



Exploiting Multi-Modal Features From Pre-trained Networks for Alzheimer’s Dementia Recognition

Junghyun Koo¹, Jie Hwan Lee¹, Jaewoo Pyo², Yujin Jo³, Kyogu Lee¹

¹Music & Audio Research Group (MARG), Seoul National University

²Electrical and Computer Engineering, Seoul National University

³College of Liberal Studies, Seoul National University

{dg22302, wiswisbus, jwpyo, tera-yujin, kglee}@snu.ac.kr

Abstract

Collecting and accessing a large amount of medical data is very time-consuming and laborious, not only because it is difficult to find specific patients but also because it is required to resolve the confidentiality of a patient’s medical records. On the other hand, there are deep learning models, trained on easily collectible, large scale datasets such as Youtube or Wikipedia, offering useful representations. It could therefore be very advantageous to utilize the features from these pre-trained networks for handling a small amount of data at hand. In this work, we exploit various multi-modal features extracted from pre-trained networks to recognize Alzheimer’s Dementia using a neural network, with a small dataset provided by the ADReSS Challenge at INTERSPEECH 2020. The challenge regards to discern patients suspicious of Alzheimer’s Dementia by providing acoustic and textual data. With the multi-modal features, we modify a Convolutional Recurrent Neural Network based structure to perform classification and regression tasks simultaneously and is capable of computing conversations with variable lengths. Our test results surpass baseline’s accuracy by 18.75%, and our validation result for the regression task shows the possibility of classifying 4 classes of cognitive impairment with an accuracy of 78.70%.

Index Terms: Multimodal Systems, Cognitive Decline Detection, Pre-trained Model

1. Introduction

Collecting a sufficient amount of electronic health records is a challenging task with various factors [1, 2]. Due to this problem, researchers in the medical field are often provided with only a small amount of data given. Owing to the fact that deep learning techniques perform better on large amounts of data, a number of studies using machine learning techniques have been conducted to solve specific medical problems, regarding a limited number of data [3, 4]. Dementia is also one of many medical symptoms facing this situation.

Dementia, a syndrome in which there is deterioration in cognitive function beyond what might be expected from normal ageing, is mostly affected by Alzheimer’s Disease [5]. There were previous researches with various approaches to recognize Alzheimer’s Dementia [6, 7, 8, 9], which has shown excellent performance. However, datasets used in these works were sufficient with quantity than the one used in this paper.

The ADReSS challenge [10] at INTERSPEECH 2020 hosts two tasks: Alzheimer’s Dementia (AD) classification and Mini Mental Status Examination (MMSE) regression, while providing a refined dataset. The dataset is equally balanced of AD and non-AD participants with the metadata of age and gender. Each

data is a conversation in which participants, in both audio and text modalities, spontaneously describes the picture given by the investigator. Participants of the challenge are suggested to solve hosted tasks using only the given data, where the numbers of train and test data are 108 and 48, respectively.

For recognizing AD with small amounts of data, we determined it would be beneficial to use both acoustic and textual features. Furthermore, we leverage models pre-trained on large scale datasets as feature extractor to get better representation. To this end, this paper focus on exploiting various multi-modal features, and design suitable network architecture. We compare 3 and 4 different acoustic and textual features, respectively, and use the hand-crafted (HC) feature and part-of-speech (POS) tagging as additional inputs. The usage of POS and HC is influenced by previous research, which has approved that using these features gained from transcript can improve the performance [8]. The proposed network is a modified version of Convolutional Recurrent Neural Network (CRNN); capable of computing conversations with variable lengths, and implemented with methods to fit with a small amount of data. Also, the model is able to compute using the acoustic feature only, without any metadata, which can be efficient considering the real-world situation. Our experimental results show using features of the pre-trained network leads to performance gain than that of raw, and regression results imply the potential of network classifying classes of cognitive impairment based on MMSE score.

2. Multi-Modal Features

This work compares 3 different acoustic and 4 different textual features. To obtain a speech signal corresponding to each utterance in the transcription, alignment of the transcription and the signal is done by using [11, 12]. Hence, the following multi-modal feature extraction in this section is applied to the aligned data.

2.1. Acoustic Features

- *openSMILE features:* The openSMILE v2.3 toolkit [13] provides multiple features from raw audio files. From the toolkit, we use the ComParE feature [14] and the eGeMAPS feature [15]. For ComParE feature, using one-way ANOVA, we select 393 features ($p \leq 0.05$), out of 6,373 concerning the efficiency of model capacity.
- *VGGish:* We use VGGish [16] which is trained with *Audio Set* [17] for audio classification. The feature is composed of 128 feature dimensions, where each feature is extracted from audio with a length of 960ms. To handle different lengths of utterance, we use the average value of the extracted VGGish features.

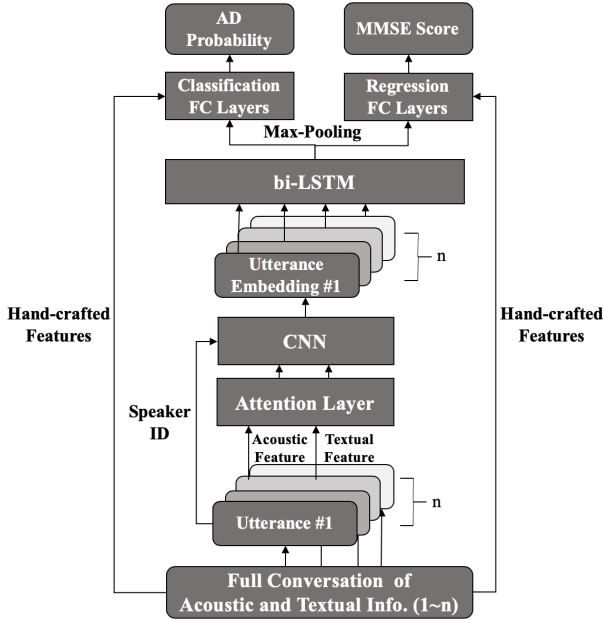


Figure 1: Overview of the proposed method. The acoustic and textual features extracted from each utterance are fed in to the CRNN network. Then, the hand-crafted features retrieved from the participant’s entire conversation are concatenated to the utterance-level features. Finally, the FC layer of each task estimates the AD probability and MMSE score of the participant.

2.2. Textual Features

- *Pre-trained language model features:* We exploit transformer [18] based language models, GPT [19], RoBERTa [20], and Transformer-XL [21]. Pre-trained on large corpora, these language models have shown the effectiveness to improve performance over a wide range of natural language processing tasks. Sentence representations are obtained by averaging word embeddings via [22]. The specific settings for the language models are as follows, GPT: openai-gpt, RoBERTa: roberta-base, Transformer-XL: transfo-xl-wt103. The feature dimensions of GPT and RoBERTa is 768, and Transformer-XL dimensions of 1024. Besides the aforementioned features, we also use 300-dimensional GloVe vectors [23].
- *Hand-crafted features:* We integrate three categories, psycholinguistic, repetitiveness, and lexical complexity features, as HC features, which reflect the features of Alzheimer’s. Psycholinguistic features and repetitiveness¹ are that suggested by [6], and lexical complexity is the Lexical Complexity Analyzer for Academic Writing (LCA-AW)². These token-level HC features are aggregated to the conversational-level by only averaging participant’s utterance. We select and use 23 features whose p-value from one-way ANOVA is less than 0.05 amongst a total of 42 features.

3. Proposed Method

While the proposed model can cope with additional inputs such as visual modality, the ADRess challenge only offers acoustic

¹https://github.com/vmasrani/dementia_classifier

²<https://github.com/Maryam-Nasseri/LCA-AW-Lexical-Complexity-Analyzer-for-Academic-Writing>

and textual modalities. Thus, we primarily focus on the network with bimodal inputs. The overview of our model is as Figure 1. In case of unimodal, the network has the same structure, except that only a single modality feature is input.

3.1. Input

An input dialogue consists of its utterances and an extracted HC feature. Each utterance comes along with an acoustic and a textual feature, and a speaker index. The speaker index is a binary feature denoting an investigator or a participant, where it is extended as the size of the largest size of input feature dimension, 1024 in our case, by a single fully connected layer. Input features smaller than 1024 are also expanded the same way by a fully connected layer.

We apply dropout [24] to the input features before they are inserted into the network. This way, the model can be provided with more opportunity to learn independent representations, because each dimension can convey significant information, especially for the features extracted from pre-trained models.

3.2. Model Architecture

The proposed network is a modified version of CRNN, where an attention layer is a forefront layer of the network, and fully connected layers followed after the recurrent layer. Here, we use a bidirectional Long Short-Term Memory Network (bi-LSTM) [25] as the recurrent network.

Each modality input is individually inserted and computed through an attention layer. Our attention layer is implemented as the *Scaled Dot-Product Attention* mechanism introduced in [18]. We use a self-attention mechanism, where an individual feature is used as a query, key, and value during the attentional computation.

Outputs of the attention layer and embedded speaker index of a single utterance are channel-wise concatenated then inserted into the one-dimensional Convolutional Neural Network (CNN). After a convolutional layer expands channel dimension to 32, 6 *Squeeze-and-Excitation* (SE) [26] blocks are followed in the CNN. Each SE block consists of 2 convolution layers with a SE layer in between them. The last convolutional layer of every 2 SE blocks reduces feature dimension by convolutional stride factor of 4 and increments channel dimension. The expanding sizes of the channel dimension are 128, 512, 1024 respectively. Ultimately, CNN outputs 1024-dimensional channels with a global max pooled value.

After every utterance from the input dialogue is each computed through the CNN, the processed utterance embeddings are sequentially inputted into the bi-LSTM. The recurrent network consists of 3 bi-LSTM layers with 512 hidden units. Ultimately, the recurrent network outputs the max-pooled state from the results of the last layer’s hidden states and is concatenated with HC.

Three fully connected (FC) layers follow after the bi-LSTM layers. Both the first two FC layers are followed by a rectified linear unit (ReLU) activation and reduce the input dimension by a factor of 4. The last activation function for classification and regression tasks are softmax and sigmoid, respectively. Ground truth MMSE score is scaled from 0 to 1 for regression loss computation.

3.3. Training and Inference

We use different numbers of utterances per batch during the training phase for the network to have opportunities to interpret

Table 1: Validation Results of Acoustic Unimodal Network

Feature	Accuracy	F1	RMSE
<i>eGeMAPS</i>	61.82%	71.98%	6.7178
<i>ComParE</i>	68.27%	74.62%	6.7852
<i>VGGish</i>	85.27%	86.28%	5.1144

various sequences of dialogue. The size is randomly selected between 5 and the minimum number of utterances among the dialogues in each batch. Since the minimum number of utterances of dialogue in the training data is 7, it was reasonable to set the minimum length to 5. If the length is too short, the network could be vulnerable to utterances with less meaningful data such as the investigator’s “okay” or “mhm”. A single batch is used during the inference phase to analyze every utterance in an input dialogue.

Our training loss for classification and regression tasks are binary cross-entropy error and mean squared error, respectively. The total cost function is a summation of these two values. We use the Adam optimizer [27] with a learning rate of 0.0002 and momentum parameters $\beta_1 = 0.5$, $\beta_2 = 0.9$.

4. Experiments

In this section, we evaluate model performances for both classification and regression tasks. Recorded performances are averaged value from measurements of 5-fold cross validation, where each fold contains 86 training and 22 validation conversations, except for the last fold containing 88 training and 20 validation conversations.

Prior experiments were conducted for optimizing several hyperparameters in the proposed network. First, we compared model performance with a one-dimensional convolutional kernel size of 3, 5, 10, 15. Through observations, larger kernel sizes led to performance gain; thus, we set the kernel size to 15. Attempts to ascertain the ideal dropout rate among 0 to 50% at 10% intervals could not be determined. Yet, we adopted a 20% dropout rate for data augmentation and prevention of overfitting. Finally, we discovered using 6 instead of 3 stacked convolutional blocks achieved better performance. Experimental results of each model shown in this section share above achieved hyperparameter values.

4.1. Feature Comparison

4.1.1. Unimodal Network

Table 1 is validation results of unimodal networks using acoustic features. The accuracy using VGGish exceeds openSMILE’s by over 17%, which conveys a significant difference in these audio features on performance. Hence, this result establishes a strong point that using an acoustic feature extracted from a pre-trained network outperforms features extracted from scratch.

Textual feature comparing experiment is further conducted by including combinations of using POS and HC features as input. Upon using POS, it is concatenated to the input textual feature to fed into the network. The best performing features for classification and regression are Transformer-XL and GloVe, respectively, according to Table 2.

4.1.2. Bimodal Network

We choose VGGish as a fixed auditory input feature for the bimodal network, considering its leading validation performance among other audio features. The use of POS and HC features is

Table 2: Validation Results of Textual Unimodal Network

Feature	Accuracy	F1	RMSE	
GloVe	+ None	90.73%	0.9158	3.9282
	+ POS	90.73%	0.9122	3.8959
	+ HC	92.55%	0.9303	3.3493
	+ POS + HC	93.55%	0.9389	3.3650
GPT	+ None	91.55%	0.9224	3.7825
	+ POS	92.55%	0.9303	4.0275
	+ HC	91.55%	0.9246	3.6695
	+ POS + HC	89.82%	0.9076	3.7684
RoBERTa	+ None	92.45%	0.9312	3.7622
	+ POS	91.64%	0.9231	3.8437
	+ HC	93.45%	0.9391	3.3852
	+ POS + HC	93.45%	0.9391	3.3773
Transformer-XL	+ None	92.55%	0.9296	4.0078
	+ POS	93.45%	0.9382	4.0588
	+ HC	94.36%	0.9469	3.4866
	+ POS + HC	92.55%	0.9325	3.6602

Table 3: Validation Results of Bimodal Network

Feature	Accuracy	F1	RMSE	
GloVe	+ None	92.55%	0.9288	4.0743
	+ POS	93.55%	0.9398	3.7091
	+ HC	90.73%	0.9122	3.9138
	+ POS + HC	93.55%	0.9382	3.4989
GPT	+ None	93.45%	0.9398	3.5503
	+ POS	93.45%	0.9398	3.9910
	+ HC	91.64%	0.9231	3.6334
	+ POS + HC	92.55%	0.9318	3.5182
RoBERTa	+ None	91.64%	0.9231	3.7842
	+ POS	92.55%	0.9318	3.6860
	+ HC	93.45%	0.9375	3.4977
	+ POS + HC	92.55%	0.9311	3.5182
Transformer-XL	+ None	91.64%	0.9201	4.0703
	+ POS	92.55%	0.9288	4.0546
	+ HC	90.73%	0.9114	3.7820
	+ POS + HC	94.45%	0.9454	3.6099

performed in the bimodal network as well. Acknowledging the results of unimodal networks, Transformer-XL feature is also well performed in the classification tasks, where RoBERTa feature scores the best root mean squared error (RMSE).

From the observation of this experimental result, it was reasonable to ascertain the best text feature. Notably, using Transformer-XL produced the highest performance in the classification task. Moreover, while comparing the average RMSE scores by feature, RoBERTa outputs the lowest score for both unimodal and bimodal networks. On the other hand, when analyzing performance between additional inputs, only little tendency could be observed. This can be an implication that the quantity of given data may not be sufficient for the additional inputs to exert influence.

4.2. Analysis of Regression Task

Figure 2 illustrates a graph comparing regression outputs from a bimodal network and the actual patient’s corresponding MMSE score during validation stage. The severity classes, each shaded area in the figure, were categorized based on the MMSE score presented in [28].

In this example, VGGish and RoBERTa were used as input features, and the RMSE and r^2 value between network outputs and ground truths are 3.5182 and 0.7361, respectively. The

Table 4: Results of Test Set

Model	Modality	Feature	Classes	Precision	Recall	F1	Accuracy	RMSE
Baseline		ComParE	non-AD	0.67	0.50	0.57	0.625	6.14
			AD	0.60	0.75	0.67		
Ours	Unimodal Network	VGGish	non-AD	0.6897	0.8333	0.7547	0.7292	5.0765
			AD	0.7895	0.6250	0.6977		
		Transformer-XL	non-AD	0.8261	0.7917	0.8085	0.8125	4.0182
	AD	0.8000	0.8333	0.8163				
	Bimodal Network	VGGish + GLoVe	non-AD	0.7407	0.8333	0.7843	0.7708	4.3301
			AD	0.8095	0.7083	0.7556		
VGGish + Transformer-XL		non-AD	0.7500	0.7500	0.7500	0.7500	3.7472	
AD	0.7500	0.7500	0.7500					
Ensembled Output	non-AD	0.7586	0.9167	0.8302	0.8125	3.7749		
	AD	0.8947	0.7083	0.7907				

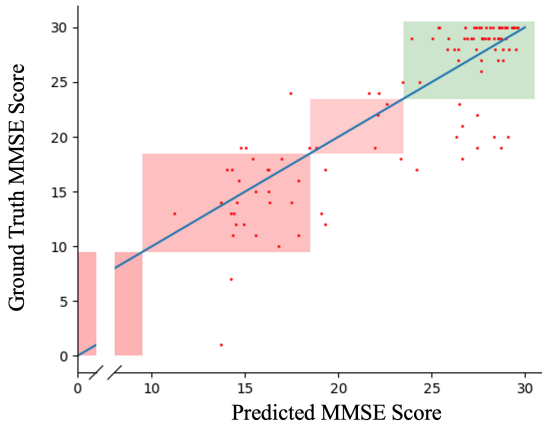


Figure 2: MMSE comparison graph between regression outputs and ground truth scores - The linear line is a representation of the ideal output with zero error for the task. Each rectangular regions represent dementia severity based on the MMSE score and is shaded respectively as classes of normal, mild, moderate, and severe ranging from high to low values.

output plot shows that the distribution of the network output decreases as the MMSE score decreases, which is inferred to follow the distribution of the given training data. Even though there was no network output below a score of 11, 78.70% of the points are included in the shaded area. This indicates that classifying severity classes of dementia is possible to some extent, based on regression outputs.

4.3. Test Set Results

The test dataset of the ADReSS challenge consists of 48 conversations and can be scored with a total of 5 different submissions. Taking this into account, we use two different models for unimodal and bimodal networks each and an ensembled output of bimodal networks to infer the test data. In the case of the unimodal network, VGGish and Transformer-XL are adopted to represent acoustic and textual modality, respectively. For the bimodal network, GloVe and Transformer-XL are adopted as textual modality regarding on their performance from the validation results. Besides, we select models using POS and HC features along with the bimodality inputs. Lastly, the outputs of the top 5 bimodal networks with high validation results are ensembled and used as the final submission.

The final result for each conversation was deduced by five different models with the same configurations used during the training and validation stage. Combining these results, the final output was concluded using majority voting for AD classification and the median value for the MMSE regression task. The

baseline and our test results are presented at Table 4. When using only audio modality, our test accuracy surpasses baseline’s by 10%, where accounting textual modality contributes another 8% performance gain. Although our textual unimodal model performed the best classification result among the single models, our bimodal network’s ensembled output indicates that other bimodal models were able to achieve better performance. Furthermore, the test RMSE implies using both modalities is more advantageous for the regression task.

We could infer from the experimental results that the auditory information led to some performance degradation compared to the textual. This matter can be attributed to the low quality of the audio files provided. In particular, the participant’s voice was barely hearable, while it was clear for the investigator’s comments in some audio files. Even so, with the methodology to infer AD possible with only recorded audio files, the proposed model can be utilized as a real-world application reflecting on the difficulty of acquiring transcriptions and the target’s metadata. The metadata is not dealt with in this work; this is because there was little difference when conditioning age and gender into our model in our prior empirical results.

5. Conclusion

This paper demonstrates extracted features from pre-trained networks are satisfactory for handling small amounts of data, to recognize Alzheimer’s Dementia. The proposed model can compute variable lengths of dialogue and also introduce productive methods to fit the network with a little amount of data. Furthermore, our model does not require any metadata and also can perform well without transcript, which may be practical in real-world situations. Our test result outperforms baseline’s with both tasks, and our regression results imply the potential of network classifying classes of cognitive impairment based on the MMSE score.

For future work, with the expectation of performance gain, mechanisms effectively fusing different modality features [29] [30] can be applied in the model architecture.

6. Acknowledgements

This work was supported partly by Institute for Information & Communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (No.2019-0-01367, BabyMind) and partly by Next-Generation Information Computing Development Program through the National Research Foundation of Korea(NRF) funded by the Ministry of Science and ICT(NRF-2017M3C4A7078548). Also, the authors would like to thank Jeonghyun Yoon, and Ayoung Choi for their fruitful comments, and inspiration.

7. References

- [1] M. Adibuzzaman, P. DeLaurentis, J. Hill, and B. D. Benneyworth, "Big data in healthcare—the promises, challenges and opportunities from a research perspective: a case study with a model database," in *AMIA Annual Symposium Proceedings*, vol. 2017. American Medical Informatics Association, 2017, p. 384.
- [2] T. Edinger, A. M. Cohen, S. Bedrick, K. Ambert, and W. Hersh, "Barriers to retrieving patient information from electronic health record data: failure analysis from the trec medical records track," in *AMIA annual symposium proceedings*, vol. 2012. American Medical Informatics Association, 2012, p. 180.
- [3] M. Chen, X. Shi, Y. Zhang, D. Wu, and M. Guizani, "Deep features learning for medical image analysis with convolutional autoencoder neural network," *IEEE Transactions on Big Data*, 2017.
- [4] T. Shaikhina, D. Lowe, S. Daga, D. Briggs, R. Higgins, and N. Khovanova, "Machine learning for predictive modelling based on small data in biomedical engineering," *IFAC-PapersOnLine*, vol. 48, no. 20, pp. 469–474, 2015.
- [5] A. Burns and S. Iliffe, "Alzheimer's disease." *BMJ*, vol. 338, p. b158, 2009.
- [6] 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.
- [7] S. Karlekar, T. Niu, and M. Bansal, "Detecting linguistic characteristics of alzheimer's dementia by interpreting neural models," *arXiv preprint arXiv:1804.06440*, 2018.
- [8] F. Di Palo and N. Parde, "Enriching neural models with targeted features for dementia detection," *arXiv preprint arXiv:1906.05483*, 2019.
- [9] T. Jo, K. Nho, and A. J. Saykin, "Deep learning in alzheimer's disease: diagnostic classification and prognostic prediction using neuroimaging data," *Frontiers in aging neuroscience*, vol. 11, p. 220, 2019.
- [10] S. Luz, F. Haider, S. de la Fuente, D. Fromm, and B. MacWhinney, "Alzheimer's dementia recognition through spontaneous speech: The ADReSS Challenge," in *Proceedings of INTERSPEECH 2020*, Shanghai, China, 2020. [Online]. Available: <https://arxiv.org/abs/2004.06833>
- [11] F. Schiel, "Automatic phonetic transcription of non-prompted speech," 1999.
- [12] T. Kislser, U. Reichel, and F. Schiel, "Multilingual processing of speech via web services," *Computer Speech & Language*, vol. 45, pp. 326–347, 2017.
- [13] F. Eyben, M. Wöllmer, and B. Schuller, "Opensmile: the munich versatile and fast open-source audio feature extractor," in *Proceedings of the 18th ACM international conference on Multimedia*, 2010, pp. 1459–1462.
- [14] F. Eyben, F. Weninger, F. Gross, and B. Schuller, "Recent developments in opensmile, the munich open-source multimedia feature extractor," in *Proceedings of the 21st ACM international conference on Multimedia*, 2013, pp. 835–838.
- [15] F. Eyben, K. R. Scherer, B. W. Schuller, J. Sundberg, E. André, C. Busso, L. Y. Devillers, J. Epps, P. Laukka, S. S. Narayanan *et al.*, "The geneva minimalistic acoustic parameter set (gemaps) for voice research and affective computing," *IEEE transactions on affective computing*, vol. 7, no. 2, pp. 190–202, 2015.
- [16] S. Hershey, S. Chaudhuri, D. P. Ellis, J. F. Gemmeke, A. Jansen, R. C. Moore, M. Plakal, D. Platt, R. A. Saurous, B. Seybold *et al.*, "Cnn architectures for large-scale audio classification," in *2017 IEEE international conference on acoustics, speech and signal processing (icassp)*. IEEE, 2017, pp. 131–135.
- [17] J. F. Gemmeke, D. P. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, "Audio set: An ontology and human-labeled dataset for audio events," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 776–780.
- [18] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [19] A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever, "Improving language understanding by generative pre-training," URL https://s3-us-west-2.amazonaws.com/openai-assets/researchcovers/languageunsupervised/language_understanding_paper.pdf, 2018.
- [20] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: A robustly optimized bert pretraining approach," *arXiv preprint arXiv:1907.11692*, 2019.
- [21] Z. Dai, Z. Yang, Y. Yang, J. Carbonell, Q. V. Le, and R. Salakhutdinov, "Transformer-xl: Attentive language models beyond a fixed-length context," *arXiv preprint arXiv:1901.02860*, 2019.
- [22] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, and J. Brew, "Huggingface's transformers: State-of-the-art natural language processing," *ArXiv*, vol. abs/1910.03771, 2019.
- [23] J. Pennington, R. Socher, and C. D. Manning, "Glove: Global vectors for word representation," in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 2014, pp. 1532–1543.
- [24] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *The journal of machine learning research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [25] Z. Huang, W. Xu, and K. Yu, "Bidirectional lstm-crf models for sequence tagging," *arXiv preprint arXiv:1508.01991*, 2015.
- [26] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7132–7141.
- [27] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [28] R. M. Crum, J. C. Anthony, S. S. Bassett, and M. F. Folstein, "Population-based norms for the mini-mental state examination by age and educational level," *Jama*, vol. 269, no. 18, pp. 2386–2391, 1993.
- [29] C. Hori, T. Hori, T.-Y. Lee, Z. Zhang, B. Harsham, J. R. Hershey, T. K. Marks, and K. Sumi, "Attention-based multimodal fusion for video description," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 4193–4202.
- [30] T. Yilmaz, A. Yazici, and M. Kitsuregawa, "Non-linear weighted averaging for multimodal information fusion by employing analytical network process," in *Proceedings of the 21st International Conference on Pattern Recognition (ICPR2012)*. IEEE, 2012, pp. 234–237.