



# Objective Assessment of Social Skills Using Automated Language Analysis for Identification of Schizophrenia and Bipolar Disorder

Rohit Voleti<sup>1</sup>, Stephanie Woolridge<sup>3</sup>, Julie M. Liss<sup>2</sup>, Melissa Milanovic<sup>3</sup>, Christopher R. Bowie<sup>3</sup>,  
Visar Berisha<sup>1,2</sup>

<sup>1</sup>School of Electrical, Computer, & Energy Engineering, Arizona State University, USA

<sup>2</sup>Department of Speech & Hearing Science, Arizona State University, USA

<sup>3</sup>Department of Psychology, Queen's University, Canada

rivoleti@asu.edu, l2smw2@queensu.ca, jmliss@asu.edu, melissa.milanovic@queensu.ca  
bowiec@queensu.ca, visar@asu.edu

## Abstract

Several studies have shown that speech and language features, automatically extracted from clinical interviews or spontaneous discourse, have diagnostic value for mental disorders such as schizophrenia and bipolar disorder. They typically make use of a large feature set to train a classifier for distinguishing between two groups of interest, *i.e.* a clinical and control group. However, a purely data-driven approach runs the risk of overfitting to a particular data set, especially when sample sizes are limited. Here, we first down-select the set of language features to a small subset that is related to a well-validated test of functional ability, the Social Skills Performance Assessment (SSPA). This helps establish the concurrent validity of the selected features. We use only these features to train a simple classifier to distinguish between groups of interest. Linear regression reveals that a subset of language features can effectively model the SSPA, with a correlation coefficient of 0.75. Furthermore, the same feature set can be used to build a strong binary classifier to distinguish between healthy controls and a clinical group (AUC = 0.96) and also between patients within the clinical group with schizophrenia and bipolar I disorder (AUC = 0.83).

**Index Terms:** computational linguistics, schizophrenia, bipolar disorder, semantic coherence, natural language processing

## 1. Introduction & Previous Work

In the United States alone, the National Institute of Mental Health (NIMH) in 2016 estimated that  $\sim 10.4$  million individuals live with a form of severe mental illness, approximately 4.2% of the adult population [1]. Among these are *schizophrenia* and *bipolar disorder*, for which diagnosis is difficult and treatment costs are disproportionately high [2]. Additionally, differential diagnosis of bipolar I disorder and schizophrenia is often difficult, with some estimating about 30% of bipolar I patients are misdiagnosed [3]. Therefore, there is a demand for effective methods with which we can classify and track the progress of treatment in these conditions. Language impairments are a well-known component of schizophrenia and bipolar disorder, including symptoms like *alogia* (poverty of speech) or development of *formal thought disorder* (FTD), including *schizophasia* ("word salad" or semantically incoherent utterances) [4]. These impairments are typically assessed by clinical interviews, but few quantitative measures exist for measuring them objectively. Recent work in computational linguistics and natural language processing (NLP) have paved the way for research into *computational psychiatry* to objectively assess the degree of language impairment [5]. Several recent studies

have made use of these tools for psychiatric evaluation, but their presence in clinical practice is still largely absent [6]. In this paper, we aim to bridge this gap by presenting an objective and interpretable panel of language features for assessment of patients with schizophrenia and bipolar disorder that is anchored to a well-validated clinical assessment of social skills.

Most existing work in this area takes a largely data-driven approach to language analysis, considering a host of semantic and lexical complexity measures over a large variety of language elicitation tasks [7, 8, 9, 10, 11, 12]. Semantic features are often captured with numerical word and sentence *embeddings*, in which words, sentences, phrases, *etc.* are represented in high-dimensional vector space; typically, words that are semantically similar are embedded close together in this vector space, *e.g.* *latent semantic analysis* (LSA) [13], *word2vec* [14], and several others. Another measure of semantics can be achieved by topic modeling, such as with *latent dirichlet allocation* (LDA) [15]. Semantic features are often combined with other lexical measures of language complexity to improve classification performance. Some examples are "surface features" (*i.e.* words per sentence, speaking rate, *etc.*) with tools like *Linguistic Inquiry and Word Count* (LIWC) [16], statistical language features (*n*-gram word likelihoods) [7], part-of-speech tag statistics [8, 9, 17], and sentiment analysis [18, 19, 20].

Despite promising early results, these tools are not currently used in clinical practice. We posit that this is because the large and varied feature space, the variability associated with the speech elicitation tasks, and the small sample sizes make it difficult to develop reliable and interpretable algorithms that generalize. As patient data is scarce, the identification of a standard set of important, interpretable, and easy-to-compute language features that clinicians can use is a significant hurdle to overcome. We address this by evaluating the language of patients with schizophrenia, patients with bipolar I disorder, and healthy control subjects on the *Social Skills Performance Assessment* (SSPA) [21], a well-validated test of social functional competence (described in Section 2). Our approach is motivated by our previous work in interpretable clinical-speech analytics [22]. First, we identify a subset of language measures that reliably model clinical SSPA scores. Next, we use only this reduced feature set to perform two classification problems: (1) distinguishing between healthy controls and clinical subjects and (2) distinguishing patients with schizophrenia/schizoaffective disorder (Sz/Sza) and bipolar I patients within the clinical group. To the best of our knowledge, this is the first study to establish a set of language measures that jointly assess social skills and uses those features to accurately classify all groups of interest.

## 2. SSPA Data Collection

Our study involves the analysis of interview transcripts collected from a total of 87 clinical subjects and 22 healthy controls that participated in the SSPA task described by Patterson *et al.* [21]. Of the clinical population, 44 had been diagnosed with bipolar I disorder and 43 had been diagnosed with schizophrenia or schizoaffective disorder (considered together in this analysis). The SSPA interviews are described by Bowie *et al.* in [23]. The transcriptions used in our analysis were completed at Queen’s University in Kingston, ON, Canada.

The task consists of three role-playing scenes: (1) 1-minute practice scene of making plans with a friend (not scored), (2) 3 minutes of greeting a new neighbor, and (3) 3 minutes of negotiation with a recalcitrant landlord over fixing an unrepaired leak. Each session was recorded and scored by trained research assistants upon reviewing the recording. Scene 2 (new neighbor) and Scene 3 (negotiation with landlord) were scored on a scale of 1 (low) to 5 (high) on several categories, *i.e.* interest/disinterest, fluency, clarity, social appropriateness, negotiation ability, *etc.* A composite score for each scene and an overall score is computed by averaging Scene 2 and Scene 3 scores.

Bowie *et al.* identified group differences between the scores of both clinical populations and healthy control subjects in [23] by evaluation on the SSPA task and several other clinical measures. In this work, we aim to automate this task with a subset of language metrics from the SSPA transcripts. As stated in Section 1, our first goal is to identify semantic and lexical features from which we can reliably predict SSPA performance. Then, we test the ability of these features to differentiate between healthy control and clinical populations, and we also test their ability to differentiate within the distinct groups in the clinical population.

## 3. Computed Language Features

In our work, we attempt to identify a comprehensive set of objective language measures from which we can model and predict SSPA performance and classify individuals using these features. Inspired by much of the previous work described in Section 1, we theorized that it is critical to consider language features that model semantic coherence through the use of word and sentence embeddings. We focused on a few pre-trained neural embedding models that are publicly available and known to model semantic similarity accurately. Additionally, we consider a set of lexical complexity features that are measures of lexical and syntactic complexity, described below.

### 3.1. Semantic Coherence

Many of the previously described studies in this area involve computing a notion of semantic coherence in language with the use of word embeddings in high-dimensional vector space, either with LSA or neural word embedding techniques [7, 8, 9, 10]. In nearly all cases, word and sentence/phrase embedding pairs, denoted by vectors  $\mathbf{a}$  and  $\mathbf{b}$ , are evaluated with the notion of *cosine similarity*, a measure of the cosine of the angle  $\theta_{\mathbf{a},\mathbf{b}}$  between the two vectors. We also use cosine similarity as a measure of pairwise sentence similarity, but with some modifications in implementation due to the difference in the nature of the SSPA task and data collection.

Our work differs from several of the previously discussed studies in that we are interested in conversational semantic similarity between the subject and clinical assessor in each of the three scenes of the SSPA task. Therefore, we sought to uti-

lize some of the latest sentence/phrase embedding methods to compute a vector representation for each assessor and subject speaking turn. Then, we used the cosine similarity to compute the similarity score between each consecutive assessor + subject speaking turn, generating a distribution of similarity scores for each embedding method for each subject in each transcribed scene. The following sentence embedding representations are used in our analysis: (1) an unweighted bag-of-words (BoW) average for all word vectors based on the pre-trained *skip-gram* implementation of *word2vec* trained on the Google News corpus [14], (2) *Smooth Inverse Frequency* (SIF) with pre-trained skip-gram *word2vec* vectors [24], and (3) *InferSent* (INF) sentence encodings based on pre-trained *FastText* vectors [25]. The BoW average of vectors and SIF embeddings showed good baseline performance in [10], and we additionally included *InferSent*, a deep neural network sentence encoder, due to its strong performance on semantic similarity tasks. Then, basic statistics for the similarity score distribution were computed for each subject and transcribed scene. These included minimum, maximum, mean, median, 90<sup>th</sup> percentile, and 10<sup>th</sup> percentile coherence.

### 3.2. Linguistic Complexity

While semantic coherence measures are often the most effective at classifying patients with schizophrenia and bipolar disorder, several other linguistic complexity measures are used for a more holistic analysis. We consider a subset of these features, computed for the entire set of subject responses across all three scene transcripts.

*Lexical diversity* refers to unique vocabulary usage for a particular subject and for which several measurement techniques exist. The *type-to-token ratio* (TTR) is a well-known measure of lexical diversity, in which the number of unique words (word *types*,  $V$ ) are compared against the total number of words (word *tokens*,  $N$ ):  $TTR = V/N$ . However, TTR is known to be negatively impacted for longer utterances, as the diversity of unique words plateaus as the number of total words increase. Hence, we consider a small selection of modified measures for lexical diversity in our work. The moving average type-to-token ratio (MATTR) [26] is one such method which aims to reduce the dependence on text length by considering TTR over a sliding window of the text. *Brunet’s Index* (BI) [27], defined in Equation (1), is another measure of lexical diversity that has a weaker dependence on text length. A smaller value indicates a greater degree of lexical diversity

$$BI = N^{V^{-0.165}} \quad (1)$$

An alternative is also provided by *Honoré’s Statistic* (HS) [28], defined in Equation (2), which emphasizes the use of words that are spoken only once (denoted by  $V_1$ ).

$$HS = 100 \log \frac{N}{1 - V_1/V} \quad (2)$$

MATTR, BI, and HS have been used successfully in computational linguistics studies for patients with Alzheimer’s disease [17, 29] and may prove to be similarly useful in our task.

Because we expect schizophrenia and bipolar patients to sometimes exhibit poverty of speech, we considered a few measures of lexical and syntactic complexity in our work.

*Lexical density*, which quantifies the degree of information packaging in a given text, is defined as the proportion of *content words* (*i.e.* nouns, verbs, adjectives, adverbs) [30]. Typically,




Figure 1: A linear regression model was fit using 25 out of the 73 semantic coherence and linguistic complexity features from the 109 subject responses to predict the SSPA scores. Correlation Coefficient = **0.752**, Mean Absolute Error = **0.330**, Root Mean Square (RMS) Error = **0.405**

these words convey more information than *function words*, e.g. prepositions, conjunctions, interjections, etc. We make use of the Stanford tagger [31] to compute POS tags to determine the number of function words (FUNC) and total words (W) and measure  $\text{FUNC}/w$ , which represents an inverse of the lexical density. A related, more granular measure is the proportion of interjections (UH) to the total words, which is given by  $\text{UH}/w$ . The mean length of sentence (MLS) is another easily computed measure which we expect to be lower for clinical subjects when compared with healthy controls. Finally, we considered parse tree statistics, computed using the Stanford Parser [32]. This includes the parse tree height and Yngve depth scores (mean, total, and maximum), a measure of embedded clause usage [33].

## 4. Results & Discussion

We first sought to determine a subset of language features (described in Section 3) from which we can accurately model the clinical SSPA scores. A total of 73 features were considered: 63 semantic features (7 statistical features  $\times$  3 sentence embedding types  $\times$  3 scenes) and 10 linguistic complexity features computed over all three scenes concatenated. Next, we aim to determine the predictive power of the selected subset of these features in separating the groups of interest (i.e. Sz/Sza, bipolar I disorder, and healthy control subjects). The regression and classification models built with these features were designed and tested using WEKA [34]. It is important to note that the SSPA itself is correlated to the clinical diagnosis and has been effective in differentiating groups of interest [23]. As a result, we note that using it to select features may result in overly-optimistic classification performance for the clinical vs. healthy control and Sz/Sza vs. bipolar disorder classification problems. However, due to the relative dearth of available data in this area, we performed this analysis on the same dataset.

Table 1: Selected features to model SSPA scores with a linear regression model, including ranking of overall importance for each feature. Italicized features were included in both the 25 feature and 15 feature classification problems.

Category	Features	Rank
Semantic Coherence	<i>BoW mean scene 3</i>	1
	<i>INF minimum scene 3</i>	2
	<i>SIF 90<sup>th</sup> percentile scene 3</i>	5
	<i>INF maximum scene 2</i>	7
	<i>INF median scene 3</i>	8
	<i>BoW median scene 3</i>	9
	<i>BoW minimum scene 2</i>	10
	<i>BoW st. dev. scene 2</i>	11
	<i>BoW maximum scene 3</i>	12
	<i>INF st. dev. scene 3</i>	13
	<i>BoW maximum scene 2</i>	18
	<i>BoW 90<sup>th</sup> percentile scene 2</i>	19
	<i>BoW st. dev. scene 3</i>	20
	<i>BoW 90<sup>th</sup> percentile scene 3</i>	21
Lexical Diversity	<i>MATTR</i>	3
	<i>Brunét's index</i>	4
	<i>Honoré's statistic</i>	25
	<i>FUNC/W</i>	6
	<i>UH/W</i>	14
Syntactic Complexity	<i>Maximum Yngve depth</i>	15
	<i>Mean length sent. (MLS)</i>	16
	<i>Parse tree height</i>	17

### 4.1. Modeling SSPA Performance

We use a greedy stepwise search (with linear regression) through the feature space to determine the optimal subset of the features which accurately model the SSPA scores for all 109 subjects without considering the group variable. We down-selected to a set of 25 computed features out of the original 73. These are briefly summarized in Table 1, and the resulting regression model (evaluated using leave-one-out) is shown in Figure 1. We notice that several of the coherence statistics for Scene 3 (negotiation with landlord) are particularly influential when tracking the assigned SSPA score with this model. Interestingly, the top three coherence statistics include a bag-of-words average of *word2vec* vectors (BoW mean scene 3), an *InferSent* sentence encoding (INF minimum scene 3), and a SIF embedding (SIF 90<sup>th</sup> percentile scene 3), indicating a variety of embeddings and range of statistics all provide useful information in predicting SSPA performance. We also note that a variety of lexical diversity (MATTR, Brunét's index), lexical density ( $\text{FUNC}/w$ ,  $\text{UH}/w$ ) and syntactic complexity (maximum Yngve depth) measures are among the most influential, confirming the benefit of a complementary set of language measures.

### 4.2. Identification of Schizophrenia and Bipolar Disorder

Next, we aim to determine the ability of this subset of language features to correctly predict which subjects fall into the groups of interest. We performed two separate classification tasks: (1) separation of the clinical and healthy control groups, (2) separation within the clinical group between Sz/Sza subjects and bipolar I subjects. Both a logistic regression (LR) and a naïve Bayes (NB) classifier were trained in each case using leave-one-out cross validation to determine model parameters and performance. Then, we further down-selected this set to a group of 15




Figure 2: Selected receiver operating characteristic (ROC) curves for both binary classification tasks. For clinical vs control classification, TPR indicates correctly classifying a clinical subject and FPR indicates falsely classifying a control subject as clinical. For Sz/Sza vs bipolar classification, TPR is correctly classifying an Sz/Sza subject and FPR is falsely classifying a bipolar subject as Sz/Sza.

features and re-evaluated the performance of both classifiers.

The confusion matrices for the clinical and control group classification task are shown in Table 2a. As we can see, LR with all 25 selected features works best, with the area under curve (AUC) in the ROC plot being 0.960 (see Figure 2). In this case, 78 of 87 (89.7%) clinical subjects and 19 of 22 (86.7%) healthy controls were correctly identified in our leave-one-out evaluation. We also see comparable performance for the NB and LR models when the feature set is reduced to only the top 15 features that model SSPA scores, though AUC is lower than both models with 25 features.

Next, we consider a classification problem within the group of 87 clinical subjects, of which 43 are diagnosed with Sz/Sza and 44 are diagnosed with bipolar I disorder. We use the same feature subsets and same binary classifier models as in the previous task, trained and evaluated using leave-one-out cross-validation. From the confusion matrices in Table 2b, we see that NB performs better than LR when either a 25 feature or 15 feature subset are used, with the best AUC = 0.826 for NB with 25 features. The ROC curve for a 25-feature NB classifier is shown in Figure 2. Interestingly, LR with 25 features had the lowest performance on this task (AUC = 0.700).

LR typically performs better than NB when more data is available for training [35]; however in clinical applications data set size is often limited. This makes sense with respect to our study, as the dataset used in the Sz/Sza vs. bipolar I classification problem is smaller than the dataset used in the clinical vs control group classification problem. In this case, the LR model is prone to overfitting, as is evident by the fact that performance improves when the feature dimension is reduced. As expected, the classifier performance is considerably worse than the clinical and control group classification problems, as

Table 2: Confusion matrices for binary classification results with logistic regression (LR) and naïve Bayes (NB) classifiers with a 25 feature and 15 feature subset. (a) For clinical vs control classification, LR with 25 features works best at differentiating groups. (b) For Sz/Sza vs bipolar classification, LR using a 25 feature subset works poorly. NB provides more consistent results, even when the feature set is reduced.

(a) Clinical vs Control				(b) Sz/Sza vs Bip.					
Log. Reg.	25 feat.	Clinical Control	True group:	Clinical	Control	Sz/Sza Bipolar	True group:	Sz/Sza	Bipolar
				78	3			30	14
				9	19			13	30
			AUC = 0.960				AUC = 0.700		
15 feat.	Clinical Control	Clinical Control		79	10	Sz/Sza Bipolar		30	10
				8	12			13	34
			AUC = 0.882				AUC = 0.796		
N. Bayes	25 feat.	Clinical Control	True group:	Clinical	Control	Sz/Sza Bipolar	True group:	Sz/Sza	Bipolar
				73	2			30	11
				14	20			13	33
			AUC = 0.908				AUC = 0.826		
15 feat.	Clinical Control	Clinical Control		76	5	Sz/Sza Bipolar		31	11
				11	17			12	33
			AUC = 0.873				AUC = 0.803		

the language differences between schizophrenia and bipolar patients are more difficult to distinguish, even for experienced clinicians. Considering this fact, we still see reasonable performance with only computed language measures and no additional clinical assessment.

## 5. Conclusion

This paper demonstrates the potential of computational linguistics to aid neuropsychiatric practice in the clinic. We believe it is critically important to tie computational methods to established clinical practice in order to bridge the gap between the latest developments in NLP, which motivated our feature selection using SSPA. Still, there are many directions in which we can take future work. The sentence embedding and coherence metrics computed in this study are by no means an exhaustive list of potential methods, and it is likely a more optimal easily computable feature set exists to model SSPA performance and classify groups of interest. In particular, we are interested in finding a more concise group of clinically relevant language features with which we can perform this analysis. Additionally, we can look at more language metrics within each subject group to further subtype and cluster individuals within each group based on language metrics. These methods can also be applied to clinical assessments beyond the SSPA tasks and for a wider variety of psychiatric conditions. Lastly, we would like to examine how classification and modeling of clinical test scores changes when computed features are used in conjunction with other clinical tests to model task performance and classification of groups.

## 6. Acknowledgment

This work was partially funded by a grant from the Boehringer Ingelheim International GmbH to ASU (PI: Berisha).

## 7. References

- [1] Center for Behavioral Health Statistics and Quality, “2016 national survey on drug use and health: Methodological summary and definitions,” *Substance Abuse and Mental Health Services Administration, Rockville, MD*, 2017.
- [2] P. R. Desai, K. A. Lawson, J. C. Barner, and K. L. Rascati, “Estimating the direct and indirect costs for community-dwelling patients with schizophrenia: Schizophrenia-related costs for community-dwellers,” *Journal of Pharmaceutical Health Services Research*, vol. 4, no. 4, pp. 187–194, Dec. 2013.
- [3] A. C. Altamura and J. M. Goikolea, “Differential diagnoses and management strategies in patients with schizophrenia and bipolar disorder,” *Neuropsychiatric Disease and Treatment*, vol. 4, no. 1, pp. 311–317, Feb. 2008.
- [4] American Psychiatric Association, *Diagnostic and Statistical Manual of Mental Disorders: DSM-5*, 5th ed. Arlington, VA: American Psychiatric Publishing, 2013.
- [5] P. R. Montague, R. J. Dolan, K. J. Friston, and P. Dayan, “Computational Psychiatry,” *Trends in Cognitive Sciences*, vol. 16, no. 1, pp. 72–80, Jan. 2012.
- [6] G. A. Cecchi, V. Gurev, S. J. Heisig, R. Norel, I. Rish, and S. R. Schrecke, “Computing the structure of language for neuropsychiatric evaluation,” *IBM Journal of Research and Development*, vol. 61, no. 2/3, pp. 1:1–1:10, Mar. 2017.
- [7] B. Elvevåg, P. W. Foltz, D. R. Weinberger, and T. E. Goldberg, “Quantifying incoherence in speech: An automated methodology and novel application to schizophrenia,” *Schizophrenia Research*, vol. 93, no. 1-3, pp. 304–316, Jul. 2007.
- [8] G. Bedi, F. Carrillo, G. A. Cecchi, D. F. Slezak, M. Sigman, N. B. Mota, S. Ribeiro, D. C. Javitt, M. Copelli, and C. M. Corcoran, “Automated analysis of free speech predicts psychosis onset in high-risk youths,” *npj Schizophrenia*, vol. 1, p. 15030, 2015.
- [9] C. M. Corcoran, F. Carrillo, D. Fernández-Slezak, G. Bedi, C. Klim, D. C. Javitt, C. E. Bearden, and G. A. Cecchi, “Prediction of psychosis across protocols and risk cohorts using automated language analysis,” *World Psychiatry*, vol. 17, no. 1, pp. 67–75, Feb. 2018.
- [10] D. Iter, J. Yoon, and D. Jurafsky, “Automatic Detection of Incoherent Speech for Diagnosing Schizophrenia,” in *Proceedings of the Fifth Workshop on Computational Linguistics and Clinical Psychology: From Keyboard to Clinic*, 2018, pp. 136–146.
- [11] I. Sekulic, M. Gjurković, and J. Šnajder, “Not Just Depressed: Bipolar Disorder Prediction on Reddit,” in *Proceedings of the 9th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*. Association for Computational Linguistics, Oct. 2018, pp. 72–78.
- [12] N. B. Mota, M. Copelli, and S. Ribeiro, “Thought disorder measured as random speech structure classifies negative symptoms and schizophrenia diagnosis 6 months in advance,” *npj Schizophrenia*, vol. 3, no. 1, Dec. 2017.
- [13] T. K. Landauer and S. T. Dumais, “A solution to Plato’s problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge,” *Psychological Review*, vol. 104, no. 2, pp. 211–240, 1997.
- [14] T. Mikolov, K. Chen, G. Corrado, and J. Dean, “Efficient estimation of word representations in vector space,” *arXiv preprint arXiv:1301.3781*, 2013.
- [15] D. M. Blei, A. Y. Ng, and M. I. Jordan, “Latent dirichlet allocation,” *Journal of Machine Learning Research*, vol. 3, no. Jan, pp. 993–1022, 2003.
- [16] Y. R. Tausczik and J. W. Pennebaker, “The psychological meaning of words: LIWC and computerized text analysis methods,” *Journal of language and social psychology*, vol. 29, no. 1, pp. 24–54, 2010.
- [17] 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, Oct. 2015.
- [18] R. Socher, A. Perelygin, J. Wu, J. Chuang, C. D. Manning, A. Ng, and C. Potts, “Recursive deep models for semantic compositionality over a sentiment treebank,” in *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, 2013, pp. 1631–1642.
- [19] E. S. Kayi, M. Diab, L. Pauselli, M. Compton, and G. Copper-smith, “Predictive Linguistic Features of Schizophrenia,” in *Proceedings of the 6th Joint Conference on Lexical and Computational Semantics (\*SEM 2017)*, 2017, pp. 241–250.
- [20] M. Mitchell, K. Hollingshead, and G. Coppersmith, “Quantifying the language of schizophrenia in social media,” in *Proc. of the 2nd Workshop on Computational Linguistics and Clinical Psychology: From Linguistic Signal to Clinical Reality*, 2015, pp. 11–20.
- [21] T. L. Patterson, S. Moscona, C. L. McKibbin, K. Davidson, and D. V. Jeste, “Social skills performance assessment among older patients with schizophrenia,” *Schizophrenia Research*, vol. 48, no. 2-3, pp. 351–360, Mar. 2001.
- [22] M. Tu, V. Berisha, and J. Liss, “Interpretable Objective Assessment of Dysarthric Speech Based on Deep Neural Networks,” in *Interspeech 2017*. ISCA, Aug. 2017, pp. 1849–1853.
- [23] C. R. Bowie, C. Depp, J. A. McGrath, P. Wolyniec, B. T. Mauts-bach, M. H. Thornquist, J. Luke, T. L. Patterson, P. D. Harvey, and A. E. Pulver, “Prediction of real-world functional disability in chronic mental disorders: A comparison of schizophrenia and bipolar disorder,” *American Journal of Psychiatry*, vol. 167, no. 9, pp. 1116–1124, 2010.
- [24] S. Arora, Y. Liang, and T. Ma, “A Simple but Tough-to-Beat Baseline for Sentence Embeddings,” in *Proc. of 5th International Conference on Learning Representations*, Toulon, France, 2017, p. 16.
- [25] A. Conneau, D. Kiela, H. Schwenk, L. Barrault, and A. Bordes, “Supervised Learning of Universal Sentence Representations from Natural Language Inference Data,” *arXiv:1705.02364 [cs]*, May 2017.
- [26] M. A. Covington and J. D. McFall, “Cutting the Gordian Knot: The Moving-Average Type-Token Ratio (MATTR),” *Journal of Quantitative Linguistics*, vol. 17, no. 2, pp. 94–100, May 2010.
- [27] E. Brunet, *Le Vocabulaire de Jean Giraudoux. Structure et Evolution*. Slatkine, 1978, no. 1.
- [28] A. Honoré, “Some Simple Measures of Richness of Vocabulary,” *Association for Literary and Linguistic Computing Bulletin*, vol. 7, no. 2, pp. 172–177, 1979.
- [29] R. S. Bucks, S. Singh, J. M. Cuerden, and G. K. Wilcock, “Analysis of spontaneous, conversational speech in dementia of Alzheimer type: Evaluation of an objective technique for analysing lexical performance,” *Aphasiology*, vol. 14, no. 1, pp. 71–91, Jan. 2000.
- [30] V. Johansson, “Lexical Diversity and Lexical Density in Speech and Writing: A Developmental Perspective,” *Working Papers in Linguistics*, vol. 53, pp. 61–79, 2009.
- [31] K. Toutanova, D. Klein, C. D. Manning, and Y. Singer, “Feature-rich part-of-speech tagging with a cyclic dependency network,” in *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology - NAACL ’03*, vol. 1. Edmonton, Canada: Association for Computational Linguistics, 2003, pp. 173–180.
- [32] R. Socher, J. Bauer, C. D. Manning, and A. Y. Ng, “Parsing With Compositional Vector Grammars,” in *In Proceedings of the ACL Conference*, 2013.
- [33] V. H. Yngve, “A Model and an Hypothesis for Language Structure,” *Proceedings of the American Philosophical Society*, vol. 104, no. 5, 1960.
- [34] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten, “The WEKA data mining software: An update,” *ACM SIGKDD Explorations Newsletter*, vol. 11, no. 1, pp. 10–18, 2009.
- [35] A. Y. Ng and M. I. Jordan, “On Discriminative vs. Generative Classifiers: A Comparison of Logistic Regression and Naïve Bayes,” in *Advances in Neural Information Processing Systems*, 2002, pp. 841–848.