



BlaBla: Linguistic Feature Extraction for Clinical Analysis in Multiple Languages

Abhishek Shivkumar, Jack Weston, Raphael Lenain, Emil Fristed

Novoic Ltd

{abhishek, jack, raphael, emil}@novoic.com

Abstract

We introduce BlaBla, an open-source Python library for extracting linguistic features with proven clinical relevance to neurological and psychiatric diseases across many languages. BlaBla is a unifying framework for accelerating and simplifying clinical linguistic research. The library is built on state-of-the-art NLP frameworks and supports multithreaded/GPU-enabled feature extraction via both native Python calls and a command line interface. We describe BlaBla’s architecture and clinical validation of its features across 12 diseases. We further demonstrate the application of BlaBla to a task visualizing and classifying language disorders in three languages on real clinical data from the AphasiaBank dataset. We make the codebase freely available to researchers with the hope of providing a consistent, well-validated foundation for the next generation of clinical linguistic research.

Index Terms: Alzheimer’s disease, aphasia, BlaBla, computational linguistics, healthcare, machine learning (ML), natural language processing (NLP), Novoic, speech and language disorders, speech processing.

1. Introduction

Natural language processing (NLP) has many diverse use cases, including machine translation [1, 2], question answering [3], named entity recognition [4], sentiment analysis [5], document summarization [6] and topic modelling [7]. Some of the most important applications lie in the medical domain, where NLP is applied to tasks such as automatic processing of electronic health records (EHRs) [8]. Feature extraction from language (e.g. the number of noun phrases) is also used directly for analyzing patterns of speech to characterize different medical conditions. Speech and language changes have been observed in neurodegenerative diseases (e.g. Alzheimer’s disease), motor conditions (e.g. Parkinson’s) and affective disorders (e.g. depression) [9, 10, 11]; see Table 1 for details. Spoken language might therefore offer a universal and accessible means for measuring neurological health, but its application has been so far limited: feature extraction from language is often done manually or using a patchwork of different packages and custom scripts [12, 13]. Further, most work to date has been done in the English language. The research community is therefore faced with a pressing need to harmonize the extraction of language features and generalize this to languages other than English.

We introduce BlaBla, an easy-to-use Python library for language feature extraction, written with the aim of providing a unifying and consistent implementation of commonly used language features from the clinical literature across multiple languages. We write BlaBla with the aim of equipping every member of the NLP community (including those without prior Python knowledge) with powerful tooling by designing two distinct ways to extract features: either by calling BlaBla program-

matically in Python or via its command line interface (CLI). BlaBla was built to facilitate fast prototyping and complement current machine learning workflows using sklearn, TensorFlow or PyTorch for example.

Modern linguistic analysis, especially that in the clinical domain, has previously required stitching together multiple tools and resources. This typically includes transcript-parsing software such as CLAN [14] and NXT [15], low-level NLP libraries such as NLTK [16], Stanford CoreNLP [17], and hand-engineered features like Honoré’s statistic [18] and mean Yngve depth [19] following a literature review. Performing such analyses in languages other than, or in addition to, English further complicates the task. What lacks is a single source of truth able to operate faithfully across multiple languages.

Today, a growing number of open-source NLP toolkits with sophisticated offerings exist, with examples including CoreNLP [17], FLAIR [20] and spaCy [21]. A major limitation of these toolkits has been support of only a few languages: CoreNLP (implemented in Java) supports 6 human languages, FLAIR (Python) 12, and spaCy (Python) 10 [22]. Universal Dependencies (UD) is a framework for consistent annotation of grammar (parts of speech, morphological features and syntactic dependencies) across a variety of human languages [23]. The UD community has built a large database of treebanks in many languages, now used for training some newer toolkits. One example is UDPipe [24], a C++ toolkit with support for 61 human languages. Very recently, Stanza was released by the Stanford NLP group [22], a Python library with a fully neural pipeline trained on UD treebanks and other multilingual corpora. Based on raw text input, it produces annotations including tokenization, multi-word token expansion, lemmatization, part-of-speech (POS) and morphological feature tagging, dependency parsing, and named entity recognition [22]. It further offers a Python interface to CoreNLP, providing additional annotations such as the constituency parse tree, though only in the 6 languages supported by CoreNLP.

BlaBla is built using Stanza, which we consider the state-of-the-art NLP toolkit, and CoreNLP. Our algorithms sit on top of these packages to form a unified library for extracting clinically relevant features:

- **Phonetic and phonological features** calculated using timestamp-aligned transcripts.
- **Lexicosemantic features** derived from Stanza’s dependency parse tree.
- **Morphosyntactic and syntactic features** derived from CoreNLP’s constituency parse tree.
- **Discourse and pragmatic features** derived from Stanza’s dependency parse tree.

We describe BlaBla’s implementation and interface along with the clinical rationale behind its features, then illustrate the ability to rapidly prototype ML models using BlaBla. We take as

an example the task of visualizing and classifying aphasic patients in the AphasiaBank dataset [25] across three languages. We provide a list of reference values on the AMI dataset [26] and the clinical rationale underlying each feature. Finally, we release the BlaBla codebase¹ to the research community under an open-source license, along with notebooks containing all the code used in this paper².

2. BlaBla Architecture

Stanza utilizes deep neural networks for extracting linguistic structures, including the dependency parse tree and POS tags [22] but notably lacking the constituency parse tree, which is used to identify phrase-level structures such as noun phrases and verb phrases. This can be extracted by CoreNLP, which instead employs traditional statistical models [17]. BlaBla processes this low-level information from Stanza and CoreNLP to build a set of linguistic features tailored for the clinical research community. BlaBla takes as input a piece of text or a timestamp-aligned transcript and outputs a CSV file or a Pandas DataFrame of the derived features; timestamp-aligned transcripts are required only to calculate phonetic and phonological features. Language support for each feature is determined by its dependence on Stanza and/or CoreNLP. The languages supported by CoreNLP are English, Arabic, Chinese, French, German and Spanish [17], whereas Stanza supports 66 languages [22]. BlaBla itself has been tested on the 6 languages supported by CoreNLP and the remaining 60 languages are subject to ongoing validation. For a comprehensive account of language support for each feature, see Table 1.

We release BlaBla under the GNU GPL v3 open-source license to grant its users the freedom to use, adapt and build on top of this library. We hope that this will encourage contributors from both research and industry to offer their time, knowledge and support to the future of this project.

2.1. Motivation

The features implemented in BlaBla are chosen to reflect the most up-to-date feature sets commonly used in the diagnosis of neurodegenerative, motor and affective conditions. This was achieved through a clinical literature review drawing on a number of recent review papers [9, 10, 27, 28, 11]. The selected features were further split into four categories: *phonetic and phonological* (including features such as the hesitation ratio and the speech rate), *lexicosemantic* (including the noun rate and idea density), *morphosyntactic and syntactic* (including the proportion of auxiliary verbs and the mean Yngve depth) and *discourse and pragmatic* (including the rate of discourse markers), following the nomenclature of [9]. Further details of the features and their clinical motivation can be found in Table 1.

Also included in Table 1 is a set of reference values for each feature derived from manual transcripts from the AMI Meeting Corpus [26], excluding phonetic and phonological features. A total of 682 single-speaker transcripts were extracted from diarized meeting transcripts; each of these represents one data point in the distributions described by the reference values.

2.2. Using BlaBla

There are two distinct ways to use BlaBla. The first is the native Pythonic way, whereby the user imports and instantiates

the DocumentProcessor class with the path to the configuration file and text language. Features can then be extracted from the input via the `compute_features` method. This Python interface allows the user to easily integrate BlaBla with other frameworks such as NumPy and TensorFlow. The second way to use BlaBla is via the command line interface (CLI), which extracts features from a text/JSON file or directory of such files. The CLI reads its configuration from a YAML file to determine which features are to be extracted. Sentence parsing can be computationally burdensome and so, to address this, BlaBla’s calls to Stanza and CoreNLP are multithreaded and GPU-enabled by default. See BlaBla’s documentation for a comprehensive description of usage, structure and functionality.

3. Application: Visualizing and Classifying Aphasia Across Multiple Languages

3.1. Experimental Design

The features implemented in BlaBla are motivated by the clinical literature, with considered conditions including multiple types of aphasia (see Table 1). Aphasia broadly refers to a class of language disorders that leave their sufferer with difficulties understanding or producing language [29]. The characteristic patterns of impairment can be classified as either fluent or nonfluent aphasia [30]. Linguistic changes due to aphasia have been the subject of much research and many are well-documented. These include word-finding difficulties and naming disturbances [31, 32], changes in noun/verb rates [33, 34], changes in rates of closed-class words and the omission of determiners [35, 36], and reduced syntactic complexity/accuracy [37, 38]. These features can be used for classification of aphasia; in [39], for instance, the authors extract 58 linguistic measures (such as the inflected verb rate, the height of the parse tree and the rate of demonstratives) from the widely-used *Cinderella* story narration task [37] to classify primary progressive aphasia (PPA) patients using machine learning. They claim an accuracy of 100% distinguishing between fluent PPA and healthy controls, and 75% distinguishing between fluent PPA and nonfluent PPA on their small dataset.

In this section, we take the example of dimensionality reduction of linguistic features for aphasic patients and healthy controls in three languages to illustrate the clinical relevance of BlaBla features. We use the AphasiaBank dataset [40], a corpus of clinically diagnosed aphasic participants and healthy controls containing manually transcribed and annotated transcripts. We extract the *Cinderella* story narration task from the combined English corpora and the French corpus, along with the *Cry Wolf* story narration task from the Mandarin corpus, where the protocol was adapted for cultural reasons. Extraction of annotation-free transcripts was performed using the CLAN tool [14] and all linguistic analysis was performed using BlaBla v0.1. For the fluent aphasic, nonfluent aphasic and control participants in the English corpora, we take all 43 non-phonetic and -phonological BlaBla features from Table 1 and visualize them using a 2D t-SNE to visually identify groups of clinical significance. Only a subset of these features has been previously identified with aphasia but most have not been explicitly studied, so we choose to include them all for illustrative purposes. We repeat this procedure for aphasic participants versus controls in each of the three languages using a restricted feature list that excludes directly length-dependent features (those indicated by a dagger in Table 1) and features which are undefined for Mandarin (those

¹<https://github.com/novoic/blabla>

²<https://github.com/novoic/blabla-IS2020>

Table 1: A description of BlaBla features including their clinical rationale. This includes a set of reference values calculated using BlaBla v0.1 on the AMI corpus. The languages supported by CoreNLP and Stanza are described in the main text (see Section 2). The right half of the table is adapted from [9, 10, 27, 28, 11], summarizing clinical validation of recent review papers across indications. \uparrow = the feature increases relative to healthy controls; \downarrow = the feature decreases relative to healthy controls; * = the review has confirmed that the feature does not change; - = unknown. NF/av = nonfluent agrammatic variant primary progressive aphasia (PPA); Sv = semantic variant PPA; L/Pv = logopenic variant PPA; AD = Alzheimer’s disease; MND = motor neurone disease, synonymous with amyotrophic lateral sclerosis (ALS); PD = Parkinson’s disease; HD = Huntington’s disease; MS = multiple sclerosis; MDD = major depressive disorder; HpM = hypomania; Szo = schizophrenia. A dagger (\dagger) indicates that the feature expected to scale with the text length so the reference value (if provided) is given only for completeness. Phonetic and phonological features are calculated directly using timestamp-aligned transcripts so their implementation is language-agnostic.

Feature	Base (#languages)	AMI reference	AD	Sv	L/Pv	NF/av	MND	PD	LDB	HD	MS	MDD	HpM	Szo
Phonetic & phonological features														
Number of pauses \dagger	-	-	-	-	-	-	-	\uparrow [9]	-	-	\uparrow [28]	-	-	-
Total pause time \dagger	-	-	*[9]	\uparrow [9]	-	\uparrow [9]	\uparrow [9]	*[9]	\uparrow [9]	-	-	\uparrow [10]	\uparrow [10]	\uparrow [10]
Mean pause duration	-	-	*[9]	\downarrow [9]	-	-	-	-	-	-	\uparrow [28]	*[10]	\uparrow [10]	\uparrow [10]
Between-utterance pause duration	-	-	-	-	-	-	-	\uparrow [9]	-	-	-	-	-	-
Hesitation ratio	-	-	\uparrow [9]	-	-	\uparrow [9]	-	-	-	-	-	-	-	-
Speech rate	-	-	\downarrow [9]	\downarrow [9]	\downarrow [9]	\downarrow [9]	*[9]	*[9]	\downarrow [9]	-	\downarrow [28]	\downarrow [10]	-	\downarrow [10]
Maximum speech rate	-	-	-	-	\downarrow [9]	\downarrow [9]	-	-	-	-	-	-	-	-
Total phonation time \dagger	-	-	*[9]	-	-	-	-	-	-	-	\downarrow [28]	-	-	\downarrow [10]
Standardized phonation time \dagger	-	-	\downarrow [9]	-	-	-	-	-	-	-	-	-	-	-
Total locution time \dagger	-	-	*[9]	*[9]	*[9]	\uparrow [9]	*[9]	*[9]	*[9]	-	-	-	-	-
Lexicosemantic features														
Noun rate	Stanza (66)	0.123 \pm 0.0223	*[9]	\downarrow [9]	*[9]	*[9]	*[9]	*[9]	*[9]	-	-	-	-	-
Verb rate	Stanza (66)	0.108 \pm 0.0168	\uparrow [9]	*[9]	*[9]	*[9]	*[9]	-	-	*[9]	-	-	-	-
Demonstrative rate	Stanza (66)	0.0263 \pm 0.00869	*[9]	\uparrow [9]	-	\uparrow [9]	-	-	-	-	-	-	-	-
Adjective rate	Stanza (66)	0.0479 \pm 0.0116	*[9]	*[9]	-	*[9]	-	-	-	*[9]	-	-	-	-
Pronoun rate	Stanza (66)	0.144 \pm 0.0201	*[9]	\uparrow [9]	\uparrow [9]	*[9]	*[9]	-	-	-	-	-	-	-
Adverb rate	Stanza (66)	0.0742 \pm 0.0157	*[9]	\uparrow [9]	-	*[9]	-	-	-	*[9]	-	-	-	-
Conjunction rate	Stanza (66)	0.0339 \pm 0.00946	*[9]	-	-	-	-	-	-	-	-	-	-	-
Possessive rate	Stanza (66)	0.192 \pm 0.0215	-	*[9]	-	\downarrow [9]	-	-	-	-	-	-	-	-
Noun-verb ratio	Stanza (66)	1.17 \pm 0.313	-	\downarrow [9]	-	*[9]	-	-	-	-	-	-	-	-
Noun ratio	Stanza (66)	0.531 \pm 0.0621	*[9]	\downarrow [9]	-	*[9]	-	-	-	-	-	-	-	-
Pronoun-noun ratio	Stanza (66)	1.22 \pm 0.362	\uparrow [9]	-	-	-	-	-	-	-	-	-	-	-
Closed-class word rate	Stanza (66)	0.312 \pm 0.0244	\uparrow [9]	-	-	-	-	-	-	-	-	-	-	-
Open-class word rate	Stanza (66)	0.353 \pm 0.0315	*[9]	\downarrow [9]	\downarrow [9]	*[9]	-	*[9]	*[9]	-	-	-	-	-
Content density	Stanza (66)	1.14 \pm 0.122	*[9]	\downarrow [9]	-	-	-	-	-	-	-	-	-	-
Idea density	Stanza (66)	0.324 \pm 0.0310	\downarrow [9]	-	-	-	-	-	-	-	-	-	-	-
Honoré’s statistic	Stanza (66)	(1.61 \pm 0.156) \times 10 ³	*[9]	-	-	-	-	-	-	-	-	-	-	-
Brunet’s index	Stanza (66)	14.9 \pm 1.55	*[9]	-	-	-	-	-	-	-	-	-	-	-
Type-token ratio	Stanza (66)	0.247 \pm 0.0886	*[9]	*[9]	-	*[9]	-	-	-	-	-	-	-	-
Word length	Stanza (66)	3.42 \pm 0.187	*[9]	\downarrow [9]	-	\downarrow [9]	-	-	-	-	-	-	-	-
Morphosyntactic & syntactic features														
Proportion of inflected verbs	Stanza (66)	0.337 \pm 0.0930	*[9]	*[9]	*[9]	*[9]	*[9]	*[9]	*[9]	-	-	-	-	-
Proportion of auxiliary verbs	Stanza (66)	0.798 \pm 0.205	*[9]	-	-	*[9]	-	-	-	-	-	-	-	-
Proportion of gerund verbs	Stanza (66)	0.0393 \pm 0.0307	\downarrow [9]	-	-	-	-	-	-	-	-	-	-	-
Proportion of participles	Stanza (66)	0.141 \pm 0.0669	\downarrow [9]	-	-	-	-	-	-	-	-	-	-	-
Number of clauses \dagger	CoreNLP (6)	311 \pm 250	*[9]	*[9]	-	*[9]	-	-	-	*[9]	-	-	-	-
Clause rate	CoreNLP (6)	3.38 \pm 1.28	-	*[9]	-	\downarrow [9]	-	-	-	-	-	-	-	-
Proportion of nouns with determiners	Stanza (66)	0.544 \pm 0.0999	*[9]	*[9]	*[9]	-	-	-	-	-	-	-	-	-
Proportion of nouns with adjectives	Stanza (66)	0.215 \pm 0.0635	-	-	-	-	-	-	-	-	-	-	-	-
Number of noun phrases \dagger	CoreNLP (6)	450 \pm 342	\downarrow [9]	-	-	\uparrow [9]	-	-	-	-	-	-	-	-
Noun phrase rate	CoreNLP (6)	4.89 \pm 1.64	-	-	-	-	-	-	-	-	-	-	-	-
Number of verb phrases \dagger	CoreNLP (6)	302 \pm 234	\downarrow [9]	*[9]	-	\downarrow [9]	-	-	-	-	-	-	-	-
Verb phrase rate	CoreNLP (6)	3.31 \pm 1.23	-	-	-	-	-	-	-	-	-	-	-	-
Number of infinitive phrases \dagger	Stanza (66)	16.2 \pm 12.7	-	-	-	\downarrow [9]	-	-	-	-	-	-	-	-
Infinitive phrase rate	Stanza (66)	0.189 \pm 0.120	-	-	-	-	-	-	-	-	-	-	-	-
Number of prepositional phrases \dagger	CoreNLP (6)	89.0 \pm 74.4	-	-	-	\downarrow [9]	-	-	-	-	-	-	-	-
Prepositional phrase rate	CoreNLP (6)	0.979 \pm 0.470	-	-	-	-	-	-	-	-	-	-	-	-
Number of dependent clauses \dagger	CoreNLP (6)	72.8 \pm 67.5	*[9]	*[9]	*[9]	\downarrow [9]	-	*[9]	\downarrow [9]	*[9]	-	-	-	-
Dependent clause rate	CoreNLP (6)	0.780 \pm 0.400	-	-	-	-	-	-	-	-	-	-	-	-
Max Yngve depth	CoreNLP (6)	4.98 \pm 0.707	*[9]	*[9]	-	*[9]	-	-	-	-	-	-	-	-
Mean Yngve depth	CoreNLP (6)	2.69 \pm 0.345	-	-	-	-	-	-	-	-	-	-	-	-
Total Yngve depth	CoreNLP (6)	58.0 \pm 27.6	-	-	-	-	-	-	-	-	-	-	-	-
Parse tree height	CoreNLP (6)	10.4 \pm 1.71	*[9]	*[9]	-	*[9]	-	-	-	-	-	-	-	-
Discourse & pragmatic features														
Number of discourse markers \dagger	Stanza (66)	75.6 \pm 62.7	\uparrow [9]	-	-	-	-	-	-	-	-	-	-	-
Discourse marker rate	Stanza (66)	0.834 \pm 0.465	-	-	-	-	-	-	-	-	-	-	-	-

relating to verb forms) to facilitate a fairer comparison between different languages and tasks.

English tends to dominate clinical language datasets [25, 41, 42], so understanding the extent to which feature-based linguistic analysis generalizes across languages is of great interest. We train a simple aphasia/control classifier using a linear SVM on an English train set and validate it on English, French and Mandarin test sets as a first step towards this. To encourage robustness to task and language variation, we again

use the restricted feature list. To account for data scarcity and feature redundancy, we first select the five most important features through recursive feature elimination and use only these to perform the classification. For simplicity, we randomly select balanced train and test sets.

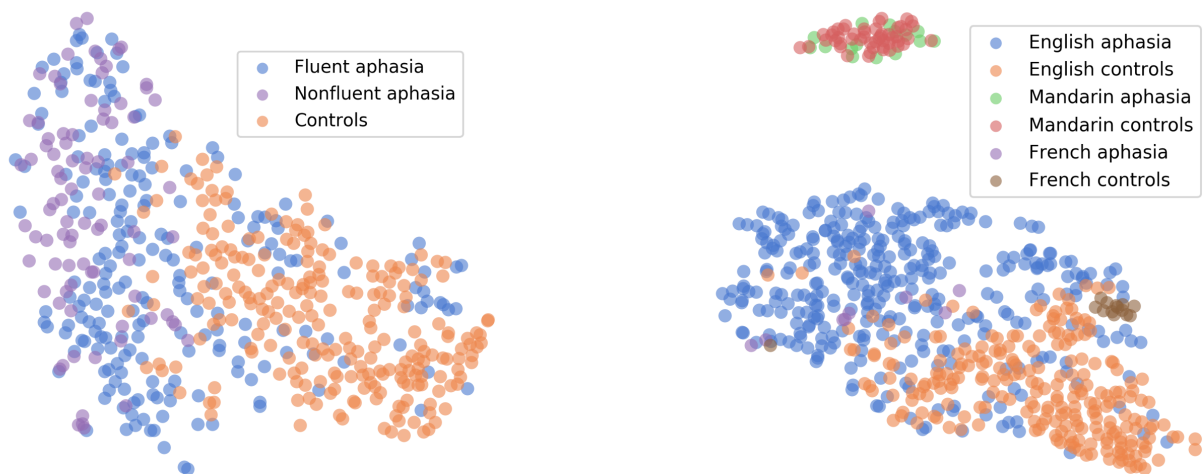


Figure 1: *t*-SNE visualizations of a subset of BlaBla features on story narration tasks from English, French and Mandarin corpora. Left: visualization of fluent aphasia, nonfluent aphasia and controls derived from the English corpora only. Right: visualization of aphasic and control patients from corpora spanning English, Mandarin and French.

3.2. Results

The results of the *t*-SNE visualizations are shown in Figure 1. Despite being fully unsupervised, the distinction between the controls and nonfluent aphasic participants in English is clear. Fluent aphasia is visually harder to separate from both nonfluent aphasia and controls, which we note does not necessarily imply that these groups are difficult to separate in their full 43-dimensional space. Data from the French corpus appears largely consistent with the English corpora, whereas the Mandarin corpus appears less so. This could be due to innate linguistic differences or those arising from the task variation, though scarcity of data should be noted before drawing conclusions.

Table 2: Performance of the aphasia/control classifier on the English, French and Mandarin test sets.

	English	French	Mandarin
Train $N_{\text{aph}} : N_{\text{contr}}$	201:201	-	-
Test $N_{\text{aph}} : N_{\text{contr}}$	36:36	9:9	15:15
Accuracy	0.903	0.833	0.500
F1 score	0.902	0.833	0.333

Table 2 shows the results of the aphasia/control classifier trained on English and tested on each of the three languages. The five features chosen through recursive feature elimination were: the noun-verb ratio (\uparrow), the pronoun-noun ratio (\uparrow), the mean Yngve depth (\downarrow), the pronoun rate (\downarrow) and the content density (\downarrow), where the arrows indicate the change direction indicating aphasia as implied by the linear SVM coefficient sign. The classifier achieves an accuracy of 90.3% on the English test set and 83.3% on the French test set but does not perform better than random on the Mandarin test set, suggesting that this simple approach may generalize better to other European languages than more distant languages. Domain adaptation techniques [43] may be necessary to achieve the cross-language and -task generalizability required here. Other complicating factors

include the small data set size and the diverse subtypes of aphasia within the ‘aphasia’ class. Finally, we note that we chose not to optimize model parameters in this illustrative example.

4. Conclusions

In this paper we introduced BlaBla, a Python package for NLP with application to the clinical domain. A multilingual study of BlaBla features on aphasic discourse revealed significant separation between aphasic and control participants, demonstrating the clinical significance of its linguistic features. A classifier trained on English discourse and validated on French performed significantly better than baseline, showing resilience to domain shift, but the same model was unable to generalize to Mandarin, suggesting a need for more proactive methods for dealing with larger shifts. We hope that BlaBla will facilitate the next wave of clinical NLP analyses and prove to be an invaluable, evolving resource for both experts and newcomers to the field.

5. Acknowledgements

The authors would like to acknowledge and thank the participants, organizers and sponsors who contribute to and maintain the AphasiaBank³ database, without whom the aphasia illustration used in this paper would not have been possible.

6. References

- [1] I. Sutskever, O. Vinyals, and Q. V. Le, “Sequence to sequence learning with neural networks,” in *Advances in neural information processing systems*, 2014, pp. 3104–3112.
- [2] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” *arXiv preprint arXiv:1409.0473*, 2014.
- [3] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.

³<https://aphasia.talkbank.org/>

- [4] D. Nadeau and S. Sekine, "A survey of named entity recognition and classification," *Linguisticae Investigationes*, vol. 30, no. 1, pp. 3–26, 2007.
- [5] B. Pang, L. Lee *et al.*, "Opinion mining and sentiment analysis," *Foundations and Trends® in Information Retrieval*, vol. 2, no. 1–2, pp. 1–135, 2008.
- [6] X. Zhang, F. Wei, and M. Zhou, "Hibert: Document level pre-training of hierarchical bidirectional transformers for document summarization," *arXiv preprint arXiv:1905.06566*, 2019.
- [7] D. M. Blei, "Probabilistic topic models," *Communications of the ACM*, vol. 55, no. 4, pp. 77–84, 2012.
- [8] L. Ohno-Machado, "Realizing the full potential of electronic health records: the role of natural language processing," *Journal of the American Medical Informatics Association*, vol. 18, no. 5, pp. 539–539, 2011.
- [9] V. Boschi, E. Catricala, M. Consonni, C. Chesi, A. Moro, and S. F. Cappa, "Connected speech in neurodegenerative language disorders: a review," *Frontiers in psychology*, vol. 8, p. 269, 2017.
- [10] D. M. Low, K. H. Bentley, and S. S. Ghosh, "Automated assessment of psychiatric disorders using speech: A systematic review," *Laryngoscope Investigative Otolaryngology*, vol. 5, no. 1, pp. 96–116, 2020.
- [11] J. C. Chan, J. C. Stout, and A. P. Vogel, "Speech in prodromal and symptomatic huntington's disease as a model of measuring onset and progression in dominantly inherited neurodegenerative diseases," *Neuroscience & Biobehavioral Reviews*, 2019.
- [12] M. De Choudhury, M. Gamon, S. Counts, and E. Horvitz, "Predicting depression via social media," in *Seventh international AAAI conference on weblogs and social media*, 2013.
- [13] K. C. Fraser, J. A. Meltzer, and F. Rudzicz, "Linguistic features identify alzheimer's disease in narrative speech," *Journal of Alzheimer's Disease*, vol. 49, no. 2, pp. 407–422, 2016.
- [14] B. MacWhinney and J. Wagner, "Transcribing, searching and data sharing: The clan software and the talkbank data repository," *Gesprächsforschung: Online-Zeitschrift zur verbalen Interaktion*, vol. 11, p. 154, 2010.
- [15] S. Calhoun, J. Carletta, J. M. Brenier, N. Mayo, D. Jurafsky, M. Steedman, and D. Beaver, "The nxt-format switchboard corpus: a rich resource for investigating the syntax, semantics, pragmatics and prosody of dialogue," *Language resources and evaluation*, vol. 44, no. 4, pp. 387–419, 2010.
- [16] E. Loper and S. Bird, "Nltk: the natural language toolkit," *arXiv preprint cs/0205028*, 2002.
- [17] C. D. Manning, M. Surdeanu, J. Bauer, J. R. Finkel, S. Bethard, and D. McClosky, "The stanford corenlp natural language processing toolkit," in *Proceedings of 52nd annual meeting of the association for computational linguistics: system demonstrations*, 2014, pp. 55–60.
- [18] A. Honoré, "Some simple measures of richness of vocabulary," *Association for literary and linguistic computing bulletin*, vol. 7, no. 2, pp. 172–177, 1979.
- [19] V. Yngve, *The depth hypothesis*, 1972.
- [20] A. Akbik, T. Bergmann, D. Blythe, K. Rasul, S. Schweter, and R. Vollgraf, "Flair: An easy-to-use framework for state-of-the-art nlp," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, 2019, pp. 54–59.
- [21] spacy github repository. [Online]. Available: <https://github.com/explosion/spaCy>
- [22] P. Qi, Y. Zhang, Y. Zhang, J. Bolton, and C. D. Manning, "Stanza: A python natural language processing toolkit for many human languages," *arXiv preprint arXiv:2003.07082*, 2020.
- [23] J. Nivre, M.-C. De Marneffe, F. Ginter, Y. Goldberg, J. Hajic, C. D. Manning, R. McDonald, S. Petrov, S. Pyysalo, N. Silveira *et al.*, "Universal dependencies v1: A multilingual treebank collection," in *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, 2016, pp. 1659–1666.
- [24] M. Straka, "Udpipe 2.0 prototype at conll 2018 ud shared task," in *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, 2018, pp. 197–207.
- [25] M. M. Forbes, D. Fromm, and B. MacWhinney, "Aphasiabank: A resource for clinicians," in *Seminars in speech and language*, vol. 33, no. 03. Thieme Medical Publishers, 2012, pp. 217–222.
- [26] J. Carletta, S. Ashby, S. Bourban, M. Flynn, M. Guillemot, T. Hain, J. Kadlec, V. Karaiskos, W. Kraaij, M. Kronenthal *et al.*, "The ami meeting corpus: A pre-announcement," in *International workshop on machine learning for multimodal interaction*. Springer, 2005, pp. 28–39.
- [27] B. Tomik and R. J. Guiloff, "Dysarthria in amyotrophic lateral sclerosis: A review," *Amyotrophic Lateral Sclerosis*, vol. 11, no. 1-2, pp. 4–15, 2010.
- [28] G. Noffs, T. Perera, S. C. Kolbe, C. J. Shanahan, F. M. Boonstra, A. Evans, H. Butzkueven, A. van der Walt, and A. P. Vogel, "What speech can tell us: A systematic review of dysarthria characteristics in multiple sclerosis," *Autoimmunity reviews*, vol. 17, no. 12, pp. 1202–1209, 2018.
- [29] D. G. Clark and J. L. Cummings, "Aphasia," in *Neurological Disorders*. Elsevier, 2003, pp. 265–275.
- [30] J. K. Gordon and S. Clough, "How fluent? part b. underlying contributors to continuous measures of fluency in aphasia," *Aphasiology*, pp. 1–21, 2020.
- [31] W. U. Dressler and C. Pleh, "On text disturbances in aphasia," in *Linguistic analyses of aphasic language*. Springer, 1988, pp. 151–178.
- [32] C. Larfeuil and G. L. Dorze, "An analysis of the word-finding difficulties and of the content of the content of the discourse of recent and chronic aphasic speakers," *Aphasiology*, vol. 11, no. 8, pp. 783–811, 1997.
- [33] E. Bates, S. Chen, O. Tzeng, P. Li, and M. Opie, "The noun-verb problem in chinese aphasia," *Brain and language*, vol. 41, no. 2, pp. 203–233, 1991.
- [34] A. E. Hillis, S. Oh, and L. Ken, "Deterioration of naming nouns versus verbs in primary progressive aphasia," *Annals of Neurology: Official Journal of the American Neurological Association and the Child Neurology Society*, vol. 55, no. 2, pp. 268–275, 2004.
- [35] A. D. Friederici, "Syntactic and semantic processes in aphasic deficits: The availability of prepositions," *Brain and Language*, vol. 15, no. 2, pp. 249–258, 1982.
- [36] B. T. M. Hofstede, *Agrammatic speech in Broca's aphasia: strategic choice for the elliptical register*. Nijmegen: NICI, Nijmeegs Instituut voor Cognitie en Informatie, 1992.
- [37] H. Bird and S. Franklin, "Cinderella revisited: A comparison of fluent and non-fluent aphasic speech," *Journal of neurolinguistics*, vol. 9, no. 3, pp. 187–206, 1996.
- [38] M. F. Schwartz, E. M. Saffran, D. E. Bloch, and G. S. Dell, "Disordered speech production in aphasic and normal speakers," *Brain and language*, vol. 47, no. 1, pp. 52–88, 1994.
- [39] K. C. Fraser, J. A. Meltzer, N. L. Graham, C. Leonard, G. Hirst, S. E. Black, and E. Rochon, "Automated classification of primary progressive aphasia subtypes from narrative speech transcripts," *cortex*, vol. 55, pp. 43–60, 2014.
- [40] B. MacWhinney, D. Fromm, M. Forbes, and A. Holland, "Aphasiabank: Methods for studying discourse," *Aphasiology*, vol. 25, no. 11, pp. 1286–1307, 2011.
- [41] J. Gratch, R. Artstein, G. M. Lucas, G. Stratou, S. Scherer, A. Nazarian, R. Wood, J. Boberg, D. DeVault, S. Marsella *et al.*, "The distress analysis interview corpus of human and computer interviews," in *LREC*. Citeseer, 2014, pp. 3123–3128.
- [42] J. T. Becker, F. Boiler, O. L. Lopez, J. Saxton, and K. L. McGonigle, "The natural history of alzheimer's disease: description of study cohort and accuracy of diagnosis," *Archives of Neurology*, vol. 51, no. 6, pp. 585–594, 1994.
- [43] S. J. Pan, I. W. Tsang, J. T. Kwok, and Q. Yang, "Domain adaptation via transfer component analysis," *IEEE Transactions on Neural Networks*, vol. 22, no. 2, pp. 199–210, 2010.