

Utterance Classification with Logical Neural Network: Explainable AI for Mental Disorder Diagnosis

Yeldar Toleubay IBM Research, Japan yeldar.toleubay@gmail.com	Don Joven Agravante IBM Research, Japan don.joven.r.agravante@ibm.com	Daiki Kimura IBM Research, Japan daiki@jp.ibm.com
Baihan Lin Columbia University, USA baihan.lin@columbia.edu	Djallel Bouneffouf IBM Research, USA djallel.bouneffouf@ibm.com	Michiaki Tatsubori IBM Research, Japan mich@jp.ibm.com

Abstract

In response to the global challenge of mental health problems, we propose a Logical Neural Network (LNN) based Neuro-Symbolic AI method for the diagnosis of mental disorders. Due to the lack of effective therapy coverage for mental disorders, there is a need for an AI solution that can assist therapists with the diagnosis. However, current Neural Network models lack explainability and may not be trusted by therapists. The LNN is a Recurrent Neural Network architecture that combines the learning capabilities of neural networks with the reasoning capabilities of classical logic-based AI. The proposed system uses input predicates from clinical interviews to output a mental disorder class, and different predicate pruning techniques are used to achieve scalability and higher scores. In addition, we provide an insight extraction method to aid therapists with their diagnosis. The proposed system addresses the lack of explainability of current Neural Network models and provides a more trustworthy solution for mental disorder diagnosis.

1 Introduction

A mental disorder is a significant deterioration of human thinking, emotional control, or behavior, which is diagnosed clinically and can affect key areas of life. Due to the COVID-19 pandemic, the number of people who suffer from anxiety and depressive illnesses greatly increased in 2020. Initial projections indicate a 26% and 28% increase in anxiety and major depressive disorders respectively during the first year of the pandemic (who). Moreover, every year, 703 000 people commit suicide, with many more attempting to do so. Although people of all ages commit suicides, it is alarming that in 2019 suicide was one of the leading causes of death among young people worldwide (Sui). Furthermore, around 24 million people, or 1 in 300 persons (0.32%), globally suffer from schizophrenia. Although it is not common as other mental

disorders, schizophrenia produces psychosis, is associated with significant disability, and may have an impact on all aspects of life, including personal, family, social, educational, and occupational functioning (Sch).

Diagnosis of mental disorders is accomplished through a clinical interview, where a therapist evaluates the mental health of the patient and identifies possible disorders based on symptoms. However, although many mental health issues may be properly treated at low cost, there is still a wide gap between those who need care and those who have access to it. Despite the progress in some countries, there is still a severe lack of effective therapy coverage. Therefore, there is a need for an AI solution that can assist therapists with a diagnosis of mental disorders.

Although current Neural Network (NN) models are powerful and can operate in a wide range of tasks, their effectiveness in mental disorder classification is questionable due to their black-box nature. In this regard, the model explainability is a vital property, which is required to make a diagnosis of mental disorders. While Neural Network models can achieve high scores, therapists may be hesitant to trust such tools and accept classification results if proper explanations are not provided. Because of NN nature, it is impossible to tell whether their predictions are the result of robust features or some spurious clues (Ribeiro et al., 2020). There are attempts to provide interpretable insights in mental disorder diagnosis, such as using topic modeling to extract concepts (Lin et al., 2023a) or inferring psychological properties such as working alliance (Lin et al., 2023b). Although such approaches can enable explainable AI systems for passive assistance (Lin et al., 2023c; Lin, 2022b) or interventional recommendations (Lin et al., 2023e,d) to the therapists, applying these insights directly to the classification problem yields suboptimal performance (Lin et al., 2022). Furthermore, despite

being able to provide global explanations for the prediction (Mowery et al., 2017), traditional ML models lack scalability and they are not generalizable for broader tasks.

In this regard, Logical Neural Network (Riegel et al., 2020) might be a good solution to the problem. It is a Neuro-Symbolic AI method (NSAI), that combines the learning capabilities of neural networks with the reasoning capabilities of classical logic-based AI. The LNN is a Recurrent neural network architecture in which neurons represent a precisely defined notion of weighted real-valued logic. It has a 1-to-1 relationship to a system of logical formulae. The main problem related to this approach is that it has not been implemented for the supervised learning utterance classification task. Therefore, this work proposes an LNN-based explainable NSAI utterance classification method for mental disorder diagnosis. The model was trained with different predicate pruning techniques to achieve scalability and higher scores. The advantages of the proposed system can be summarized via the following points:

- We propose design of the supervised NSAI method for utterance classification task, where input to the model is predicates from clinical interviews and output is a mental disorder class. After the training the system outputs weighted logical rule to make classifications.
- We propose a predicate pruning methods to improve scalability and generalizability of the model.
- We propose an insight extraction methods which can aid therapists with their mental disorder diagnosis.

1) This paper is organized as follows: Section II details the proposed system, Section III contains experiment results , Section IV provides discussions and future work, and the paper ends with a Conclusion.

2 Supervised learning with LNN

Although NSAI supports data driven training of the network, it encodes knowledge into logic rules with predicates as inputs, where predicates represents a property or a relation. Therefore, NSAI method requires special preprocessing of the dataset to generate predicates and data samples for training and testing purposes. The proposed system consists of

two parts Abstract Meaning Representation (AMR) (Zhou et al., 2021) semantic parser and LNN. Fig. 1 shows overall pipeline of the system, first part containing AMR parser is used to convert raw text into classifier input data, and second part is an LNN model which performs rule-based classification.

2.1 Dataset preparation and preprocessing

Counseling and Psychotherapy Transcripts (ale) is a unique and fully anonymized online series of clinical interviews that allows students and researchers to dive deeply into the patient-therapist relationship and track the progress and setback of patients over multiple therapy sessions. These materials bring the mental disorder diagnosis process to life and provide unprecedented levels of access to the widest possible range of clients. Therefore, transcripts of 4 types of mental disorders, which are anxiety, depression, suicidal thoughts, and schizophrenia, from this dataset are used in our training and evaluation of the model. Table 1 shows the details of the dataset; in our simulations, only 12 sessions of clinical interviews have been used due to the computational constraints of semantic parsing. An example from the transcript has shown in Fig 2. In our experiments, a transcript is a full clinical interview between a patient and a therapist, while an utterance represents a full response of the patient to a specific question from the therapist.

Class	Number of total sessions	Number of used session
<i>Anxiety</i>	498	12
<i>Depression</i>	377	12
<i>Suicidal</i>	12	12
<i>Schizophrenia</i>	71	12

Table 1: Details of the dataset.

As mentioned before, LNN requires a special data structure to function. AMR parser is used for generation of predicates by extracting the semantics of the utterance and converting semantics into a graph, where nodes (keys) represent concepts and edges (values) represent relations to concepts. Example of AMR Representation is shown in Fig 1. AMR Representation keys and values are combined to generate predicates as shown in Table 2.

Moreover, a training and a testing sample is input to the model and is obtained by using AMR parser over an utterance. Furthermore, a sample contains all predicates that has been mined from dataset and

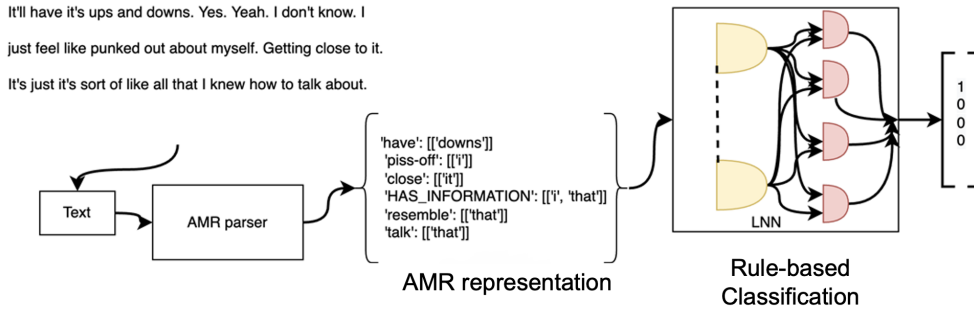


Figure 1: The overview of the proposed system.

THERAPIST: So, how're you doing?
 PATIENT: I'm all right.
 THERAPIST: So, what's new?
 PATIENT: Please take me off the Prozac.
 THERAPIST: It's not working, or you
 having problems with it?
 PATIENT: It is bad.
 THERAPIST: Oh, no. What are you
 noticing?
 PATIENT: Like, I don't even wanna get
 out of bed. I don't even wanna do my
 house.

Figure 2: Examples of a dataset transcript.

the corresponding output of parser as groundings. The values of the groundings are assigned according to the presence of the particular predicates in the parsed utterance, which means only predicates that results from that particular utterance assigned with *TRUE* grounding for that particular sample. In this regard, certain combinations of predicates might repeat over multiple classes and the proposed design takes into account this issue.

2.2 Proposed system details

LNN is a core of the model, which has only few differences from regular neural network. The main difference of LNN is that its neural parameters are limited such that the truth functions of the relevant logical gates govern the behavior of the neurons. Moreover, LNN neuron has more parameters compared to dense neuron, since it keeps both upper and lower bounds to the corresponding subformula or predicate.

The proposed LNN architecture has 4 *AND* logic gates that act as binary classifiers for each mental disorder class. Predicates are inputs to the logic gates, while model is trained by samples generated from utterances. Those samples show truth values for formulae. After the training model outputs set of weight for each predicate and outputs a tensor

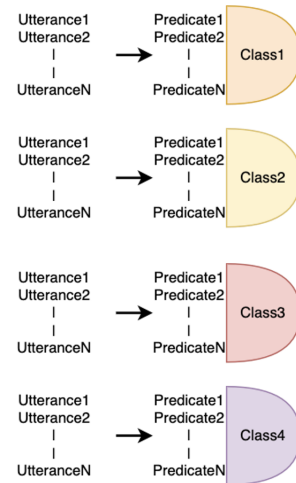


Figure 3: Proposed LNN architecture for mental disorder diagnosis.

of lower and upper bounds as a score for a particular input. In our experiments the each logic gate is evaluated as a binary classifier that classifies according to some threshold, thus the upper and lower bounds are averaged to obtain a single score. The score S for each class is obtained via following equation:

$$S = w_1 P_1(x_1) + w_2 P_2(x_2) + \dots + w_N P_N(x_N) \quad (1)$$

where P is a predicate, w is a weight obtained from training and x_i is a grounding for each predicate in a sample.

The proposed system will be evaluated as a separate binary classification models for each gate by True Positive Rate (TPR) and False Positive Rate (FPR) metrics. The TPR indicates the proportion of all available positive samples that contain correct positive results. In contrast, FPR quantifies the pro-

Input	Predicates			Output
	AMR Representation			
	Keys	Values	Output of Parser	Class
sample0	HAS_POSSESSION	your medication	TRUE	Depression
	HAS_POSSESSION	any details	FALSE	
	HAS_POSSESSION	downs	FALSE	
	HAS_POSSESSION	just awkward thing	FALSE	
	have	your medications	FALSE	
	have	any details	FALSE	
	have	downs	TRUE	
	have	just awkward thing	FALSE	
	talk	your medication	FALSE	
	talk	any details	FALSE	
	talk	downs	FALSE	
	talk	just awkward thing	FALSE	
sample1	HAS_POSSESSION	your medications	TRUE	Anxiety
	HAS_POSSESSION	any details	TRUE	
	HAS_POSSESSION	downs	FALSE	
	HAS_POSSESSION	just awkward thing	TRUE	
	have	your medications	FALSE	
	have	any details	FALSE	
	have	downs	FALSE	
	have	just awkward thing	FALSE	
	talk	your medications	FALSE	
	talk	any details	TRUE	
	talk	downs	FALSE	
	talk	just awkward thing	FALSE	

Table 2: LNN for supervised learning inputs and outputs – predicates, data samples and class

portion of available negative samples that contain incorrectly positive results. Moreover, the receiver operating characteristic (ROC) curve is created by plotting the TPR against the FPR at various threshold values.

$$TPR = \frac{True\ Positives}{True\ Positives + False\ Negatives} \quad (2)$$

$$FPR = \frac{False\ Positives}{True\ Negatives + False\ Positives} \quad (3)$$

2.3 Predicate pruning methods

Predicates play a crucial role in LNN training and can greatly affect the accuracy of the model. Table 4 shows that 48 transcripts result in more than 19000 predicates. However, according to Table 3 a preliminary simulation results show that for a linear

# of predicates	Training time (s)
710	4.49
1415	16.54

Table 3: Results of training time with different number of predicates for an LNN model with 2 Logic gates.

increase in number of predicates, LNN requires exponential increase in training time. Therefore, there is a need for predicate pruning methods, which will help to chose predicates that contribute the most towards the correct diagnosis. Thus, similarity, exclusivity and frequency based predicate pruning methods has been proposed to reduce number of predicates.

Similarity pruning. Simulations has shown that AMR Parser returns multiple variants of values per one key. Often, those values contain repeating phrases. Thus, it is possible to group all those

Class	Original predicates	Similarity pruning	Exclusive Pruning	F=1	F=2	2<F<10	F>9
<i>Anxiety</i>	5529	2773	245	3152	216	150	14
<i>Depression</i>	7227	3532	472	2174	454	133	12
<i>Suicidal</i>	6067	3213	230	2839	197	160	17
<i>Schizophrenia</i>	3746	1914	96	1718	102	87	7

Table 4: Number of original predicates and number of predicates after similarity, exclusive and frequency pruning methods.

lookalike predicates by taking a predicate that contains possible repetitions, e.g. instead of taking both “HAS_POSSESSION_my sister’s birthday” and “HAS_POSSESSION_sister’s birthday”, one can take only the first one.

Frequency pruning. In traditional ML word count can show the importance of some features for a specific class. Using the same logic, it was assumed that predicates that are encountered frequently in sessions will have higher impact on model training. Thus, predicates has been analyzed in terms of repetitions across sessions and have been grouped according to the specified frequencies.

Exclusive pruning. Since transcripts are conversations between patients and therapists, there are many repeating predicates between classes. Thus, it was suggested that predicates belonging only to a class will avoid contradictions in the model as well as will have higher correlation to a specific class. Therefore, predicates repeating between classes predicates that are repeated only once have been removed.

3 Experiment results

In this section experiment results for predicate pruning and LNN model evaluation will be provided. Table 4 shows number of predicates for a particular pruning method. Similarity pruning method prunes almost half of the original predicates. Furthermore, Exclusive and Frequency pruning methods have been applied on top of the similarity pruning method. Results for the Exclusive pruning shows that Depression class has twice of Anxiety and Suicidal predicates and 5-times of Schizophrenia predicates. Moreover, results for the different frequencies show that majority of the predicates (43%) repeat just once, while the higher frequency rates have fewer predicates.

The LNN models have been trained using different pruning methods and have been compared with

Deep Learning (DL) and LNN baselines. The number of predicates and training samples are shown in the Table 5. The LNN models have been trained with supervised loss, which targets the labels with learning rate of 0.05, for 50 epochs. The main difference between LNN models is in the predicates. The details of each model are summarized below:

- *DL baseline.* As a DL baseline pre-trained BERT (Devlin et al., 2018) model and Bert tokenizer with a maximum sequence length of 256 inputs have been selected for finetuning. The model has been trained for 10 epochs using Adam optimizer with a learning rate of 10^{-5} .
- *LNN baseline.* The predicates for the LNN baseline have been selected randomly from Similarity predicates. The number of predicates for each class varies from 340 to 380 predicates.
- *Frequency pruning models.* Several models with different frequencies have been trained to examine the effectiveness of the frequency pruning methods. $F > Threshold$ stands for the model with predicates repeating with a frequency higher than the threshold value. The $F > 5$ balanced ensures that classes are balanced in terms of predicates. The remaining predicates have been chosen from a lower frequency.
- *Exclusive pruning models.* The exclusive pruning method is used in combination with similarity and frequency pruning methods. In the simulations, Frequency pruning prunes predicates that repeat only once. Then the exclusive pruning removes all the repeating predicates between classes.

According to the Table 5, Frequency predicates does not have a significant effect on model performance when they are applied alone, since the are

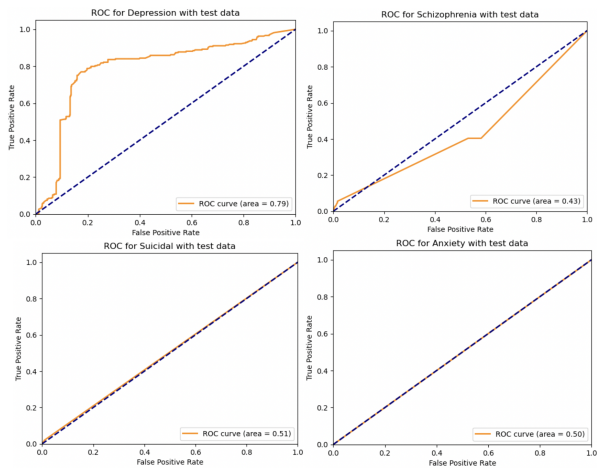


Figure 4: AUC ROC curves for each class in testing.

under (AUC) the ROC curve is around 0.5, which is close to the random classifier. Moreover, LNN baseline with 1000 predicates and 10000 training samples performed surprisingly well for the Anxiety class, achieving AUC of 0.76. Baseline DL model has AUC scores higher than 0.72 for all classes when treated as a binary classifier for each class. However, since therapist cannot use this data explicitly, the accuracy for the multi-class classification has higher importance for this case and DL model can provide only 58% accuracy in such setting. Exclusive predicates model have shown a good performance overall. It reached AUC of 0.79 for the depression class and 0.57 for schizophrenia.

4 Discussions and Future Work

Scaling of the LNN is a significant issue which requires selection of the right predicates. Pruning of the predicates essentially limits the knowledge base of the LNN, thus it is important to understand the effect of the predicates on model performance. Frequency predicate models have not shown great results, the possible explanation for that behavior can be found in predicate analysis. The analysis shows that predicates with higher frequencies also tend to be inclusive for several classes. Such predicates might be extracted from common dialogue phrases, that are common to regular conversations. Thus, it is more difficult to learn for LNN in such circumstances and it might lead to a behavior similar to the random classifiers'. Moreover, variation in the frequency thresholds did not affect the overall performance of the LNN model. Thus, it can be concluded that frequency predicates cannot provide a quality selection of the predicates when they are

applied alone. Furthermore, in the case of exclusive predicates, the model has learned depression class better than others. It can be explained by the depression class possessing more exclusive predicates compared to other classes. Interestingly, the model has learned to identify non-schizophrenia samples better than schizophrenia samples. Possible reasoning for that is fewer predicates for schizophrenia compared to other classes. Furthermore, some mental disorders have the same symptoms, and exclusive pruning eliminate such predicates from the training, which might lead to limited diagnostic abilities. Thus, exclusive predicates should be combined with other methods to provide trade-off between generalization and exclusivity of predicates.

Another challenge of this line of work is the usage of AI for mental disorder diagnosis. As pointed out in (Lin, 2022a), one significant challenge is related to the privacy and security of patient data. To train the model, the system requires access to sensitive patient data, which must be protected from unauthorized access or misuse. There is also a concern that the use of AI in mental health diagnosis may lead to the stigmatization of individuals with mental disorders. In this work, we have de-identified all the sessions and all the transcripts are obtained under proper license and consent. We would also like to point out that the system may not work for all individuals, which could lead to misdiagnosis or lack of diagnosis, leading to harm to the patient. Therefore, the ethical challenge lies in ensuring the system's reliability, fairness, and transparency and balancing the use of AI with the need for human involvement in mental health diagnosis and treatment, as part of the future work.

The main advantage of the LNN over DL is in its explainability. It is possible to extract predicates with high weights for the each class and to examine which predicates contribute to the result significantly. Table 6 shows the predicate semantics analysis for each class after the training. Predicates of depression and anxiety suicidal classes are mostly related to the first-person and third-person actions respectively, while people with anxiety tend to talk about feelings more. In addition, predicates of the schizophrenia class tend to relate to the medical terms. This overlaps with overall content of the transcripts and predicates that possess high weights can be used to give insights to therapist during the diagnosis of the patients.

	# of training samples	# of predicates	Suicidal (AUC)	Depression (AUC)	Anxiety (AUC)	Schizophrenia (AUC)
<i>BaselineLNN</i>	10000	1000	0.50	0.50	0.76	0.52
<i>BaselineDL</i>	10000	N/A	0.73	0.83	0.81	0.72
			Accuracy for multiclass = 0.58			
<i>F > 5</i>	3947	87	0.55	0.58	0.55	0.52
<i>F > 5. Balanced</i>	3947	141	0.55	0.59	0.52	0.55
<i>F > 3</i>	3605	349	0.54	0.56	0.53	0.52
<i>F > 6</i>	3947	81	0.55	0.53	0.56	0.50
<i>Exclusive predicates and F > 1</i>	3947	981	0.51	0.79	0.50	0.43

Table 5: AUC ROC scores for DL baseline, LNN baseline and proposed pruning methods .

	Grouping of predicates	Top 1 weight	Top 2 weight	Top 3 weight
Depression	Related to first-person actions	Do_i	Come_they	Resemble_what
Anxiety	Related to feelings	get_it	look_it	have-rel-role_my mom
Schizophrenia	Related to medical terms	give_me	resemble_things	do_it
Suicidal	Related to third-person actions	have-manner_sense	put_it	do_everything

Table 6: Analysis of the semantics of the predicates for the each class.

4.1 Future work

Overall, it is evident that predicates are too specific from the number of predicates with a frequency of 1, which might be a possible explanation for the poor performance of the model overall. Therefore, they might require some generalization of the predicates. One of the promising methods for that is to use synonym-based predicates. By using thesaurus dictionaries, it is possible to cluster all the keys and values of the AMR representations and use only one variants for the synonyms. That way, it might be possible to reduce the number of predicates significantly and achieve their generalization.

Another possible way to enhance the model is the explore LNN and DL hybrid approach. By using LNN scores it is possible to train some dense layers with SoftMax to predict classes in the multiclass setting. In a such way it will be more convenient to compare LNN results with DL solutions

while keeping the explainability of the LNN.

5 Conclusion

Mental disorders are a significant issue that is affecting more people every year. Therefore, explainable AI mental disorder diagnosis through utterance classification can aid the therapist in their practice. In this work, a supervised learning setting for the LNN has been proposed to address this issue. Moreover, predicate pruning methods based on the similarity, frequency, and exclusivity of the predicates are analyzed in terms of training performance. Overall, the model trained with exclusive predicates shows the best results among the pruning methods, and achieved AUC ROC of 0.79 for the depression disorder. Finally, explainability of the LNN diagnosis has been shown by analyzing significant predicates for each class and extracting the predicates with high weights.

References

Counseling and psychotherapy transcripts series.

Mental health and covid-19: Early evidence of the pandemic's impact: Scientific brief.

Schizophrenia - key facts.

Suicide - key facts.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [Bert: Pre-training of deep bidirectional transformers for language understanding](#).

Baihan Lin. 2022a. Computational inference in cognitive science: Operational, societal and ethical considerations. *arXiv preprint arXiv:2210.13526*.

Baihan Lin. 2022b. Voice2alliance: automatic speaker diarization and quality assurance of conversational alignment. In *INTERSPEECH*.

Baihan Lin, Djallel Bouneffouf, Guillermo Cecchi, and Ravi Tejwani. 2023a. Neural topic modeling of psychotherapy sessions. In *International Workshop on Health Intelligence*. Springer.

Baihan Lin, Guillermo Cecchi, and Djallel Bouneffouf. 2022. Working alliance transformer for psychotherapy dialogue classification. *arXiv preprint arXiv:2210.15603*.

Baihan Lin, Guillermo Cecchi, and Djallel Bouneffouf. 2023b. Deep annotation of therapeutic working alliance in psychotherapy. In *International Workshop on Health Intelligence*. Springer.

Baihan Lin, Guillermo Cecchi, and Djallel Bouneffouf. 2023c. Helping therapists with nlp-annotated recommendation. In *Joint Proceedings of the ACM IUI Workshops*.

Baihan Lin, Guillermo Cecchi, and Djallel Bouneffouf. 2023d. Psychotherapy AI companion with reinforcement learning recommendations and interpretable policy dynamics. In *Proceedings of the Web Conference 2023*.

Baihan Lin, Guillermo Cecchi, and Djallel Bouneffouf. 2023e. SupervisorBot: NLP-Annotated Real-Time Recommendations of Psychotherapy Treatment Strategies with Deep Reinforcement Learning. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*. International Joint Conferences on Artificial Intelligence Organization.

Danielle L. Mowery, Hilary Smith, Tyler Cheney, Greg Stoddard, Glen A. Coppersmith, Craig J. Bryan, and Mike Conway. 2017. Understanding depressive symptoms and psychosocial stressors on twitter: A corpus-based study. *Journal of Medical Internet Research*, 19.

Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. [Beyond accuracy: Behavioral testing of nlp models with checklist](#).

Ryan Riegel, Alexander Gray, Francois Luus, Naweed Khan, Ndivhuwo Makondo, Ismail Yunus Akhalwaya, Haifeng Qian, Ronald Fagin, Francisco Barahona, Udit Sharma, Shajith Ikbal, Hima Karanam, Sumit Neelam, Ankita Likhyan, and Santosh Srivastava. 2020. [Logical neural networks](#).

Jiawei Zhou, Tahira Naseem, Ramón Fernandez As-tudillo, Young-Suk Lee, Radu Florian, and Salim Roukos. 2021. [Structure-aware fine-tuning of sequence-to-sequence transformers for transition-based AMR parsing](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6279–6290, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.