, Przemysław Kazienko, and Maciej Piasecki. "WordNet2Vec: Corpora agnostic word vectorization method." *Neurocomputing* 326 (2019): 141-150.

Buckner, R.L., 1998. Event-related fMRI and the hemodynamic response, in: Human Brain Mapping. [https://doi.org/10.1002/\(SICI\)1097-0193\(1998\)6:5/6<373::AID-HBM8>3.0.CO;2-P](https://doi.org/10.1002/(SICI)1097-0193(1998)6:5/6<373::AID-HBM8>3.0.CO;2-P)

Cer, Daniel, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant et al. "Universal sentence encoder for English." In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 169-174. 2018.

Collobert, R., Weston, J., Bottou, L., Karlen, M., Kavukcuoglu, K., and Kuksa, P. (2011). Natural language processing (almost) from scratch. *The Journal of Machine Learning Research*, 12:2493-2537.

Cooney, C., Folli, R., Coyle, D., 2018. Neurolinguistics Research Advancing Development of a Direct-Speech Brain-Computer Interface. *iScience* 8, 103–125.
<https://doi.org/10.1016/j.isci.2018.09.016>

Dai, Zihang, Zhilin Yang, Yiming Yang, William W. Cohen, Jaime Carbonell, Quoc V. Le, and Ruslan Salakhutdinov. "Transformer-xl: Attentive language models beyond a fixed-length context." *arXiv preprint arXiv:1901.02860* (2019).

Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding." In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 4171-4186. 2019.

Ettinger, A., Feldman, N., Resnik, P., and Phillips, C. (2016). Modeling N400 amplitude using vector space models of word representation. In *CogSci.*, 1445-1450.

Fergadiotis, G., Gorman, K., Bedrick, S., 2016. Algorithmic classification of five characteristic types of paraphasias. *Am. J. Speech-Language Pathol.*
https://doi.org/10.1044/2016_AJSLP-15-0147

Fiebach, C. J., & Friederici, A. D. (2004). Processing concrete words: fMRI evidence against a specific right-hemisphere involvement. *Neuropsychologia*, 42(1), 62-70.

Frank Stefan L, 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.

Fraser, K.C., Fors, K.L., Kokkinakis, D., 2019. Multilingual word embeddings for the assessment of narrative speech in mild cognitive impairment I. *Comput. Speech Lang.* 53, 121–139.
<https://doi.org/10.1016/j.csl.2018.07.005>

Friederici, A. D., Opitz, B., & Von Cramon, D. Y. (2000). Segregating semantic and syntactic aspects of processing in the human brain: an fMRI investigation of different word types. *Cerebral cortex*, 10(7), 698-705.

Fyshe, Alona, Partha P. Talukdar, Brian Murphy, and Tom M. Mitchell. "Interpretable semantic vectors from a joint model of brain-and text-based meaning." In *Proceedings of the conference. Association for Computational Linguistics. Meeting*, vol. 2014, p. 489. NIH Public Access, 2014.

Fyshe, A., Sudre, G., Wehbe, L., Rafidi, N., Mitchel, T. M. The lexical semantics of adjective-noun phrases in the human brain. (2019). *Hum Brain Mapp*. 40:4457-4469. doi: 10.1002/hbm.24714

Jain, S. and Huth, A. (2018). Incorporating context into language encoding models for fmri. *bioRxiv*, page 327601.

Harris, Zellig S. "Distributional structure." *Word* 10, no. 2-3 (1954): 146-162.

Hanke, M., Baumgartner, F.J., Ibe, P., Kaule, F.R., Pollmann, S., Speck, O., Zinke, W., Stadler, J., 2014. A high-resolution 7-Tesla fMRI dataset from complex natural stimulation with an audio movie. *Sci. Data* 1, 1–18. <https://doi.org/10.1038/sdata.2014.3>

Hollenstein, N., de la Torre, A., Langer, N., Zhang, C., 2019. CogniVal: A Framework for Cognitive Word Embedding Evaluation. <https://doi.org/10.18653/v1/k19-1050>

Huster, R.J., Debener, S., Eichele, T., Herrmann, C.S., 2012. Methods for simultaneous EEG-fMRI: An introductory review. *J. Neurosci*.
<https://doi.org/10.1523/JNEUROSCI.0447-12.2012>

Huth, A.G., de Heer, W.A., Griffiths, T.L., Theunissen, F.E., Gallant, J.L., 2016. Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature* 532, 453–458.
<https://doi.org/10.1038/nature17637>

Ivanova, A.A., Mineroff, Z., Zimmerer, V., Kanwisher, N., Varley, R., Fedorenko, E., 2019. The language network is recruited but not required for non-verbal semantic processing. *bioRxiv*.
<https://doi.org/10.1101/696484>

Kiros, Ryan, Yukun Zhu, Ruslan R. Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. "Skip-thought vectors." In *Advances in neural information processing systems*, pp. 3294-3302. 2015.

Kuperberg, G.R., 2007. Neural mechanisms of language comprehension: Challenges to syntax. *Brain Res*. <https://doi.org/10.1016/j.brainres.2006.12.063>

Le, Quoc, and Tomas Mikolov. "Distributed representations of sentences and documents." In *International conference on machine learning*, pp. 1188-1196. 2014

Lan, Zhenzhong, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. "ALBERT: A Lite BERT for Self-supervised Learning of Language Representations." In *International Conference on Learning Representations*. 2019

Linz, N., Tröger, J., Alexandersson, J., König, A., 2017. Using Neural Word Embeddings in the Analysis of the Clinical Semantic Verbal Fluency Task, in: *IWCS 2017 --- 12th International Conference on Computational Semantics --- Short Papers*.

Liu, Yinhan, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. "Roberta: A robustly optimized bert pretraining approach." *arXiv preprint arXiv:1907.11692* (2019).

McKinney-Bock, K., Bedrick, S., 2019. Classification of Semantic Paraphrasias: Optimization of a Word Embedding Model. <https://doi.org/10.18653/v1/w19-2007>

Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean. "Distributed representations of words and phrases and their compositionality." In *Advances in neural information processing systems*, pp. 3111-3119. 2013.

Mitchell, T. M., Shinkareva, S. V., Carlson, A., Chang, K.-M., Malave, V. L., Mason, R. A., and Just, M. A. (2008). Predicting Human Brain Activity Associated with the Meanings of Nouns. *Science*, 320:1191-1195

Murphy, K., Birn, R.M., Bandettini, P.A., 2013. Resting-state fMRI confounds and cleanup. *Neuroimage* 80, 349–359. <https://doi.org/10.1016/j.neuroimage.2013.04.001>

Murphy, B., Talukdar, P., and Mitchell, T. (2012). Selecting corpus-semantic models for neuroin- guistic decoding. In *Proceedings of the First Joint Conference on Lexical and Computational Semantics-Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation*, pages 114-123. Association for Computational Linguistics.

Oota, S. R., Manwani, N., & Bapi, R. S. (2018, December). fMRI Semantic Category Decoding Using Linguistic Encoding of Word Embeddings. In *International Conference on Neural Information Processing* (pp. 3-15). Springer, Cham.

Pakhomov, S.V.S., Eberly, L., Knopman, D., 2016. Characterizing cognitive performance in a large longitudinal study of aging with computerized semantic indices of verbal fluency. *Neuropsychologia*. <https://doi.org/10.1016/j.neuropsychologia.2016.05.031>

Pelevina, Maria, Nikolay Arefiev, Chris Biemann, and Alexander Panchenko. "Making Sense of Word Embeddings." In *Proceedings of the 1st Workshop on Representation Learning for NLP*, pp. 174-183. 2016.

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

Pereira, F., Lou, B., Pritchett, B., Ritter, S., Gershman, S. J., Kanwisher, N., Botvinick, M., and Fedorenko, E. (2018). Toward a universal decoder of linguistic meaning from brain activation. *Nature communications*, 9(1), 963.

Peters, Matthew E., Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. "Deep contextualized word representations." arXiv preprint arXiv:1802.05365 (2018).

Raffel, Colin, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. "Exploring the limits of transfer learning with a unified text-to-text transformer." arXiv preprint arXiv:1910.10683 (2019)

Robert, P.H., Lafont, V., Medecin, I., Berthet, L., Thauby, S., Baudu, C., Darcourt, G., 1998. Clustering and switching strategies in verbal fluency tasks: Comparison between schizophrenics and healthy adults. *J. Int. Neuropsychol. Soc.* <https://doi.org/10.1017/S1355617798466025>

Rodrigues, J. A., Branco, R., Silva, J., Saedi, C., and Branco. A. (2018). Predicting brain activation with WordNet embeddings. In *Proceedings of the Eight Workshop on Cognitive Aspects of Computational Language Learning and Processing*, pages 1–5.

Roelofs, A., 2019. Phonological cueing of word finding in aphasia: insights from simulations of immediate and treatment effects. *Aphasiology* 00, 1–17.
<https://doi.org/10.1080/02687038.2019.1686748>

Roelofs, A., 1997. The WEAVER model of word-form encoding in speech production. *Cognition*.
[https://doi.org/10.1016/S0010-0277\(97\)00027-9](https://doi.org/10.1016/S0010-0277(97)00027-9)

Rustandi Indrayana, Marcel Adam Just, and Tom M Mitchell. 2009. Integrating Multiple-Study Multiple-Subject fMRI Datasets Using Canonical Correlation Analysis. In *MICCAI 2009 Workshop: Statistical modeling and detection issues in intra- and inter-subject functional MRI data analysis*.

Salmelin, R., 2007. Clinical neurophysiology of language: The MEG approach. *Clin. Neurophysiol.* <https://doi.org/10.1016/j.clinph.2006.07.316>

Sejnowski, T. J., Churchland, P. S., & Movshon, J. A. (2014). Putting big data to good use in neuroscience. *Nature neuroscience*, 17(11), 1440.

Søgaard, A. (2016). Evaluating word embeddings with fMRI and eye-tracking. In *Proceedings of the 1st Workshop on Evaluating Vector-Space Representations for NLP*, pages 116-121.

Sharaev, M., Andreev, A., Artemov, A., Burnaev, E., Kondratyeva, E., Sushchinskaya, S., Samotaeva, I., Gaskin, V., Bernstein, A., 2018a. Pattern Recognition Pipeline for Neuroimaging Data, in: 8th IAPR TC3 Workshop on Artificial Neural Networks in Pattern Recognition, September 19-21, 2018, Siena, Italy. pp. 306–319.
https://doi.org/10.1007/978-3-319-99978-4_24

Sharaev, M., Artemov, A., Kondrateva, E., Sushchinskaya, S., Burnaev, E., Bernstein, A., Akzhigitov, R., Andreev, A., 2018b. MRI-Based Diagnostics of Depression Concomitant with Epilepsy: In Search of the Potential Biomarkers, in: 2018 IEEE 5th International Conference on Data Science and Advanced Analytics (DSAA). IEEE, pp. 555–564.
<https://doi.org/10.1109/DSAA.2018.00071>

Schwartz, Dan, and Tom Mitchell. "Understanding language-elicited EEG data by predicting it from a fine-tuned language model." In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 43-57. 2019.

Schwartz, Dan, Mariya Toneva, and Leila Wehbe. "Inducing brain-relevant bias in natural language processing models." *Advances in Neural Information Processing Systems*. 2019.
Toneva M., and Leila Wehbe. "Interpreting and improving natural-language processing (in machines) with natural language-processing (in the brain)" . *Advances in Neural Information Processing Systems*. 2019.

Troyer, A.K., Moscovitch, M., Winocur, G., 1997. Clustering and switching as two components of verbal fluency: Evidence from younger and older healthy adults. *Neuropsychology*.
<https://doi.org/10.1037/0894-4105.11.1.138>

Troyer, A.K., Moscovitch, M., Winocur, G., Alexander, M.P., Stuss, D., 1998. Clustering and switching on verbal fluency: The effects of focal frontal- and temporal-lobe lesions. *Neuropsychologia*. [https://doi.org/10.1016/S0028-3932\(97\)00152-8](https://doi.org/10.1016/S0028-3932(97)00152-8)

Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E. J., Yacoub, E., & Ugurbil, K. (2013). The WU-Minn Human Connectome Project: An overview. *NeuroImage*.
<https://doi.org/10.1016/j.neuroimage.2013.05.041>

Wise, R., Hadar, U., Howard, D., Patterson, K., 1991. Language activation studies with positron emission tomography. *Ciba Found Symp* 163, 218–234.

Wehbe, L., Murphy, B., Talukdar, P., Fyshe, A., Ramdas, A., Mitchell, T., 2014a. Simultaneously uncovering the patterns of brain regions involved in different story reading Subprocesses. PLoS One 9, 1–19. <https://doi.org/10.1371/journal.pone.0112575>

Wehbe, L., Vaswani, A., Knight, K., Mitchell, T., 2014b. Aligning context-based statistical models of language with brain activity during reading, in: EMNLP 2014 - 2014 Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference. <https://doi.org/10.3115/v1/d14-1030>

Yang, Zhilin, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. "XLNet: Generalized Autoregressive Pretraining for Language Understanding." arXiv preprint arXiv:1906.08237 (2019).

Vu, M. A. T., Adalı, T., Ba, D., Buzsáki, G., Carlson, D., Heller, K., ... & Mayberg, H. S. (2018). A shared vision for machine learning in neuroscience. *Journal of Neuroscience*, 38(7), 1601-1607.