



ParaCLAP – Towards a general language-audio model for computational paralinguistic tasks

Xin Jing*^{1,2}, Andreas Triantafyllopoulos*^{1,2}, Björn Schuller^{1,2,3}

¹Chair of Embedded Intelligence for Health Care & Wellbeing, University of Augsburg, Germany

²CHI – Chair of Health Informatics, MRI, Technical University of Munich, Germany

³GLAM – Group on Language, Audio, & Music, Imperial College London, UK

xin.jing@tum.de

Abstract

Contrastive language-audio pretraining (CLAP) has recently emerged as a method for making audio analysis more generalisable. Specifically, CLAP-style models are able to ‘answer’ a diverse set of language queries, extending the capabilities of audio models beyond a closed set of labels. However, CLAP relies on a large set of (audio, query) pairs for pretraining. While such sets are available for general audio tasks, like captioning or sound event detection, there are no datasets with matched audio and text queries for computational paralinguistic (CP) tasks. As a result, the community relies on generic CLAP models trained for general audio with limited success. In the present study, we explore training considerations for ParaCLAP, a CLAP-style model suited to CP, including a novel process for creating audio-language queries. We demonstrate its effectiveness on a set of computational paralinguistic tasks, where it is shown to surpass the performance of open-source state-of-the-art models. Our code and resources are publicly available at: <https://github.com/KeiKinn/ParaCLAP>

Index Terms: computational paralinguistic, speech emotion recognition, contrastive learning, zero-shot learning

1. Introduction

Computational paralinguistics (CP) is a subfield of affective computing which corresponds to the analysis of the paralinguistic information in speech signals for the prediction of speaker states and traits [1]. As such, it encompasses a large gamut of phenomena that can be predicted: from the now ‘classic’ emotions, to (among many others) personality, likability, sincerity, deception, and even health-related tasks. The typical machine learning (ML) workflow requires the collection of representative data for each of those tasks and the subsequent training of models – oftentimes relying on transfer learning or unsupervised pretraining to improve performance while reducing the dependency on data quantity [2]. Nonetheless, the challenge remains that at least *some* amount of data has to be acquired for each task, which presents a bottleneck for practitioners who want to benefit from advances in CP but are interested in phenomena where no existing data or public models exist.

A similar challenge is faced by the computational audiotory scene analysis (CASA) community, which also aims to analyse a wide spectrum of scenes and events [3, 4, 5]. However, that field has seen the recent emerge of contrastive language-audio pretraining (CLAP) as a general-purpose method which can overcome this lack of data and result in a generic model that can be employed for a wide variety of tasks. The initial inspiration for CLAP came from computer vision [6]. CLAP is

trained to compute the similarity of linguistic and audio embeddings (each generated by the respective linguistic/auditory encoders) [7]. During inference, CLAP relies on a set of linguistic queries (provided by the user) which are scored with respect to their compatibility with the input audio. Importantly, this allows the application of the model on *completely new* tasks never seen during training. On the downside, its success is heavily based on the large amount of pretraining data formed of (audio, query) pairs. In the CASA domain, the community has amassed a large set of data that can be used for this purpose, e. g., by using audio captioning datasets [3, 5] or by co-opting the tags and descriptions associated with an audio sample in large data platforms like Freesound [8]. As no such public data exist for CP tasks, prior research has primarily used CP tasks in downstream evaluations of zero-shot performance. For example, Elizalde *et al.* [7] evaluated their CLAP model on two speech emotion recognition (SER) datasets and obtained results slightly above chance-level performance. This shows how general CLAP models pretrained on CASA tasks may struggle in the CP domain.

On a related, more promising note, a similar setup has been recently used to improve the expressivity of affective speech synthesis. For example, PromptTTS-2[9] constructed a set of linguistic “prompts” used to control their synthesis system. This set is created using templates derived from acoustic parameters and showed improved effectiveness in synthesising the target attributes. While such approaches illustrate the potential of generating linguistic prompts suitable to describe paralinguistic attributes, to the best of our knowledge they have not been used to improve zero-shot recognition for CP tasks. Moreover, previous work on zero-shot learning has shown that this is possible for SER [10], and in particular using semantic embeddings derived directly from the labels [11]. Similarly, these findings have not been extended to a more broader range of CP tasks.

Our work aims to bridge this gap by presenting a novel templating method based on interpretable, expert features known to capture paralinguistic attributes. Starting from those features and the set of labels available for a large SER dataset (MSP-Podcast), we construct a diverse set of (audio, query) samples. These pairs are given as input to an audio (wav2vec2.0) and text (BERT) encoder, which procure the corresponding embeddings. During training, these embeddings are *aligned* using a contrastive loss. During inference, the similarity of the input audio and the target linguistic queries is computed, and the query with the highest similarity is assigned as the predicted class. We demonstrate the effectiveness of our method on a set of SER datasets to first test their out-of-domain performance, as well as two broader CP tasks to gauge their generalisability.

The remainder of this paper is organised as follows: Our proposed templating process and model are introduced in Section 3. Section 4 presents our results and analysis. Finally, con-

* These authors contributed equally

clusion and future work are discussed in Section 5.

2. Methodology

2.1. ParaCLAP

Our overall workflow follows the standard procedure from CLAP models [7]. As shown in Fig. 1, ParaCLAP establishes a link between audio and text representations within a shared multimodal space by applying two encoders and aligning their output embeddings using contrastive learning.

In detail, let $\{X_i^a, X_i^t\}$ be the input audio-text pairs, in which $i \in [0, N]$ and N is the batch size. The audio-text pairs will be processed by an audio and text encoder, as shown in Fig. 1, to extract an audio and a text embedding. Additionally, every encoder incorporates a projection layer which serves the purpose of projecting both the audio and text embeddings into a shared feature space of dimensionality d . Every projection layer is a module consisting of two linear transformations followed by a GELU activation and layer normalisation. It initially maps the input to larger dimensions and then downsamples to its original dimensions, preserving a consistent structure. In our study, the output size of the projection layer is 768. The final audio and text embeddings are computed as $A_i = \text{proj}_A(f(X_i^a))$, $A_i \in \mathbb{R}^{N \times d}$, and $T_i = \text{proj}_T(g(X_i^t))$, $T_i \in \mathbb{R}^{N \times d}$, where the $f(\cdot)$ and $g(\cdot)$ are the audio encoder and text encoder, respectively. In this shared feature space, the text embedding T_i and audio embedding A_i can be compared by measuring their similarity. For this purpose, we use the scaled cosine similarity:

$$S = \tau * (A_i \cdot T_i^{\text{tr}}), \quad (1)$$

where τ is a temperature parameter to scale the range of logits. The similarity matrix S contains N positive pairs in the diagonal and $N^2 - N$ negative pairs in the off-diagonal.

To compute the final loss for training, the following symmetric cross-entropy loss is applied:

$$\text{Loss} = 0.5 * (h_t(S) + h_a(S)), \quad (2)$$

where the $h_k = \frac{1}{N} \sum_{n=0}^N \log(\text{diag}(\text{softmax}(S)))$ along the audio and text axis. The CLAP model is thus optimised by maximising the normalised similarity of positive pairs and minimising the similarity of negative pairs, which means the asymptotic limit of diagonal elements approaching 1 ($\lim_{i \rightarrow \infty} \text{diag}(S)_{ii} = 1$), and non-diagonal elements converging towards zero ($\lim_{i \neq j \rightarrow \infty} \text{off}(S)_{ij} = 0$).

In this study, we adopt the wav2vec 2.0 large model [12] as the audio encoder, initialised with a pre-trained state which has been finetuned for dimensional SER¹ [2]. The pre-trained model was pruned from 24 to 12 transformer layers before finetuning on the MSP-Podcast (v1.7) dataset [13]. The pre-trained wav2vec 2.0 large model consists of 160 million parameters. The audio embedding is obtained by extracting the pooled states from the last transformer layer which has a size of 1024. For the text encoder, we apply the BERT [14] base model (uncased)² implemented by HuggingFace [15]. The model has 110 million parameters and the [CLS] token extracted from the final layer is utilised as the text embedding.

¹<https://huggingface.co/auddeering/wav2vec2-large-robust-12-ft-emotion-msp-dim>

²<https://huggingface.co/bert-base-uncased>

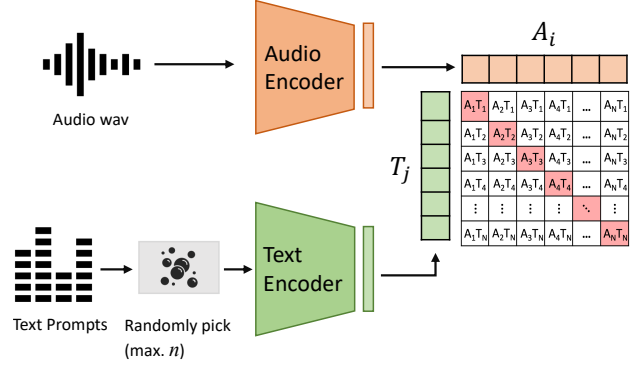


Figure 1: Diagram of our proposed ParaCLAP model. ParaCLAP jointly trains two parallel encoders to establish a link between the input audio-text pairs in a shared multimodal space. In the training phase, a maximum of n text prompts are randomly selected from the pool of generated text queries, combining both label and pseudo-caption information, while during inference, the query comprises only the actual labels.

2.2. Query generation

We generate text queries matching each sample by relying on two main sources of information: a) The labels already included in the dataset. b) Expert acoustic and prosodic features extracted for each sample. Each source of query is extended to a complete sentence as described below. A comprehensive table is provided as supplementary material³.

Dataset labels: Our training dataset (Section 3.1) is labelled for emotion, both categorical and dimensional (arousal, valence, dominance), as well as gender. Categorical emotion is expanded by constructing sentences like “this is a [EMOTION] instance” or c) “speaker is [EMOTION]”, with [EMOTION] coming in the form of adjective (e. g., “angry”, “happy”, etc.). The gender label (“male” is expanded in two versions: a) a [GENDER] is speaking; b) the speaker is [GENDER]. Finally, the dimensional emotional attributes are first binned according to their distribution (bottom 30%, middle 40%, top 30%) and templates are constructed accordingly (e. g., arousal is low/mid/high).

Pseudo-captions using expert features: We extract the extended Geneva Minimalistic Acoustic Parameter Set (eGeMAPS) [16] using openSMILE [17]. It includes 88 parameters and is known for its effectiveness in several CP tasks. In this study, we focus on the following features which are easier to interpret: a) mean (μ) and standard deviation (σ) of *pitch*; b) sound *intensity* (over whole utterance); c) *jitter* (mean over utterance); d) *shimmer* (mean over utterance). In addition, we compute the total *duration* of each utterance. This gives us 5 numerical features, which we proceed to bin according to their distribution as for the dimensional variables. Following that, we generate queries using those bins (e. g., “pitch is low/normal/high”).

Combinations: Finally, during training we combine several queries at once using an “and” conjunction (e. g., “speaker is happy and pitch is high”). This further expands the diversity of data that the model sees during training, and is aimed to improve training performance and prevent overfitting.

³The process can be more easily understood by inspecting the code submitted as supplementary material.

Table 1: *The Unweighted Average Recall (UAR) results on 7 datasets, comprising both English and German language-based datasets. ‘emo’ indicates the use of emotion queries during training. The notation $rand = n$ means a maximum of n queries are randomly selected and concatenated. If ‘emo’ is mentioned, it implies emotion queries’ involvement in the data. **Bold** marks the best performance while underline represents the second best.*

Dataset	Class	Lang.	Unweighted Average Recall (UAR)					
			CLAP (CNN14 + Bert)	Pengi (HTSAT + GPT2)	ParaCLAP (no emo, $rand = 5$)	ParaCLAP (only emo)	ParaCLAP ($rand = 1$)	ParaCLAP ($rand = 5$)
IEMOCAP	4cl	en	.353	.345	.309	.567	.307	<u>.560</u>
RAVDESS	8cl	en	.199	.148	.170	.302	.116	<u>.234</u>
CREMA-D	6cl	en	.230	.245	.201	.332	.202	<u>.291</u>
TESS	7cl	en	.232	.177	.212	.484	.219	<u>.389</u>
FAU Aibo	2cl	de	.500	.470	<u>.538</u>	.535	.468	.604
FAU Aibo	5cl	de	.211	.185	<u>.225</u>	.216	.216	.232
ALC	2cl	de	<u>.511</u>	.473	.490	.512	.501	.503
SLD	2cl	de	.472	.485	.472	.554	.443	<u>.507</u>

3. Experiments

3.1. Training Dataset

In this work, MSP-podcast Release 1.9 [13] is utilised to train our ParaCLAP models. The MSP-podcast data comprises natural English speeches extracted from podcast recordings. The dataset encompasses 55 283 utterances spoken by over 1 200 speakers, totalling more than 110 hours of speech. The corpus is annotated using crowd-sourcing with 9 emotion types (**neutral, fear, sadness, disgust, happiness, other, anger, contempt, surprise**), but nearly 17.5% of the data does not have majority agreed labels (labelled as ‘no_agreement’). Moreover, the data is heavily imbalanced towards the neutral class (35%), followed by happiness (21%) and anger (6%). During the training phase, the standard splits for training and testing are employed. However, data labelled as ‘no_agreement’ is excluded, resulting in a final amount of 45 619 utterances for training. To construct the audio-text pairs for training, we dynamically pick texts from the text queries outlined in Section 2.2, and additional details for text-audio pairs will be presented in Section 3.3.

3.2. Test Datasets

IEMOCAP: The Interactive Emotional Dyadic Motion Capture [18] is a prominent standard multi-modal database for emotion studies. We follow the most frequently used category selection (angry, happy+excited, sad, and neutral) to build the test dataset. Thus, the dataset contains 5 531 utterances (**1 103 angry, 1 636 happy, 1 708 neutral, and 1 084 sad**).

RAVDESS: The Ryerson Audio-Visual Database of Emotional Speech and Song [19] is a multimodal emotion database containing speech and songs. The speech data consists of 1 440 utterances with 8 expressions (**neutral, calm, happy, sad, angry, fearful, surprise, and disgust**).

CREMA-D: Crowd Sourced Emotional Multimodal Actors Dataset [20] is a crowd-sourced audio-visual data set tailored for emotion studies. It comprises 7 442 clips from 91 actors. The dataset consists of facial and vocal emotional expressions in sentences, conveying a diverse range of basic emotional states (**happy, sad, anger, fear, disgust, and neutral**).

TESS: Toronto Emotional Speech Set [21] is an audio emotion data set recorded by two female speakers, aged 26 and 64 years. It contains 2 800 clips featuring seven emotions: **anger, disgust, fear, happiness, pleasant surprise, sadness, and neutral**.

FAU-Aibo: The FAU-Aibo Emotion Corpus [22] is a German speech emotion database containing children speech in the ages 6 to 10 years. It is used in the INTERSPEECH 2009 Emotion Challenge [23], and includes a training set of 9 959 speech chunks and a test set of 8 257 chunks. The original 11 labels are mapped to a) a five-category classification problem, with labels being merged into **angry, emphatic, neutral, positive, and rest**, and b) a two-category classification problem, with ‘**non-negative**’ and ‘**negative**’ as the emotion labels.

ALC: The Alcohol Language Corpus [24] contains German speech collected under a systematic intoxication test. The INTERSPEECH 2011 Speaker State Challenge [25] selected part of the ALC to obtain a gender and age balanced dataset. 1 620 data with the label ‘**not intoxicated with alcohol**’ and 1 380 labelled as ‘**intoxicated with alcohol**’ are contained in the test set.

SLD: The Speaker Likability Database is a subset of the German telephone speech dataset aGender [26]. It was used in the INTERSPEECH 2012 Speaker Trait Challenge [27] and contains 800 audio chunks each for the labels ‘**likable**’ and ‘**non-likable**’.

3.3. Experimental setup

We use the raw waveform after resampling to 16 kHz for training and testing. During training, all audio sequences are randomly clipped or padded to achieve a consistent 5-second duration and a maximum of n text prompts are randomly selected from the pool of the text queries and concatenated into a single unit. We additionally control whether the emotion label is part of the training query, resulting in a total of four alternatives:

- ParaCLAP (no emo, $rand = 5$): emotion query is not part of training, a maximum of 5 random queries concatenated
- ParaCLAP (only emo): only emotion query is used as caption
- ParaCLAP ($rand = 1$): emotion query is part of training, a maximum of 1 random queries
- ParaCLAP ($rand = 5$): emotion query is part of training, a maximum of 5 random queries

All models are trained with an Adam optimiser and a batch size of 64. The number of training epochs is set to 50. Both the audio and text branches remain unfrozen, utilising a learning rate of $1e - 5$. Meanwhile, the projection layers and other parameters employ a learning rate set of $1e - 3$. The models




Figure 2: The confusion matrices illustrate the results on the IEMOCAP dataset. Models share identical training settings, differing in the inclusion of emotion queries within the audio-text pairs (left: without, right: with)




Figure 3: The confusion matrices illustrate the results on the TESS dataset. Models share identical training settings, differing in the inclusion of emotion queries within the audio-text pairs (left: without, right: with)

are implemented with PYTORCH-v1.13.1 on a single NVIDIA A40, 48 GB GPU.

The ParaCLAP model from the epoch that yielded the best performance on the MSP-Podcast test set is selected as the best-performing model and used for the downstream evaluations. During inference, we use the annotated labels provided by the datasets as the text query. As baselines, we use the original CLAP model⁴, which employs a CNN14 audio encoder and a BERT text encoder, and has been trained on larger data featuring a wider gamut of sound tasks, and Pengi [28], a large audio-language model trained for a variety of tasks. Pengi generates free-form text which can be matched to each query by computing the similarity of query and output embeddings, as in the original work [28].

4. Results & Discussion

The results are shown in Table 1. The datasets involved in the experiments are characterised by the number of classes (Class), the language they contain (Lang.), and specific training conditions, including the use of emotion queries (‘emo’) and concatenation of random queries.

Our first observation is that our models generally outperform both the original CLAP baseline and Pengi, showcasing how our model improves upon the state-of-the-art zero-shot method used in most recent works. In terms of ParaCLAP alternatives, it is evident that a model only trained with emotion queries (i. e., “emotion is [EMOTION]”) is showing the best overall performance. This is especially true for datasets which show a big overlap in their emotion classes with MSP-Podcast, i. e., standard categorical SER datasets. For example, the sets

⁴<https://huggingface.co/microsoft/msclap>

of labels in MSP-Podcast and IEMOCAP are identical. In this case, our method is tantamount to doing standard classification, and we expect the optimal performance when optimising specifically for this task.

It is only when turning to alternative SER formulations, such as the one in FAU-AIBO, that the addition of template queries shows its strength. This helps improve performance to a UAR of .604 and .232, over .535 and .216 for the 2- and 5-class tasks, respectively, compared to ParaCLAP (only emo). On a further positive note, this shows how our model can generalise beyond English data, even though both the upstream training of the speech encoder and CLAP fine-tuning only contained English data. However, these gains do not translate further to tasks beyond emotions (ALC/SLD), where ParaCLAP (only emo) still shows the best performance.

Further insight into the impact of the use of emotion labels to generate queries can be found in Fig. 2 and Fig. 3, which show the confusion matrix of ParaCLAP (no emo, $rand = 5$) and ParaCLAP ($rand = 5$) on IEMOCAP and TESS. In the absence of emotion queries during training, the model tends to categorise all utterances as “sadness”⁵. This is partially rectified through the inclusion of emotion queries, but still remains a challenge – especially for TESS. We hypothesise that this is caused by a discrepancy in the type of data included in MSP-Podcast compared to the other emotional datasets considered here; MSP-Podcast is a naturalistic corpus collected ‘in-the-wild’ whereas all four emotional datasets are acted – and therefore contain more ‘archetypal’ depictions of emotions (i. e., high arousal). The depiction of sadness in MSP-Podcast may be less pronounced and more similar to the (overrepresented) neutral class, thus causing misclassifications towards that when the model is applied on other data.

Overall, our results show that increasing the diversity of the training queries is important to improve performance. The inclusion of the labelled emotion in these queries additionally plays a key role – especially when evaluating on related tasks. This illustrates how the quest to identify good strategies for creating queries that lend themselves to generalisation across different CP tasks is still open.

5. Conclusions

In this study, we conducted a preliminary investigation of different factors that can influence the performance of a CLAP-style model specifically designed for computational paralinguistics tasks. In this process, we have created the ParaCLAP model that can substantially outperform the state-of-the-art generic CLAP model widely used for zero-shot evaluation in recent works. Furthermore, by exploring a templating method to generate text queries derived from handcrafted features and expert knowledge, we have shown how diversity in training can be beneficial for performance.

In summary, this first attempt to create ParaCLAP shows great promise, albeit demonstrating that there is still a lot of ground to be covered. More diversity in the training queries is required, perhaps augmented by the use of large language models (LLMs), as done for audio captioning [29], or even a more seamless integration with LLMs [28, 30, 31]. Moreover, the role of the training data still remains to be investigated, with data quantity and diversity expected to be another key factor.

⁵Similar behaviour is observed for the other datasets, whose confusion matrices are included as supplementary material.

6. Acknowledgements

This work was funded by the China Scholarship Council (CSC), Grant #202006290013, and by the DFG, Reinhart Koselleck-Project AUDIIONOMOUS (Grant No. 442218748).

7. References

- [1] B. Schuller and A. Batliner, *Computational Paralinguistics: Emotion, Affect and Personality in Speech and Language Processing*. Wiley, Nov. 2014, ISBN: 978-1119971368.
- [2] J. Wagner, A. Triantafyllopoulos, H. Wierstorf, M. Schmitt, F. Burkhardt, F. Eyben, and B. W. Schuller, “Dawn of the transformer era in speech emotion recognition: Closing the valence gap,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 02, pp. 10 745–10 759, 2023.
- [3] K. Drossos, S. Lipping, and T. Virtanen, “Clotho: An audio captioning dataset,” in *Proc. ICASSP, IEEE*, 2020, pp. 736–740.
- [4] H. Xie and T. Virtanen, “Zero-shot audio classification via semantic embeddings,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1233–1242, 2021.
- [5] X. Mei, X. Liu, M. D. Plumbley, and W. Wang, “Automated audio captioning: An overview of recent progress and new challenges,” *EURASIP journal on audio, speech, and music processing*, vol. 2022, no. 1, pp. 1–18, 2022.
- [6] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, *et al.*, “Learning transferable visual models from natural language supervision,” in *International conference on machine learning*, PMLR, 2021, pp. 8748–8763.
- [7] B. Elizalde, S. Deshmukh, M. Al Ismail, and H. Wang, “Clap learning audio concepts from natural language supervision,” in *Proc. the International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, Rhodes Island, Greece, 2023, pp. 1–5.
- [8] E. Fonseca, J. Pons Puig, X. Favory, F. Font Corbera, D. Bogdanov, A. Ferraro, S. Oramas, A. Porter, and X. Serra, “Freesound datasets: A platform for the creation of open audio datasets,” in *Proc. ISMIR*, 2017, pp. 486–493.
- [9] Y. Leng, Z. Guo, K. Shen, X. Tan, Z. Ju, Y. Liu, Y. Liu, D. Yang, L. Zhang, K. Song, *et al.*, “Prompts 2: Describing and generating voices with text prompt,” *arXiv preprint arXiv:2309.02285*, 2023.
- [10] X. Xu, J. Deng, Z. Zhang, X. Fan, L. Zhao, L. Devillers, and B. W. Schuller, “Rethinking auditory affective descriptors through zero-shot emotion recognition in speech,” *IEEE Transactions on Computational Social Systems*, vol. 9, no. 5, pp. 1530–1541, 2021.
- [11] X. Xu, J. Deng, N. Cummins, Z. Zhang, L. Zhao, and B. W. Schuller, “Exploring zero-shot emotion recognition in speech using semantic-embedding prototypes,” *IEEE Transactions on Multimedia*, vol. 24, pp. 2752–2765, 2021.
- [12] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “Wav2vec 2.0: A framework for self-supervised learning of speech representations,” *Advances in neural information processing systems*, vol. 33, pp. 12 449–12 460, 2020.
- [13] R. Lotfian and C. Busso, “Building naturalistic emotionally balanced speech corpus by retrieving emotional speech from existing podcast recordings,” *IEEE Transactions on Affective Computing*, vol. 10, no. 4, pp. 471–483, 2017.
- [14] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [15] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, *et al.*, “Huggingface’s transformers: State-of-the-art natural language processing,” *arXiv preprint arXiv:1910.03771*, 2019.
- [16] 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.
- [17] F. Eyben, M. Wöllmer, and B. Schuller, “openSMILE: The Munich versatile and fast open-source audio feature extractor,” in *Proc. the International Conference on Multimedia*, 2010, pp. 1459–1462.
- [18] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan, “Iemocap: Interactive emotional dyadic motion capture database,” *Language resources and evaluation*, vol. 42, pp. 335–359, 2008.
- [19] S. R. Livingstone and F. A. Russo, “The ryerson audio-visual database of emotional speech and song (ravdess): A dynamic, multimodal set of facial and vocal expressions in north american english,” *PLoS one*, vol. 13, no. 5, e0196391, 2018.
- [20] H. Cao, D. G. Cooper, M. K. Keutmann, R. C. Gur, A. Nenkova, and R. Verma, “Crema-d: Crowd-sourced emotional multimodal actors dataset,” *IEEE transactions on affective computing*, vol. 5, no. 4, pp. 377–390, 2014.
- [21] M. K. Pichora-Fuller and K. Dupuis, “Toronto emotional speech set (TESS),” Borealis, 2020. DOI: 10.5683/SP2/E8H2MF. [Online]. Available: <https://doi.org/10.5683/SP2/E8H2MF>.
- [22] A. Batliner, S. Steidl, and E. Nöth, “Releasing a thoroughly annotated and processed spontaneous emotional database: The fau aibo emotion corpus,” 2008.
- [23] B. Schuller, S. Steidl, and A. Batliner, “The interspeech 2009 emotion challenge,” in *Proc. International Speech Communication Association (INTERSPEECH)*, Brighton, UK, 2009.
- [24] F. Schiel, C. Heinrich, and S. Barfüsser, “Alcohol language corpus: The first public corpus of alcoholized german speech,” *Language resources and evaluation*, vol. 46, pp. 503–521, 2012.
- [25] B. Schuller, A. Batliner, S. Steidl, F. Schiel, and J. Krajewski, “The interspeech 2011 speaker state challenge,” in *Proc. International Speech Communication Association (INTERSPEECH)*, Florence, Italy, 2011.
- [26] F. Burkhardt, M. Eckert, W. Johannsen, and J. Stegmann, “A database of age and gender annotated telephone speech,” in *Proc. The International Conference on Language Resources and Evaluation*, Malta, 2010.
- [27] B. Schuller, S. Steidl, A. Batliner, E. Nöth, A. Vinciarelli, F. Burkhardt, R. Van Son, F. Weninger, F. Eyben, T. Bocklet, *et al.*, “The interspeech 2012 speaker trait challenge,” in *Proc. International Speech Communication Association (INTERSPEECH)*, Portland, USA, 2012.
- [28] S. Deshmukh, B. Elizalde, R. Singh, and H. Wang, “Pengi: An audio language model for audio tasks,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [29] X. Mei, C. Meng, H. Liu, Q. Kong, T. Ko, C. Zhao, M. D. Plumbley, Y. Zou, and W. Wang, “Wavcaps: A ChatGPT-assisted weakly-labelled audio captioning dataset for audio-language multimodal research,” *arXiv preprint arXiv:2303.17395*, 2023.
- [30] Y. Chu, J. Xu, X. Zhou, Q. Yang, S. Zhang, Z. Yan, C. Zhou, and J. Zhou, “Qwen-audio: Advancing universal audio understanding via unified large-scale audio-language models,” *arXiv preprint arXiv:2311.07919*, 2023.
- [31] Z. Kong, A. Goel, R. Badlani, W. Ping, R. Valle, and B. Catanzaro, “Audio flamingo: A novel audio language model with few-shot learning and dialogue abilities,” *arXiv preprint arXiv:2402.01831*, 2024.