



# Cognitive Insights Across Languages: Enhancing Multimodal Interview Analysis

David Ortiz-Perez<sup>1</sup>, Jose Garcia-Rodriguez<sup>1</sup>, David Tomás<sup>2</sup>

<sup>1</sup> Department of Computer Technology, University of Alicante, Alicante, Spain

<sup>2</sup> Department of Software and Computing Systems, University of Alicante, Alicante, Spain

dortiz@dtic.ua.es, jgarcia@dtic.ua.es, dtomas@dlsi.ua.es

## Abstract

Cognitive decline is a natural process that occurs as individuals age. Early diagnosis of anomalous decline is crucial for initiating professional treatment that can enhance the quality of life of those affected. To address this issue, we propose a multimodal model capable of predicting Mild Cognitive Impairment and cognitive scores. The TAUADIAL dataset is used to conduct the evaluation, which comprises audio recordings of clinical interviews. The proposed model demonstrates the ability to transcribe and differentiate between languages used in the interviews. Subsequently, the model extracts audio and text features, combining them into a multimodal architecture to achieve robust and generalized results. Our approach involves in-depth research to implement various features obtained from the proposed modalities.

**Index Terms:** Multimodal, Cognitive Estimation, Audio Processing, Natural Language Processing

## 1. Introduction

Cognitive abilities tend to decline over time. While some abilities tend to remain unaffected, others, such as processing speed, reasoning, and memory, often show signs of deterioration [1]. Although cognitive decline is a common phenomenon, some individuals may experience a more pronounced decrease in these functions, particularly those affected by diseases such as dementia [2]. However, it is important to note that not every substantial decline in cognitive function indicates dementia. Some individuals may exhibit Mild Cognitive Impairment (MCI), which represents a stage between normal age-related cognitive deterioration and dementia [3, 4]. Therefore, there are various stages and considerations within this spectrum. The decline in cognitive abilities due to aging poses a significant challenge in our society, as there is a trend towards an aging population [5].

The primary motivation behind our work focuses on the early detection of these cognitive impairments. Early detection is crucial as it facilitates prompt intervention by professionals in the field. Such interventions can greatly benefit patients by enhancing their quality of life and mitigating the decline over time, consequently stimulating positive changes in their cognitive behavior [6].

Given the aforementioned facts, we propose a multimodal and multilingual model designed for the early detection of cognitive impairment. The model aims to accurately transcribe and distinguish between various languages in audio conversations to predict the cognitive state of elderly subjects. It determines whether they exhibit signs of MCI or possess a normal cognitive state for their age, using both text transcriptions and audio data. We experiment over the TAUADIAL dataset [7], which

includes audio conversations where elderly individuals describe a set of images in both English and Chinese language. We also consider the actual image descriptions in our analysis.

In summary, we make the following contributions:

- We propose a multimodal architecture capable of differentiating between Mild Cognitive Impairment (MCI) and a normal cognitive decline due to aging. The model takes actual conversations between a clinician and the elderly as input, extracting and post-processing transcription, and obtaining textual and acoustic information from these interactions.
- Our research involves a comprehensive exploration of various modalities approaches, and fusion strategies to effectively combine these modalities for optimal performance in this task.
- Furthermore, we emphasize the robustness and generalization of our model achieved through combinations of diverse modalities. Additionally, we highlight the ease of deployment, as the model solely relies on audio conversations.

The remainder of the paper is structured as follows. In Section 2, we discuss the related works and datasets relevant to this task. Section 3 presents the employed dataset and a brief analysis. Section 4 elaborates on the various approaches employed. Section 5 introduces the experimentation and presents the obtained results. Finally, in Section 6, we discuss the conclusions drawn from this work.

## 2. Related Works

This section presents information on datasets used for predicting dementia and other cognitive impairments. It also covers research carried out on these datasets, with a specific focus on multimodality.

### 2.1. Datasets

Due to the sensitivity of this type of data, acquiring large volumes poses a significant challenge, constituting a crucial aspect for training deep learning architectures.

When predicting dementia or other forms of cognitive decline, some datasets focus on medical data, such as the OASIS [8] and ADNI [9] datasets, which incorporate medical information and Magnetic Resonance Imaging (MRI) scans of the subject's brains. Datasets that provide information in video, audio, or text formats are relevant to our specific task and research. This emphasis relies on the practicality in real-life scenarios, where information can be captured easily using a standard camera for video recording.

Among these datasets are DementiaBank [10], and the ADReSS challenge [11]. These collections share similarities, featuring audio and text transcriptions of interviews between

clinicians and subjects.

## 2.2. Works

Recently, proposals for the automatic detection of cognitive impairments have made a breakthrough due to advancements in deep learning. Studies have explored patients' abilities to articulate and convey their thoughts, particularly in relation to cognitive impairments.

Significant contributions in this field include the works of Yuan et al. [12] and Balagopalan et al. [13]. They used pre-trained transformer-based models, such as BERT [14], over the ADRess challenge to analyze text features.

Audio analysis is another relevant factor for disease detection. Studies by Hajjar et al. [15] and Laguarda et al. [16] have leveraged audio biomarkers to predict cognitive diseases, such as dementia. Integration of these modalities in multi-modal approaches can be achieved through various strategies, as demonstrated in our previous works [17] and those proposed by Sarawgi et al [18]. A current hot topic is the fusion strategies of data for models that accommodate diverse data types, as exemplified by recent works such as CLIP [19] or LLaVA [20].

## 3. Dataset

The TAUKADIAL dataset [7] consists of recordings of elderly individuals prompted by clinicians to describe images. These recordings, which are provided in audio files, involve both Chinese and English speakers, with each participant asked to articulate descriptions for three distinct images. Notably, the image sets differ between English and Chinese participants. For English speakers, the assigned images are “*The Cookie Theft*” [21], “*Cat Rescue*” [22], and “*Coming and Going*” [23]. Meanwhile, Chinese speakers describe a father caring for his baby, scenes from a traditional night market in Taiwan, and the daily activities in Taiwan's parks.

The dataset is annotated to indicate whether the subject has MCI or normal cognitive decline for their age. Furthermore, for each participant, a score in the form of a Mini-Mental Status Examination (MMSE) is provided. The MMSE is a commonly used clinical tool for assessing cognitive impairment in patients [24]. It comprises a series of eleven questions.

The training set comprises 222 samples from individuals with MCI and 165 control samples, including 74 MCI subjects and 55 control participants. For each participant, there are three samples corresponding to the described pictures. The dataset comprises 201 Chinese and 186 English samples. The participants are aged between 61 and 87, with an average of 72 years. Among the participants, 79 are female and 50 are male.

## 4. Methodology

In this section, we introduce the different proposed architectures.

### 4.1. Data preprocessing

This step serves as the initial phase in our work. As this dataset comprises original data in audio files, and our interest also lies in extracting the semantic meaning of interviews, a preprocessing step has been conducted to transform audio into text.

To this aim, we utilized the Whisper model [25] to transcribe all interviews and distinguish between languages, specifically English and Chinese. The distinction between languages is crucial in this process, as the image descriptions provided for

each language vary. Furthermore, translating from one language to another may not be an optimal choice, given the importance of detecting subtle grammatical errors and nuances in expression.

Concerning image descriptions, our research involved analyzing transcribed samples and following protocols outlined by the organization, leading to the acquisition of authentic descriptions corresponding to the images. Simultaneously, audio files underwent preprocessing to extract features, such as audio biomarkers, encompassing elements like audio intensity, jitter, and shimmer, among others. To accomplish this, the OpenDBM<sup>1</sup> library was employed. The features obtained from the openSMILE toolkit [26] were also extracted and tested but with no promising results.

### 4.2. Text-based models

The preprocessing of the data involves providing among other features, text transcriptions. These transcriptions represent the actual dialogue between the subjects and the clinician. During this interaction, the subject describes the image presented by the clinician, providing us with the subject's description of the image. Additionally, a genuine description of the image is also acquired. With these two sets of data, we propose the following architectures, which are tested later:

**Text.** The initial proposed model is the text model, which solely utilizes the descriptions provided by the subjects. Two separate frozen BERT models [14] are used to extract features based on the language information and description obtained during the preprocessing step. The use of two distinct BERT models is necessary as there is one for each language. In line with the Transformer architecture [27], a transformer encoder has been used. This encoder is trained to extract additional features and consists of two encoder layers, each with the same size as the BERT base model (768 dimensions and 8 heads). The embeddings obtained from the transformer encoder serve as an input for a final Multilayer Perceptron (MLP) used for prediction. The transformer layers, as well as the MLPs, are different for each language and also vary for the classification and regression tasks.

**Similarity.** The similarity model, as the text model, entails obtaining BERT embeddings from the subject's image descriptions. However, in this model, BERT embeddings are also acquired for the real image descriptions. Both sets of embeddings, derived from frozen models to ensure consistency during training, undergo cosine similarity analysis [19] to assess the similarity between the subject's description and the real image. The resulting cosine similarity matrix is subsequently fed into an MLP for prediction.

**Combination.** This model closely resembles the text model, with the distinction lying in the embeddings utilized to feed the transformer encoders. These embeddings result from combining the subject's descriptions with the actual descriptions. Through this amalgamation, achieved via summation, we ensure that the model duly considers the image being described.

**Combined similarity.** This ensemble closely resembles the previous one. However, in this instance, the combination of embeddings occurs between the subject's description and the similarity to the genuine description. Through this combination, the model takes into account the differences between the provided description and the actual one, along with the remainder descriptions.

<sup>1</sup>[https://github.com/AiCure/open\\_dbm](https://github.com/AiCure/open_dbm)

With all these combinations, our aim is to ensure the capture of the most relevant information from the textual modality.

### 4.3. Audio-based model

In the context of audio modality, the features extracted during the preprocessing step have been employed. This involves a straightforward model, where the obtained features are fed into the MLP.

As the audio biomarkers do not significantly depend on the language used, no differentiation has been applied among languages for this model.

### 4.4. Multimodal model

In the case of the multimodal model, we combined the most effective text model with the audio model. We achieved this by using features extracted from the intermediate layers of their respective MLPs and combining them through concatenation. This allows us to integrate grammatical and semantic features from text-based models with those that may be lost in transcriptions.

Figure 1 illustrates the primary architecture of this pipeline, where the final MLP of the multimodal model combines features from both the audio and text models.

## 5. Experiments

This section introduces the experimental configurations, metrics, and results obtained from the previously proposed models.

### 5.1. Experimental setup

This subsection introduces the experimental configuration and setup used throughout this work.

**Cross-validation.** As there is no predefined validation split, and to ensure the selection of the best and most general model, a cross-validation strategy has been employed with a fold size of 5. This entails splitting the dataset into 5 different folds, enabling the model to be evaluated on various subsets to ensure robust performance, especially with non-trainable data.

**Implementation details.** We employed the AdamW optimizer with a learning rate of  $10^{-5}$  for models based on text features and the multimodal one, and 0.01 for the audio-based ones while incorporating a weight decay of 0.01 for all of them. For the regression task, the learning rate has been increased ten times.

The models underwent training with a batch size of 16 for 100 epochs, employing an early stopping strategy with a patience of 10 epochs. The cross-entropy loss function was used for the classification task, and the mean squared error was used for the regression task.

The experiments were conducted on an NVIDIA GeForce RTX 4090. All code developed for this project is publicly available in our GitHub repository<sup>2</sup>.

**Evaluation metrics.** To assess the performance of our proposed models, we used the suggested metrics within the context of the TAUADIAL challenge. Given that this challenge encompasses two distinct tasks, it employs different metrics for each. Regarding the classification task consisting of differentiating between MCI and control groups, the main metric has been the Unweighted Average Recall (UAR), expressed in Equation 1.

$$UAR = \frac{\sigma + \rho}{2} \quad (1)$$

Where  $\sigma$  represents specificity and  $\rho$  represents sensitivity.

F1 metric is also proposed for the classification task, represented in Equation 2, where  $\pi$  stands for the precision.

$$F_1 = \frac{2\pi\rho}{\pi + \rho}, \quad \text{where } \pi = \frac{TP}{TP + FP} \quad (2)$$

For the regression tasks, the primary metric is the Root Mean Squared Error (RMSE), as shown in Equation 3.

$$RMSE = \sqrt{\frac{\sum_{i=1}^N (\hat{y}_i - y_i)^2}{N}} \quad (3)$$

The coefficient of determination ( $R^2$ ) metric is also proposed for the regression task, represented in Equation 4.

$$R^2 = 1 - \frac{\sum_{i=1}^N (\hat{y}_i - y_i)^2}{\sum_{i=1}^N (\hat{y}_i - \bar{y})^2} \quad (4)$$

### 5.2. Ablation study

To evaluate the performance of different models, we conducted an ablation study using the cross-validation strategy. We implemented the proposed metrics and calculated them for each of the 5 folds with each model to compare their effectiveness.

The results for the classification task are presented in Table 1, where values represent the mean of the 5-fold metric along with the standard deviation. We also calculated the mean and standard deviation of all metrics.

Analyzing the obtained results from the classification task, it is evident that the multimodal model, which amalgamates features from all others, achieves the best outcome, particularly excelling in the crucial metric of UAR. Furthermore, this model also obtained the second-best result in the remaining metrics, underscoring its capability and robustness against other models.

It is noteworthy that audio features do not offer as much information as text-based approaches, surpassing only the similarity model. However, the addition of this similarity to other features has shown improvement, as evidenced by combined models. Notably, the similarity model alone does not perform remarkably well.

Additionally, an ablation study on the regression task was conducted. Examining these results, we can conclude that the multimodal model also outperforms the others. In this case, the most important metric is the RMSE, where the multimodal model performs best. Similarly, for the coefficient of determination, both the multimodal and similarity models provide the best results. For this task, the most relevant information regarding the text differs from the other task, with the similarity to the original description being the most important part. In the case of the audio modality, it slightly improves the performance compared to the other text-based models.

### 5.3. Challenge results

As part of the TAUADIAL challenge, a designated test set has been provided to evaluate the performance on entirely unseen data. Our submission to this challenge included four distinct models, and their respective performances are detailed in Table 2. The multimodal model was identified as the top-performing model based on the ablation study and serves as the foundation for our proposed models.

<sup>2</sup><https://github.com/davidorp/taukadial>

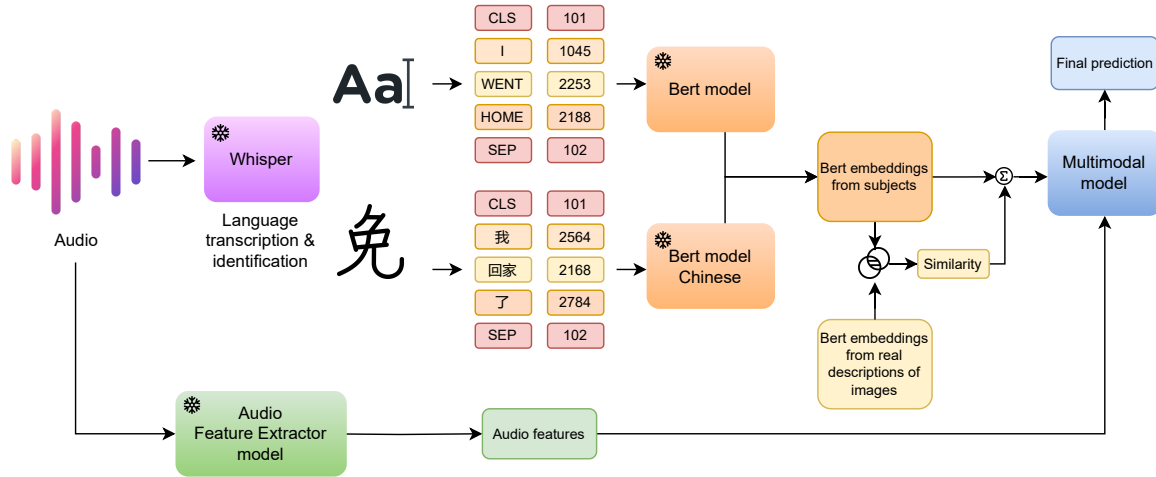


Figure 1: Overview of the multimodal architecture.

Table 1: Results of Classification Task using 5-Fold Cross-Validation.

Model	UAR	F1	$\sigma$	$\rho$	$\pi$	RMSE	$R^2$
Text	73.07 ± 1.5	76.92 ± 3.6	68.14 ± 8.4	78.0 ± 6.1	76.72 ± 7.1	3.14 ± 0.3	0.09 ± 0.0
Audio	70.75 ± 4.0	76.22 ± 5.9	61.41 ± 6.0	80.1 ± 9.2	73.32 ± 6.5	3.11 ± 0.2	0.1 ± 0.1
Similarity	68.57 ± 3.3	72.52 ± 5.1	63.67 ± 9.8	73.47 ± 11.8	73.44 ± 6.0	2.95 ± 0.4	<b>0.2 ± 0.1</b>
Combination	74.58 ± 2.2	78.47 ± 5.5	<b>70.44 ± 12.6</b>	78.72 ± 11.4	<b>79.44 ± 4.3</b>	3.14 ± 0.3	0.09 ± 0.0
Combined similarity	74.82 ± 2.5	<b>79.43 ± 1.4</b>	67.98 ± 7.3	<b>81.67 ± 4.3</b>	77.65 ± 3.6	3.14 ± 0.3	0.08 ± 0.0
Multimodal	<b>75.09 ± 1.8</b>	78.49 ± 3.4	70.0 ± 12.7	80.17 ± 10.6	78.95 ± 8.5	<b>2.93 ± 0.3</b>	<b>0.2 ± 0.1</b>

To improve the robustness of our approach, we used the multimodal model from the best fold of cross-validation as one model. Additionally, we fine-tuned this same model with data from the worst fold, resulting in the *MultimodalFull* model presented in the table. This process aims to ensure that our model is exposed to the entire training dataset while maintaining a validation set.

In addition, we have presented two extra models to calculate the mean predictions for each subject, considering that each participant has three samples. This strategy ensures consistent predictions for a given participant. The models marked with an asterisk in the table represent this approach.

Table 2: Results from the challenge’s test set

Model	UAR	RMSE
Multimodal (ours)	56.18	2.59
MultimodalFull (ours)	53.68	2.64
Multimodal* (ours)	55.39	2.58
MultimodalFull* (ours)	52.01	<b>2.50</b>
w2v+eGeMAPs [7]	<b>59.18</b>	13.28
ling. [7]	53.59	2.89
w2v+ling [7]	50	4.39
eGeMAPs [7]	44.95	17
w2v [7]	46.05	4.48

Examining the results obtained in the challenge. The model’s performance is less effective compared to the validation folds for the classification task. Despite this, the model outper-

forms the results obtained in the validation set for the regression task. All models demonstrate similar performance, with the *multimodal* model performing exceptionally well in classification and the *Mean MultimodalFull* model standing out in the regression task.

Compared to State-Of-The-Art models, our approach achieves superior results in the regression task, surpassing others with a 2.50 RMSE score. However, in the classification task, we attain the second-best performance with a 56.18 UAR. It is important to note that our model demonstrates greater robustness and stability. It is worth mentioning that the model that outperforms ours in classification tasks exhibits significantly poorer results in regression tasks compared to our model.

## 6. Conclusions

In this work, we have presented a multimodal model capable of predicting MCI or normal cognitive decline due to aging, along with a cognitive score based on MMSE. The early detection of these declines is crucial for the well-being of patients who experience them. We have employed the TAUADIAL dataset, which consists of audio recordings from interviews between a clinician and a subject.

The research conducted throughout this work focuses on the combination of features derived from both text transcriptions and audio modalities. Consequently, we have developed a strong and easily deployable model by combining modalities.

In future works, we aspire to persist in investigating this line of research and delve into additional cognitive impairment diseases, besides their impact on people’s behavior.

## 7. Acknowledgements

We would like to thank “A way of making Europe” European Regional Development Fund (ERDF) and MCIN/AEI/10.13039/501100011033 for supporting this work under the “CHAN-TWIN” project (grant TED2021-130890B-C21. HORIZON-MSCA-2021-SE-0 action number: 101086387, REMARKABLE, Rural Environmental Monitoring via ultra wide-ARea networKs And distriButed federated Learning. CIAICO/2022/132 Consolidated group project “AI4Health” funded by Valencian government and International Center for Aging Research ICAR funded project “IASISTEM”. This work has also been supported by a regional grant for PhD studies, CIACIF/2022/175.

## 8. References

- [1] I. J. Deary, J. Corley, A. J. Gow, S. E. Harris, L. M. Houlihan, R. E. Marioni, L. Penke, S. B. Rafnsson, and J. M. Starr, “Age-associated cognitive decline,” *British Medical Bulletin*, vol. 92, no. 1, pp. 135–152, 09 2009. [Online]. Available: <https://doi.org/10.1093/bmb/ldp033>
- [2] S. Duong, T. Patel, and F. Chang, “Dementia: What pharmacists need to know,” *Canadian Pharmacists Journal / Revue des Pharmaciens du Canada*, vol. 150, no. 2, pp. 118–129, 2017. [Online]. Available: <https://doi.org/10.1177/1715163517690745>
- [3] F. Portet, P. J. Ousset, P. J. Visser, G. B. Frisoni, F. Nobili, P. Scheltens, B. Vellas, J. Touchon, and the MCI Working Group of the European Consortium on Alzheimer’s Disease (EADC), “Mild cognitive impairment (mci) in medical practice: a critical review of the concept and new diagnostic procedure. report of the mci working group of the european consortium on alzheimer’s disease,” *Journal of Neurology, Neurosurgery & Psychiatry*, vol. 77, no. 6, pp. 714–718, 2006. [Online]. Available: <https://jnnp.bmj.com/content/77/6/714>
- [4] P. Celsis, “Age-related cognitive decline, mild cognitive impairment or preclinical alzheimer’s disease?” *Annals of Medicine*, vol. 32, no. 1, pp. 6–14, 2000, pMID: 10711572. [Online]. Available: <https://doi.org/10.3109/07853890008995904>
- [5] World Health Organization, “Mental health action plan 2013-2020,” *WHO Library Cataloguing-in-Publication DataLibrary Cataloguing-in-Publication Data*, pp. 1–44, 2023.
- [6] M. Di Luca, D. Nutt, W. Oertel, P. Boyer, J. Jaarsma, F. Destrebecq, G. Esposito, and V. Quoidbach, “Towards earlier diagnosis and treatment of disorders of the brain,” *Bulletin of the World Health Organization*, vol. 96, no. 5, p. 298, 2018.
- [7] S. Luz, S. d. l. F. Garcia, F. Haider, D. Fromm, B. MacWhinney, A. Lanzi, Y.-N. Chang, C.-J. Chou, and Y.-C. Liu, “Connected speech-based cognitive assessment in chinese and english,” 2024.
- [8] D. S. Marcus, T. H. Wang, J. Parker, J. G. Csernansky, J. C. Morris, and R. L. Buckner, “Open access series of imaging studies (oasis): cross-sectional mri data in young, middle aged, nondemented, and demented older adults,” *Journal of cognitive neuroscience*, vol. 19, no. 9, pp. 1498–1507, 2007.
- [9] R. C. Petersen, P. S. Aisen, L. A. Beckett, M. C. Donohue, A. C. Gamst, D. J. Harvey, C. R. Jack, W. J. Jagust, L. M. Shaw, A. W. Toga *et al.*, “Alzheimer’s disease neuroimaging initiative (adni): clinical characterization,” *Neurology*, vol. 74, no. 3, pp. 201–209, 2010.
- [10] A. M. Lanzi, A. K. Saylor, D. Fromm, H. Liu, B. MacWhinney, and M. L. Cohen, “Dementiabank: Theoretical rationale, protocol, and illustrative analyses,” *American Journal of Speech-Language Pathology*, vol. 32, no. 2, pp. 426–438, 2023.
- [11] S. Luz, F. Haider, S. de la Fuente, D. Fromm, and B. MacWhinney, “Alzheimer’s dementia recognition through spontaneous speech: The adress challenge,” 2020.
- [12] J. Yuan, Y. Bian, X. Cai, J. Huang, Z. Ye, and K. Church, “Disfluencies and fine-tuning pre-trained language models for detection of alzheimer’s disease,” in *Interspeech*, vol. 2020, 2020, pp. 2162–6.
- [13] A. Balagopalan, B. Eyre, F. Rudzicz, and J. Novikova, “To bert or not to bert: comparing speech and language-based approaches for alzheimer’s disease detection,” *arXiv preprint arXiv:2008.01551*, 2020.
- [14] J. Devlin, M. Chang, K. Lee, and K. Toutanova, “BERT: pre-training of deep bidirectional transformers for language understanding,” *CoRR*, vol. abs/1810.04805, 2018. [Online]. Available: <http://arxiv.org/abs/1810.04805>
- [15] I. Hajjar, M. Okafor, J. D. Choi, E. Moore, A. Abrol, V. D. Calhoun, and F. C. Goldstein, “Development of digital voice biomarkers and associations with cognition, cerebrospinal biomarkers, and neural representation in early alzheimer’s disease,” *Alzheimer’s & Dementia: Diagnosis, Assessment & Disease Monitoring*, vol. 15, no. 1, p. e12393, 2023.
- [16] J. Laguarda and B. Subirana, “Longitudinal speech biomarkers for automated alzheimer’s detection,” *Frontiers in Computer Science*, vol. 3, 2021. [Online]. Available: <https://www.frontiersin.org/articles/10.3389/fcomp.2021.624694>
- [17] D. Ortiz-Perez, P. Ruiz-Ponce, D. Tomás, J. Garcia-Rodriguez, M. F. Vizcaya-Moreno, and M. Leo, “Deep learning-based dementia prediction using multimodal data,” *Neurocomputing*, p. 126413, 2023.
- [18] U. Sarawgi, W. Zulfikar, N. Soliman, and P. Maes, “Multimodal inductive transfer learning for detection of alzheimer’s dementia and its severity,” 2020.
- [19] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, “Learning transferable visual models from natural language supervision,” *CoRR*, vol. abs/2103.00020, 2021. [Online]. Available: <https://arxiv.org/abs/2103.00020>
- [20] H. Liu, C. Li, Q. Wu, and Y. J. Lee, “Visual instruction tuning,” 2023.
- [21] G. H. and K. E., “Boston diagnostic aphasia examination booklet,” *Lea & Febiger*, 1983.
- [22] L. E. Nicholas and R. H. Brookshire, “A system for quantifying the informativeness and efficiency of the connected speech of adults with aphasia,” *Journal of Speech, Language, and Hearing Research*, vol. 36, no. 2, pp. 338–350, 1993. [Online]. Available: <https://pubs.asha.org/doi/abs/10.1044/jshr.3602.338>
- [23] N. Rockwell, “Going and coming [oil on canvas],” 1947.
- [24] I. Arevalo-Rodriguez, N. Smailagic, M. R. i Figuls, A. Ciapponi, E. Sanchez-Perez, A. Giannakou, O. L. Pedraza, X. B. Cosp, and S. Cullum, “Mini-mental state examination (mmse) for the detection of alzheimer’s disease and other dementias in people with mild cognitive impairment (mci),” *Cochrane Database of Systematic Reviews*, no. 3, 2015.
- [25] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” 2022.
- [26] 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*, 01 2010, pp. 1459–1462.
- [27] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” *CoRR*, vol. abs/1706.03762, 2017. [Online]. Available: <http://arxiv.org/abs/1706.03762>