

Speech-Text Dialog Pre-training for Spoken Dialog Understanding with Explicit Cross-Modal Alignment

Tianshu Yu^{123*}, Haoyu Gao^{134*}, Ting-En Lin³, Min Yang^{1†},
Yuchuan Wu³, Wentao Ma³, Chao Wang³, Fei Huang³, Yongbin Li^{3†}

¹Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences

²University of Chinese Academy of Sciences

³Alibaba Group ⁴University of Science and Technology of China

{ts.yu, hy.gao, min.yang}@siat.ac.cn

{ting-en.lte,shuide.lyb}@alibaba-inc.com

Abstract

Recently, speech-text pre-training methods have shown remarkable success in many speech and natural language processing tasks. However, most previous pre-trained models are usually tailored for one or two specific tasks, but fail to conquer a wide range of speech-text tasks. In addition, existing speech-text pre-training methods fail to explore the contextual information within a dialogue to enrich utterance representations. In this paper, we propose **Speech-text dialog Pre-training for spoken dialog understanding with ExpliCiT cRoss-Modal Alignment (SPECTRA)**, which is the first-ever speech-text dialog pre-training model. Concretely, to consider the temporality of speech modality, we design a novel temporal position prediction task to capture the speech-text alignment. This pre-training task aims to predict the start and end time of each textual word in the corresponding speech waveform. In addition, to learn the characteristics of spoken dialogs, we generalize a response selection task from textual dialog pre-training to speech-text dialog pre-training scenarios. Experimental results on four different downstream speech-text tasks demonstrate the superiority of SPECTRA in learning speech-text alignment and multi-turn dialog context.¹

1 Introduction

In recent years, speech-text pre-training, which learns universal feature representations from a large training corpus (Chen et al., 2018; Li et al., 2021; Bapna et al., 2021), has achieved significant success in both uni-modal (Schneider et al., 2019; Dosovitskiy et al., 2020) and multi-modal (Lu et al., 2019; Radford et al., 2021) downstream tasks. Ex-

* Equal contribution. This work was conducted when Tianshu Yu and Haoyu Gao were interning at Alibaba.

† Min Yang and Yongbin Li are corresponding authors.

¹For reproducibility, we release our code and pre-trained model at: <https://github.com/AlibabaResearch/DAMO-ConvAI/tree/main/SPECTRA>.

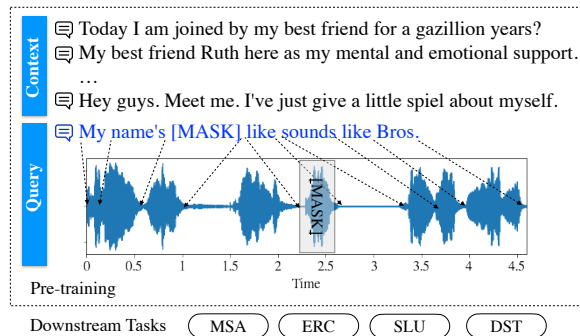


Figure 1: An illustration of SPECTRA, which considers dialogue context and explicit alignment between text and speech during pre-training, and generalizes well on various downstream tasks.

isting speech-text pre-training works mainly employed multi-modal self-supervised pre-training objectives, such as cross-modal masked data modeling (Li et al., 2021; Kang et al., 2022a) and cross-modal contrastive learning (Sachidananda et al., 2022; Elizalde et al., 2022), which align the speech utterance representation to the corresponding text sentence representation.

Despite the remarkable progress of previous speech-text pre-training models, there are still several technical challenges to constructing an effective and unified speech-text pre-training model for spoken dialog understanding, which are not addressed well in prior works. First, previous models are mainly tailored for specific speech-text tasks, such as speech-to-text translation (Liu et al., 2020b) and speech-language understanding (Chung et al., 2021), failing to conquer a wide range of speech-text tasks. Although Tang et al. (2022) proposed a unified speech-text pre-training for speech translation and recognition, it fails to exploit the temporality of an input speech sequence and cannot learn the fine-grained speech-text alignment.

Second, limited exploration has been attempted to bridge the gap between plain speeches/texts and human conversations. In particular, existing speech-

text pre-training methods fail to explore the context information within a dialog. Nevertheless, spoken dialog understanding needs to effectively process context information so as to help the system better understand the current utterance, since humans may omit previously mentioned entities/constraints and introduce substitutions to what has already been mentioned.

In this paper, we propose **Speech-text dialog Pre-training for spoken dialog understanding with ExpliCiT cROSS-Modal Alignment (SPECTRA)**, which is the first-ever speech-text dialog pre-training model. We illustrate the framework of our method in Figure 1 and details in Figure 2. The backbone of SPECTRA is composed of a text encoder, a speech encoder, and a fusion module, learning semantic/acoustic information and the interaction between them, and pre-trained on a large-scale real-world multi-modal (speech-text) dialog corpus. We propose two pre-training objectives to learn better context-aware speech/text representations for spoken dialog understanding (Dai et al., 2022; Zhang et al., 2022b). Specifically, to consider the temporality of speech modality, we design a novel temporal position prediction task to capture the speech-text alignment by predicting the start and end time of each textual word in the corresponding speech waveform. In addition, to learn the characteristics of spoken dialogs (Gao et al., 2023; Qian et al., 2023), we devise a cross-modal response selection objective to consider the context information within each dialog.

Our contributions are summarized as follows:

- To the best of our knowledge, we are the first to propose a speech-text dialog pre-training model for spoken dialog understanding, which fully exploits the characteristics of multi-modal (speech/text) dialogs.
- We introduce two pre-training objectives (temporal position prediction and multi-modal response selection) to effectively learn speech-text alignment and dialog context information.
- We conduct extensive experiments on five benchmark datasets belonging to four downstream speech-text tasks, including emotion recognition in conversation (ERC), multi-modal sentiment analysis (MSA), spoken language understanding (SLU), and dialog state tracking (DST). We believe that the release of

the pre-trained model and source code would push forward the research in this area.

2 Related Work

Uni-modal Pre-training In recent years, pre-trained language models (PLMs), such as BERT (Kenton and Toutanova, 2019), RoBERTa (Liu et al., 2019), and GPT (Radford et al., 2019a) have been proposed and applied to many NLP tasks, yielding impressive performances. PLMs benefit from the rich linguistic knowledge in large-scale corpora (He et al., 2022b,a). Inspired by the success of PLMs in NLP tasks, several speech pre-training models, such as Wav2vec (Schneider et al., 2019), HuBERT (Hsu et al., 2021), and WavLM (Chen et al., 2022), were proposed to learn high-quality universal speech representations from massive speech data.

Multimodal Pre-training Compared to multi-modal pre-training for vision-and-language tasks, speech-text pre-training is relatively less explored. SpeechBERT (Chuang et al., 2020) jointly trained multimodal representations based on a single BERT for spoken question-answering. CTAL (Li et al., 2021) extended the original Transformer to cross-modal by modifying the attention mechanism of the Transformer decoder. ST-BERT (Kim et al., 2021) combined a pre-trained acoustic model with BERT and took phoneme posterior and subword-level tokenized text as input. Kang et al. (2022b) explored multimodal pre-training model in extremely low-resource data scenarios. CLAM (Sachidananda et al., 2022) employed contrastive and multirate information inherent in audio and lexical inputs to align acoustic and lexical information. STPT (Tang et al., 2022) proposed a multi-task learning framework to integrate different modalities in speech-text pre-training.

Multimodal Dialog Systems The demand for multimodal dialog systems (Lin et al., 2022) is increasing due to the ubiquitous multimodal data. Liao et al. (2018) presented a knowledge-aware multimodal dialog (KMD) model, which leveraged reinforcement learning to generate human-like responses given multimodal (text-image) dialog context. Cui et al. (2019) considered the explicit user requirements in the attribute level and dynamically encoded the multimodal (text-image) dialog context based on users' attention. Sunder et al. (2022) proposed an end-to-end spoken language

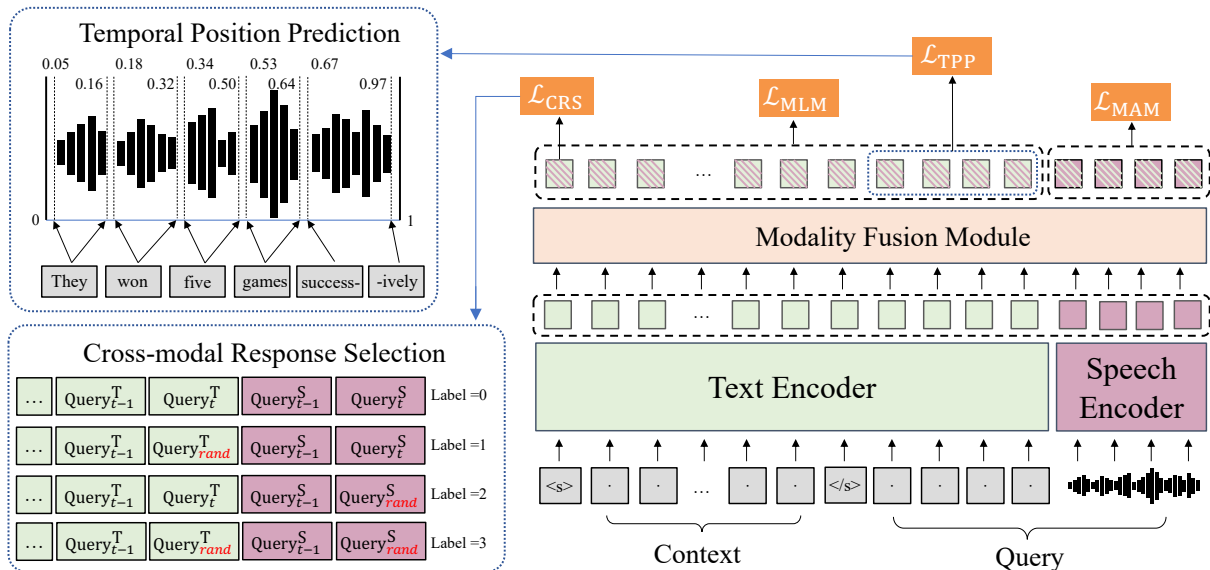


Figure 2: The overview of SPECTRA. The left part shows the illustration of the temporal position prediction task and the cross-modal response selection task. The right part shows the overall structure of the pre-trained model.

understanding model, which trained a semantically rich BERT-based conversation model along with a speech-based model.

Different from previous works, SPECTRA is the first-ever speech-text dialog pre-training model, which bridges the gap between plain texts/speeches and human conversations.

3 Method

In this section, we introduce the model architecture and pre-training objectives of SPECTRA.

3.1 The Backbone Architecture

Figure 2 shows the overall structure of our model SPECTRA, which consists of a text encoder, a speech encoder, and a modality fusion module. During pre-training, we first convert paired text and speech inputs into uni-modal embeddings, which are then fed into the text encoder and speech encoder respectively to obtain uni-modal representations. Finally, we concatenate text representations and speech representations as input of our modality fusion module to get fused representations for speech-text pre-training.

3.1.1 Data Preparation

Before diving into our model, we first prepare input text and speech sequences for our model. Let $D = \{T_1, T_2, \dots, T_n\}$ denotes a conversation with n dialog turns, where every single dialog turn T_i consists of a slice of raw speech waveform s_i and its corresponding text $\mathbf{t}_i = \{\mathbf{w}_{i1}, \mathbf{w}_{i2}, \dots, \mathbf{w}_{im}\}$.

Here, w_{ij} is the j -th word of \mathbf{t}_i , and is annotated with its corresponding start/end time in the speech, denoted as s_{ij}/e_{ij} . m is the sentence length of \mathbf{t}_i . For each dialog turn T_i where $i > 1$, we construct a sample \mathbf{X}_i with current utterance $T_i = \{\mathbf{t}_i, \mathbf{s}_i\}$, previous k ($k \geq 1$) turns of textual dialog history $\{\mathbf{t}_{i-k}, \dots, \mathbf{t}_{i-2}, \mathbf{t}_{i-1}\}$ and the previous speech dialog history \mathbf{s}_{i-1} . In this way, each sample \mathbf{X}_i consists of $k+1$ turns of text and 2 turns of speeches, where the speeches correspond to the latest 2 turns of text. Note that we only use 2 turns of speech in pre-training for efficiency, since the length of speech representation is much longer than its corresponding text representation.

3.1.2 Text Embeddings

For each input element, its vector representation is a summation of the corresponding token embedding, absolute position embedding and segment embedding. Specifically, we first concatenate all text sentences of each sample \mathbf{X}_i in temporal order to construct the text input: $I_i = \langle s \rangle \mathbf{t}_{i-k} \langle /s \rangle \mathbf{t}_{i-k+1} \langle /s \rangle \dots \langle /s \rangle \mathbf{t}_{i-1} \langle /s \rangle \mathbf{t}_i \langle /s \rangle$. Note that we use special token $\langle s \rangle$ to mark the start of the whole sequence, and $\langle /s \rangle$ to mark the end of each turn. Then, we encode each token in I_i using a pre-trained RoBERTa (Liu et al., 2019) tokenizer. We assign learnable segment embedding $\mathbf{e}_{t,1}$ to tokens of \mathbf{t}_i and the last $\langle /s \rangle$ token, and $\mathbf{e}_{t,0}$ for the rest of the tokens. The detailed tokenizing and encoding process is described in Appendix A. We denote \mathbf{x}_i as the input text embeddings of I_i .

3.1.3 Uni-modal Encoders

Text Encoder Inspired by the remarkable success of uni-modal pre-trained models on various downstream tasks, we employ RoBERTa (Liu et al., 2019) as our text encoder. We pass \mathbf{x}_i into text encoder to obtain the sequence representations:

$$\mathbf{H}_{t,i} = \text{RoBERTa}(\mathbf{x}_i) \quad (1)$$

where $\mathbf{H}_{t,i} \in \mathbb{R}^{n \times d_h}$ denotes the output hidden states of the last layer of RoBERTa, n is the length of input I_i , and d_h is the dimension of hidden state.

Speech Encoder We design our speech encoder based on the WavLM structure (Chen et al., 2022) with three key modules: a feature extractor, a feature projection module and a Transformer encoder module. The feature extractor consists of 8 temporal convolutional layers and a layer normalization. We implemented the first seven convolutional layers to be the same as WavLM, and added another convolutional layer with 512 channels, 5 strides and 5 kernels size, in order to shorten the length of the output speech features. As a result, each output token of speech features represents approximately 200ms of speech with a stride of 100ms.

The feature projection layer is a layer normalization followed by a fully connected layer converting the size of speech features from 512 to d_h . The Transformer encoder module is equipped with a convolution-based relative position embedding layer and 12 WavLM Transformer layers. For each sample, we directly input speech waveforms \mathbf{s}_{i-1} and \mathbf{s}_i into our speech encoder, and denote the outputs of the feature projection layer for \mathbf{s}_{i-1} and \mathbf{s}_i as f_{i-1} and f_i :

$$f_{i-1} = \text{Proj}(\text{Conv}(\mathbf{s}_{i-1})) \quad (2)$$

$$f_i = \text{Proj}(\text{Conv}(\mathbf{s}_i)) \quad (3)$$

Then, we obtain a speech sequence a_i by concatenating f_{i-1} and f_i together with a separation token [SEP] and a starting token [CLS]:

$$a_i = [\text{CLS}]f_{i-1}[\text{SEP}]f_i \quad (4)$$

where $a_i \in \mathbb{R}^{(m_{i-1}+m_i+2) \times d_h}$ denotes the concatenated sequence. m_{i-1} and m_i are the lengths of \mathbf{s}_{i-1} and \mathbf{s}_i , respectively. We pass a_i as the input of the Transformer encoder module to get the speech sequence representations:

$$\mathbf{H}_{s,i} = \text{WavLM}(a_i) \quad (5)$$

where $\mathbf{H}_{s,i} \in \mathbb{R}^{(m_{i-1}+m_i+2) \times d_h}$ denotes the hidden states of the last Transformer layer.

3.1.4 Modality Fusion Module

To integrate two modalities, we employ a single self-attention Transformer layer as our modality fusion module. We first concatenate the text sequence representation $\mathbf{H}_{t,i}$ and the speech sequence representation $\mathbf{H}_{s,i}$ together. Then, we assign text and speech representations with learnable modality embeddings $\mathbf{e}_{m,0}$ and $\mathbf{e}_{m,1}$ respectively, and add the modality embeddings to the concatenated representations as the input of our modality fusion module. Finally, we obtain output hidden representations of modality fusion module $\mathbf{H}_i \in \mathbb{R}^{(n+m_{i-1}+m_i+2) \times d_h}$ as the speech-text joint representations.

3.2 Pre-training Tasks

We introduce two novel pre-training objectives for our SPECTRA model, empowering SPECTRA to capture speech-text alignment and multimodal dialog context effectively.

3.2.1 Temporal Position Prediction

Existing speech-text pre-training works mainly learn from prior visual-text pre-training models. These works ignore that speeches are temporal sequences, and thus fail to learn fine-grained speech-text alignment. In this work, we propose a novel temporal position prediction (TPP) objective, which utilizes the textual part of the hidden representations \mathbf{H}_i to predict the starting and ending time of each word in the speech waveform.

In particular, for each word \mathbf{w}_{ij} in utterance \mathbf{t}_i with its start/end time annotations s_{ij}/e_{ij} , we denote its first/last token in \mathbf{H}_i as $\mathbf{h}_{s_{ij}}/\mathbf{h}_{e_{ij}}$. The goal of the TPP pre-training objective is to predict its starting and ending time in \mathbf{s}_i with $\mathbf{h}_{s_{ij}}$ and $\mathbf{h}_{e_{ij}}$, respectively. We use squared error loss to optimize the TPP task:

$$\mathcal{L}_{\text{TPP}}(t_i) = \frac{1}{2} \left(\left(\mathbf{W}_{\text{start}} \mathbf{h}_{s_{ij}} - \frac{s_{ij}}{L_a} \right)^2 + \left(\mathbf{W}_{\text{end}} \mathbf{h}_{e_{ij}} - \frac{e_{ij}}{L_a} \right)^2 \right) \quad (6)$$

where $\mathbf{W}_{\text{start}}, \mathbf{W}_{\text{end}} \in \mathbb{R}^{d_h \times 1}$ are learnable parameters. L_a is the maximum speech length limit. By normalizing s_{ij} and e_{ij} over L_a , we guarantee that the starting and ending time falls into $[0,1]$. Here, we only calculate the TPP loss for the words in the last two turns of dialog (i.e., \mathbf{t}_{i-1} and \mathbf{t}_i) for each sample \mathbf{X}_i . We calculate the average TPP loss

over all words within those two turns as the TPP loss of dialog \mathbf{X}_i :

$$\mathcal{L}_{\text{TPP}} = \frac{1}{l_{i-1} + l_i} \left[\sum_j \mathcal{L}_{\text{TPP}}(\mathbf{w}_{i-1,j}) + \sum_j \mathcal{L}_{\text{TPP}}(\mathbf{w}_{i,j}) \right] \quad (7)$$

where l_{i-1} and l_i denote the total lengths of transcripts \mathbf{t}_{i-1} and \mathbf{t}_i in sample \mathbf{X}_i .

3.2.2 Cross-modal Response Selection

Inspired by the success of response selection tasks in textual dialog systems (Bao et al., 2019), we design a cross-modal response selection objective. For each sample \mathbf{X}_i , we randomly replace the text query \mathbf{t}_i or speech query \mathbf{s}_i with the utterances or speech from other dialogs in the dataset. In this way, for each sample \mathbf{X}_i , we can obtain three kinds of corrupted samples as negatives: (1) only the speech query is randomly substituted; (2) only the text query is randomly substituted; (3) both text and speech queries are randomly substituted. Note that both text and speech queries remain unchanged as positive as illustrated in Figure 2

Since the output of the first <s> token can be viewed as the representation of the whole speech-text sample, we apply a softmax function following a fully connected layer on top of the hidden state of token <s> as a four-way classifier, predicting which case the current example belongs to. We utilize the cross-entropy loss to optimize the cross-modal response selection task, denoted as \mathcal{L}_{CRS} .

3.2.3 Cross-modal Masked Data Modeling

Following previous works (Li et al., 2021), we also adopt the cross-modal representations \mathbf{H}_f for cross-modal masked language modeling (CMLM) and cross-modal masked acoustic modeling (CMAM) objectives. For masked language modeling, we follow the setup of RoBERTa (Liu et al., 2019) to dynamically mask out textual input tokens with a probability of 15%. For masked acoustic modeling, we follow Baevski et al. (2020) and Liu et al. (2020a) to mask continuous speech frames.

We modify the implementation of the original masked acoustic modeling method in previous works to increase the average number of masked speech frames in each sample. We provide the details of masked acoustic modeling in Algorithm 1 in Appendix B. The speech token masking step is performed between the feature extractor and feature projection. We employ the cross-entropy loss for the CMLM task ($\mathcal{L}_{\text{CMLM}}$) and the mean absolute error loss for the CMAM task ($\mathcal{L}_{\text{CMAM}}$).

3.2.4 Joint Pre-training Objective

We combine four pre-training objectives to form a joint pre-training objective for speech-text pre-training:

$$\mathcal{L} = \alpha \mathcal{L}_{\text{TPP}} + \mathcal{L}_{\text{CRS}} + \mathcal{L}_{\text{CMLM}} + \mathcal{L}_{\text{CMAM}} \quad (8)$$

3.3 Fine-tuning on Downstream Tasks

We fine-tune SPECTRA on four downstream tasks, including multimodal sentiment analysis (MSA), emotion recognition in conversation (ERC), spoken language understanding (SLU), and dialog state tracking (DST).

We use the hidden state of <s> token in \mathbf{H}_i , denoted as \mathbf{h}_i , and pass it through a prediction head with two fully-connected layers and a GELU activation (Hendrycks and Gimpel, 2016) between them to get the prediction:

$$\mathbf{y}_i = \mathbf{W}^{(2)} \sigma(\mathbf{W}^{(1)} \mathbf{h}_i + \mathbf{b}^{(1)}) + \mathbf{b}^{(2)} \quad (9)$$

where σ denotes the GELU activate function, $\mathbf{W}^{(1)} \in \mathbb{R}^{d_h \times d_h}$, $\mathbf{W}^{(2)} \in \mathbb{R}^{d_o \times d_h}$, $\mathbf{b}^{(1)} \in \mathbb{R}^{d_h}$, $\mathbf{b}^{(2)} \in \mathbb{R}^{d_o}$ are new learnable parameters in the fine-tuning stage. The output size d_o for MSA task is 1, and for ERC and SLU it is the corresponding number of classes. We adopt the squared error loss as the fine-tuning loss function for MSA. The cross-entropy loss is utilized for the rest of tasks.

4 Experiments

4.1 Pre-training Data

In this paper, we adopt Spotify100K (Clifton et al., 2020) to pre-train SPECTRA, which is a real-world scene speech-text dialog dataset. Spotify100K contains 105,360 podcast episodes, with nearly 60,000 hours of speeches covering a variety of genres, subject matter, speaking styles, and structure formats. The corpus also provides automatically-generated word-level textual transcripts, marking the starting and ending time in the speech for each word.

For a fair comparison with previous speech-text pre-training studies, we only use the first 960 hours of speech as well as the corresponding transcripts to pre-train our SPECTRA model.

4.2 Experimental Setup

Baselines In addition to state-of-the-art downstream models tailored for MSA, ERC, SLU and DST (see Section 4.3-4.6), we also compare SPECTRA with three types of pre-training models, including the text modality pre-training model

| Task | Dataset | Metric | Previous SOTA | SPECTRA |
|---|-----------|-------------------|--|----------------------|
| Multimodal Sentiment Analysis (MSA) | MOSI | Acc ₂ | 84.40 (MIB (Mai et al., 2022)) | 87.50 (+3.10) |
| | MOSEI | Acc ₂ | 86.20 (BBFN (Han et al., 2021)) | 87.34 (+1.14) |
| Emotion Recognition in Conversation (ERC) | IEMOCAP | Acc | 66.52 (M2FNET (Chudasama et al., 2022)) | 67.94 (+1.42) |
| Spoken Language Understanding (SLU) | MIntRec | Acc ₂₀ | 72.16 (MAG-BERT (Rahman et al., 2020)) | 73.48 (+1.32) |
| Dialog State Tracking (DST) | SpokenWoz | JGA | 20.90 (SPACE+WavLM+TripPy (Si et al., 2023)) | 21.96 (+1.06) |

Table 1: The comparison between the key metrics of our model and the previous SOTA method on five datasets.

RoBERTa (Liu et al., 2019), speech modality pre-training model WavLM (Chen et al., 2022), and speech-text multimodal pre-training model CTAL (Li et al., 2021).

Experimental Settings during Pre-training We use the first 960 hours of speech and textual transcripts of Spotify100K dataset for pre-training. We cut the speech waveform into slices of a maximum length of 10 seconds and view each slice with the corresponding transcripts as a single dialog turn, forming 356,380 dialog turns in total. By using these dialogs and setting k to a maximum of 7, we construct 350,784 samples, where each sample consists of 2~8 dialog turns of texts and 2 turns of speeches.

Besides, we use pre-trained models **RoBERTa-base** and **WavLM-base+** to initialize our text and speech encoder, respectively. Since our speech encoder has one more convolution layer than **WavLM-base+**, we only initialize the first seven convolution layers with pre-trained parameters and randomly initialize the last layer. Both text and speech encoders have 12 Transformer layers with a hidden size d_h of 768. We pre-train our SPECTRA model for 100 epochs on 8 Tesla-A100 GPUs with a batch size of 20 per GPU. We use AdamW (Loshchilov and Hutter, 2018) to optimize our model with a peak learning rate of 1×10^{-4} and a linear warmup for the first 1% of updates.

Experimental Settings during Fine-tuning For SpokenWoz dataset, each dialog turn consist of two utterances, one from the user and the other from the system. For other datasets, each dialog turn is a single utterance. For all datasets we truncate the speech length of each dialog turn to a maximum of 10 seconds. We fine-tune our pre-trained checkpoint on each downstream dataset using an AdamW (Loshchilov and Hutter, 2018) optimizer with a peak learning rate of 2×10^{-5} and a cosine annealing warmup.

4.3 Fine-tuning on MSA

For MSA task (Hu et al., 2022), our model aims to predict the positive or negative sentiment polarities of the given multi-modal input. We conduct experiments on two multi-modal datasets MOSI (Zadeh et al., 2016) and MOSEI (Zadeh et al., 2018) to evaluate the effectiveness of our model for the MSA task. We adopt the accuracy over positive/negative sentiments classification (denoted as Acc₂) as the evaluation metric for our model and baselines. The experimental results are reported in Table 1.

From the results, we can observe that our model achieves substantially better performance than previous state-of-the-art (SOTA) methods on both datasets. In particular, for the MOSI dataset, the accuracy increases by 3.10% over the strongest baseline MIB (Mai et al., 2022). In addition, as shown in Table 2, our SPECTRA also significantly outperforms the speech modality pre-training model WavLM and speech-text pre-training model CTAL.

4.4 Fine-tuning on ERC

ERC task requires the model to predict the emotion category of an utterance given a speech clip with its transcripts and dialog history. Here, we fine-tune our model with the widely-used IEMOCAP dataset (Busso et al., 2008), and follow the settings with Chudasama et al. (2022) to perform a 6-way classification task. For each sample, we construct 11 turns of text and 2 turns of speech with a maximum text length of 512.

In Table 1, we report the accuracy of six-way classification for our model and previous SOTA method M2FNET (Chudasama et al., 2022). In addition, from Table 2, we can observe that our method outperforms uni-modal pre-training models, as well as speech-text pre-training baseline CTAL. Compared with the uni-modal baselines RoBERTa and WavLM, our model benefits from multi-modal pre-training tasks that capture interactions and alignments between modalities. Compared with CTAL, our model is equipped with

| Settings | MLM & MAM | TPP | CRS | Pre-training Data | Turns of Textual Dialog History | MSA | | ERC | SLU | DST |
|----------|-----------------|-----|-----|----------------------|--|------------------|------------------|--------------|-------------------|--------------|
| | | | | | | MOSI | MOSEI | IEMOCAP | MIntRec | SpokenWoz |
| | | | | | | Acc ₂ | Acc ₂ | Acc | Acc ₂₀ | JGA |
| RoBERTa | - | - | - | - | - | 85.67 | 85.88 | 64.53 | 71.24 | 20.76 |
| WavLM | - | - | - | - | - | 65.85 | 77.90 | 46.90 | 16.63 | - |
| CTAL | ✓ | - | - | 960h | - | 72.56 | 80.77 | 55.12 | 53.26 | 15.79 |
| SPECTRA | ✓ | ✓ | ✓ | 960h | 7 | 87.50 | 87.34 | 67.94 | 73.48 | 21.96 |
| (a) | | | | 0 | 1 | 82.16 | 84.30 | 33.17 | 69.21 | 17.59 |
| (b) | ✓ | ✓ | ✓ | 360h | 7 | 85.98 | 86.02 | 66.01 | 72.36 | 20.34 |
| (c) | ✓ | | ✓ | 960h | 7 | 85.52 | 86.19 | 66.15 | 71.69 | 20.87 |
| (d) | ✓ | ✓ | | 960h | 7 | 87.35 | 86.85 | 65.94 | 72.58 | 20.45 |
| (e) | ✓ | ✓ | ✓ | 960h | 1 | 87.20 | 86.93 | 64.98 | 73.03 | 19.78 |

Table 2: Ablation test results. Here, setting (a) and (b) mean w/o multi-modal pre-training and using less pre-training data. Setting (c), (d) and (e) indicate w/o TPP task, w/o CRS task and w/o full turns of textual history, respectively.

better speech-text alignment and multi-turn dialog context information with the help of TPP and CRS pre-training tasks.

4.5 Fine-tuning on SLU

We also conduct experiments on the spoken language understanding (SLU) task, which aims to predict the user intent (Lin and Xu, 2019) given a spoken utterance with the textual transcript. We use MIntRec (Zhang et al., 2022a) as the experimental dataset for SLU and adopt classification accuracy for the evaluation metric.

From Table 1 and 2, we can observe that SPECTRA obtains significantly better results than previous methods. In particular, our SPECTRA model improves the results of RoBERTa and the previous SOTA method MAG-BERT (Rahman et al., 2020) by 1.55% and 2.47%, respectively. Compared to WavLM and CTAL, our model can capture semantic information in textual data and the context information within each dialog.

4.6 Fine-tuning on DST

For dialogue state tracking, we use a large-scale, cross-modal dataset called SpokenWoz (Si et al., 2023). The dataset was collected by crowdsourcing recordings through phone calls using the Appen platform². Transcriptions were obtained using a commercial ASR system, and speech-text pairs were annotated using a schema similar to MultiWoz (Eric et al., 2019). SpokenWoz consists of 204k turns, 5.7k dialog, and 249 hours of recordings. We adopt joint goal accuracy (JGA) as the evaluation metric, which compares the predicted and ground-

truth dialogue states at each turn. We follow Trippy (Heck et al., 2020) and substitute its context model BERT with our SPECTRA model.

As shown in Table 1, our model outperforms the previous SOTA method, SPACE+WavLM+TripPy. In addition, our model also surpasses the three pre-training baselines by a noticeable margin. This demonstrates better speech-text alignment is critical to tackling complicated conversations.

5 Analysis

5.1 Ablation Study

To better understand the effectiveness of our SPECTRA pre-training method, we investigate the influence of pre-training components and dialog history on the overall performance of SPECTRA. We report the ablation test results in Table 2.

Impact of Pre-training To demonstrate the efficiency of multi-modal pre-training, we directly use uni-modal encoders and randomly initialize the modality fusion module. We observe a significant performance drop by comparing (a) “w/o multi-modal pre-training” to other pre-training settings on all five datasets. In particular, setting (a) directly collapses on the ERC task, which is a complicated and conversational scenario. This verifies the necessity of cross-modal pre-training and aligning speech-text modalities. In addition, by comparing SPECTRA and setting (b) “using less pre-training data”, we can find that using more pre-training data can further improve the performance of our model.

Impact of TPP and CRS By comparing the setting (c) “w/o TPP” to SPECTRA, the performances on all five datasets drop to different extents, which

²<https://appen.com/>

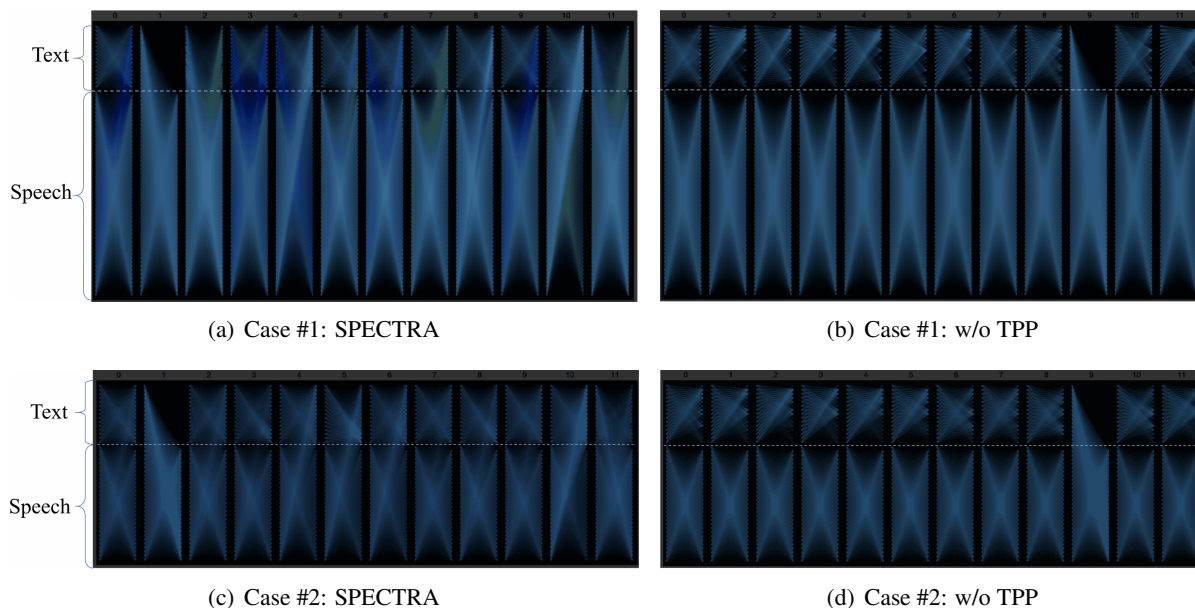


Figure 3: Visualization of self-attention weights of the fusion module in our model and the model pre-trained without TPP (w/o TPP). The upper and lower tokens stand for text and speech tokens, respectively.

| Case | Ground-Truth | By SPECTRA | By “w/o TPP” | Given Text |
|------|--------------|------------|--------------|--|
| #1 | Leave | ✓ Leave | ✗ Complain | I am going to sit in traffic for 45 minutes and return this. |
| #2 | Inform | ✓ Inform | ✗ Praise | That’s Tommy. He’s lead organizer, total badass. |

Table 3: Intent prediction results on test samples from the MIntRec dataset.

verifies the generalization and effectiveness of our TPP pre-training task. Specifically, the performance drops significantly on SpokenWoz, which requires the model to have a stronger ability to align two modalities. This demonstrates that our TPP pre-training task empowers the model with stronger alignment modeling ability. For setting (d) “w/o CRS” with SPECTRA, the performance drops significantly on multi-turn dialog tasks such as ERC and DST. This suggests that the CRS task is essential to model multi-turn dialog context.

Impact of Dialog History In setting (e) “using 1 turn of textual dialog history”, each instance consists of 2 turns of paired speech and text. The model performance drops substantially on ERC and DST downstream tasks by comparing it with SPECTRA. This demonstrates that increasing dialog history in the pre-training stage is beneficial to the tasks that require multi-turn dialog context.

5.2 Case Study

To have a straightforward understanding of how we learn cross-modal interaction in our proposed

SPECTRA model, we conduct a case study by providing two cases sampled from the MIntRec dataset. These two cases are incorrectly predicted by the model pre-trained without TPP but correctly predicted by our SPECTRA model. In Figure 3, we visualize the self-attention weights of the fusion layer in our model as well as the model pre-trained without TPP (denoted as w/o TPP). From Figure 3(a) and 3(c), we observe that there are rich cross-modal interactions in the fusion layer of the proposed SPECTRA model. Our model can capture fine-grained information between text and speech for more accurate classification. In contrast, we also visualize the self-attention weights of the w/o TPP model in Figure 3(b) and 3(d). Both cases show that text and speech sequences seldom connect to each other in self-attention layers.

In Table 3, we also illustrate the intent prediction results obtained by SPECTRA and w/o TPP. From the results, we can observe that our model can attend to both text and speech sequences effectively to predict correct intent results. However, w/o TPP is confused by the wrong labels since it hardly attends to speech tokens, which indicates that it

has the propensity to omit useful information that exists in speech exclusively.

6 Conclusion

In this paper, we proposed our model SPECTRA, the first speech-text dialog pre-training model. Considering the temporality of speech and text modalities, we introduced a novel temporal position prediction pre-training task to learn word-level speech-text alignment. To capture multi-modal dialog context in our model, we generalized the response selection task into multi-modal scenarios. Extensive experiments show that our pre-training method can learn better cross-modal interactions as well as multi-modal contextual information and significantly outperformed other strong baselines. In the future, we would like to extend speech-text dialog pre-training to more modalities or generative tasks.

Limitations

We analyze the limitations of this work, so as to further improve the performance of our model in future work. Based on our empirical observation, we reveal several limitations, which can be divided into two primary categories. (1) First, our proposed SPECTRA method relies on large-scale spoken dialog corpora with explicit word-level speech-text alignment annotation, such as Spotify100K. This limits the generality of our model on more spoken dialog corpora.

In the future, we would like to develop a semi-supervised pre-training method to leverage both labelled and unlabeled datasets. (2) Second, our method is mainly designed for speech-text understanding and has not been fully explored for generative tasks. We plan to devise dialog generation pre-training objective to empower the model with better generation ability. (3) Third, the work only involves speech and text modalities. We are interested in handling more modalities, such as images or videos, to enrich cross-modal information in joint representations.

Acknowledgements

Min Yang was partially supported by National Key Research and Development Program of China (2022YFF0902100), Shenzhen Science and Technology Innovation Program (KQTD20190929172835662), Shenzhen Basic Research Foundation (JCYJ20210324115614039 and JCYJ20200109113441941), and NSFC (no.

92270122). This work was supported by Alibaba Group through Alibaba Innovative Research Program.

References

- Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in Neural Information Processing Systems*, 33:12449–12460.
- Siqi Bao, Huang He, Fan Wang, Hua Wu, and Haifeng Wang. 2019. Plato: Pre-trained dialogue generation model with discrete latent variable. *arXiv preprint arXiv:1910.07931*.
- Ankur Bapna, Yu-an Chung, Nan Wu, Anmol Gulati, Ye Jia, Jonathan H Clark, Melvin Johnson, Jason Riesa, Alexis Conneau, and Yu Zhang. 2021. Slam: A unified encoder for speech and language modeling via speech-text joint pre-training. *arXiv preprint arXiv:2110.10329*.
- Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeanette N Chang, Sungbok Lee, and Shrikanth S Narayanan. 2008. Iemocap: Interactive emotional dyadic motion capture database. *Language resources and evaluation*, 42(4):335–359.
- Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, et al. 2022. Wavlm: Large-scale self-supervised pre-training for full stack speech processing. *IEEE Journal of Selected Topics in Signal Processing*, 16(6):1505–1518.
- Yi-Chen Chen, Chia-Hao Shen, Sung-Feng Huang, Hung-yi Lee, and Lin-shan Lee. 2018. Almost-unsupervised speech recognition with close-to-zero resource based on phonetic structures learned from very small unpaired speech and text data. *arXiv preprint arXiv:1810.12566*.
- Yung-Sung Chuang, Chi-Liang Liu, Hung-yi Lee, and Lin-shan Lee. 2020. Speechbert: An audio-and-text jointly learned language model for end-to-end spoken question answering.
- Vishal Chudasama, Purbayan Kar, Ashish Gudmalwar, Nirmesh Shah, Pankaj Wasnik, and Naoyuki Onoe. 2022. M2fnet: Multi-modal fusion network for emotion recognition in conversation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4652–4661.
- Yu-An Chung, Chenguang Zhu, and Michael Zeng. 2021. Splat: Speech-language joint pre-training for spoken language understanding. In *NAACL*.
- Ann Clifton, Sravana Reddy, Yongze Yu, Aasish Pappu, Rezvaneh Rezapour, Hamed Bonab, Maria Eskevich, Gareth Jones, Jussi Karlgren, Ben Carterette, and

- Rosie Jones. 2020. 100,000 podcasts: A spoken English document corpus. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5903–5917, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Chen Cui, Wenjie Wang, Xuemeng Song, Minlie Huang, Xin-Shun Xu, and Liqiang Nie. 2019. User attention-guided multimodal dialog systems. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 445–454.
- Yinpei Dai, Wanwei He, Bowen Li, Yuchuan Wu, Zheng Cao, Zhongqi An, Jian Sun, and Yongbin Li. 2022. Cgodial: A large-scale benchmark for chinese goal-oriented dialog evaluation. *arXiv preprint arXiv:2211.11617*.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Benjamin Elizalde, Soham Deshmukh, Mahmoud Al Ismail, and Huaming Wang. 2022. Clap: Learning audio concepts from natural language supervision. *arXiv preprint arXiv:2206.04769*.
- Mihail Eric, Rahul Goel, Shachi Paul, Abhishek Sethi, Sanchit Agarwal, Shuyang Gao, and Dilek Hakkani-Tür. 2019. Multiwoz 2.1: Multi-domain dialogue state corrections and state tracking baselines.
- Haoyu Gao, Rui Wang, Ting-En Lin, Yuchuan Wu, Min Yang, Fei Huang, and Yongbin Li. 2023. Un-supervised dialogue topic segmentation with topic-aware utterance representation. *arXiv preprint arXiv:2305.02747*.
- Wei Han, Hui Chen, Alexander Gelbukh, Amir Zadeh, Louis-philippe Morency, and Soujanya Poria. 2021. Bi-bimodal modality fusion for correlation-controlled multimodal sentiment analysis. In *Proceedings of the 2021 International Conference on Multimodal Interaction*, pages 6–15.
- Wanwei He, Yinpei Dai, Binyuan Hui, Min Yang, Zheng Cao, Jianbo Dong, Fei Huang, Luo Si, and Yongbin Li. 2022a. Space-2: Tree-structured semi-supervised contrastive pre-training for task-oriented dialog understanding. *arXiv preprint arXiv:2209.06638*.
- Wanwei He, Yinpei Dai, Yinhe Zheng, Yuchuan Wu, Zheng Cao, Dermot Liu, Peng Jiang, Min Yang, Fei Huang, Luo Si, et al. 2022b. Galaxy: A generative pre-trained model for task-oriented dialog with semi-supervised learning and explicit policy injection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 10749–10757.
- Michael Heck, Carel van Niekerk, Nurul Lubis, Christian Geisshauser, Hsien-Chin Lin, Marco Moresi, and Milica Gašić. 2020. Trippy: A triple copy strategy for value independent neural dialog state tracking. In *21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, page 35.
- Dan Hendrycks and Kevin Gimpel. 2016. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*.
- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed. 2021. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3451–3460.
- Guimin Hu, Ting-En Lin, Yi Zhao, Guangming Lu, Yuchuan Wu, and Yongbin Li. 2022. Unimse: Towards unified multimodal sentiment analysis and emotion recognition. *arXiv preprint arXiv:2211.11256*.
- Yu Kang, Tianqiao Liu, Hang Li, Yang Hao, and Wenbiao Ding. 2022a. Self-supervised audio-and-text pre-training with extremely low-resource parallel data. *arXiv preprint arXiv:2204.04645*.
- Yu Kang, Tianqiao Liu, Hang Li, Yang Hao, and Wenbiao Ding. 2022b. Self-supervised audio-and-text pre-training with extremely low-resource parallel data. *arXiv preprint arXiv:2204.04645*.
- Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186.
- Minjeong Kim, Gyuwan Kim, Sang-Woo Lee, and Jung-Woo Ha. 2021. St-bert: Cross-modal language model pre-training for end-to-end spoken language understanding. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7478–7482. IEEE.
- Hang Li, Wenbiao Ding, Yu Kang, Tianqiao Liu, Zhongqin Wu, and Zitao Liu. 2021. Ctal: Pre-training cross-modal transformer for audio-and-language representations. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3966–3977.
- Lizi Liao, Yunshan Ma, Xiangnan He, Richang Hong, and Tat-seng Chua. 2018. Knowledge-aware multimodal dialogue systems. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 801–809.
- Ting-En Lin, Yuchuan Wu, Fei Huang, Luo Si, Jian Sun, and Yongbin Li. 2022. Duplex conversation: Towards human-like interaction in spoken dialogue systems. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 3299–3308.

- Ting-En Lin and Hua Xu. 2019. Deep unknown intent detection with margin loss. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5491–5496.
- Andy T. Liu, Shu wen Yang, Po-Han Chi, Po chun Hsu, and Hung yi Lee. 2020a. Mockingjay: Unsupervised speech representation learning with deep bidirectional transformer encoders. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020b. Multilingual denoising pre-training for neural machine translation. *Transactions of the Association for Computational Linguistics*, 8:726–742.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Ilya Loshchilov and Frank Hutter. 2018. Fixing weight decay regularization in adam.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems*, 32.
- Sijie Mai, Ying Zeng, and Haifeng Hu. 2022. Multimodal information bottleneck: Learning minimal sufficient unimodal and multimodal representations. *IEEE Transactions on Multimedia*.
- Yushan Qian, Bo Wang, Ting-En Lin, Yinhe Zheng, Ying Zhu, Dongming Zhao, Yuexian Hou, Yuchuan Wu, and Yongbin Li. 2023. Empathetic response generation via emotion cause transition graph. *arXiv preprint arXiv:2302.11787*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019a. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019b. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Wasifur Rahman, Md Kamrul Hasan, Sangwu Lee, Amir Zadeh, Chengfeng Mao, Louis-Philippe Morency, and Ehsan Hoque. 2020. Integrating multimodal information in large pretrained transformers. In *Proceedings of the conference. Association for Computational Linguistics. Meeting*, volume 2020, page 2359. NIH Public Access.
- Vin Sachidananda, Shao-Yen Tseng, Erik Marchi, Sachin Kajarekar, and Panayiotis Georgiou. 2022. Calm: Contrastive aligned audio-language multi-rate and multimodal representations. *arXiv preprint arXiv:2202.03587*.
- Steffen Schneider, Alexei Baevski, Ronan Collobert, and Michael Auli. 2019. wav2vec: Unsupervised pre-training for speech recognition. *arXiv preprint arXiv:1904.05862*.
- Shuzheng Si, Wentao Ma, Yuchuan Wu, Yinpei Dai, Haoyu Gao, Ting-En Lin, Hangyu Li, Rui Yan, Fei Huang, and Yongbin Li. 2023. Spokenwoz: A large-scale speech-text benchmark for spoken task-oriented dialogue in multiple domains.
- Vishal Sunder, Samuel Thomas, Hong-Kwang J Kuo, Jatin Ganhotra, Brian Kingsbury, and Eric Fosler-Lussier. 2022. Towards end-to-end integration of dialog history for improved spoken language understanding. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7497–7501. IEEE.
- Yun Tang, Hongyu Gong, Ning Dong, Changhan Wang, Wei-Ning Hsu, Jiatao Gu, Alexei Baevski, Xian Li, Abdelrahman Mohamed, Michael Auli, et al. 2022. Unified speech-text pre-training for speech translation and recognition. *arXiv preprint arXiv:2204.05409*.
- Amir Zadeh, Rowan Zellers, Eli Pincus, and Louis-Philippe Morency. 2016. Mosi: multimodal corpus of sentiment intensity and subjectivity analysis in online opinion videos. *arXiv preprint arXiv:1606.06259*.
- AmirAli Bagher Zadeh, Paul Pu Liang, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. 2018. Multimodal language analysis in the wild: Cmu-mosei dataset and interpretable dynamic fusion graph. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2236–2246.
- Hanlei Zhang, Hua Xu, Xin Wang, Qianrui Zhou, Shaojie Zhao, and Jiayan Teng. 2022a. Mintrec: A new dataset for multimodal intent recognition. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 1688–1697.
- Sai Zhang, Yuwei Hu, Yuchuan Wu, Jiaman Wu, Yongbin Li, Jian Sun, Caixia Yuan, and Xiaojie Wang. 2022b. A slot is not built in one utterance: Spoken language dialogs with sub-slots. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 309–321, Dublin, Ireland. Association for Computational Linguistics.

Algorithm 1 Method to mask speech tokens

Require: the output of temporal convolution layer $\mathbf{f} = \text{Conv}(\mathbf{s})$ with the shape of $l \times 512$, where \mathbf{s} is the input speech waveform.

Ensure: Masked convolutional feature output $\tilde{\mathbf{f}}$ and masked indices $m_{\mathbf{f}}$.

```
1:  $\tilde{\mathbf{f}} = \mathbf{s}; m_{\mathbf{f}} = [0, 0, \dots, 0]$  with the length of  $l$ .
2: Randomly picks an integer  $n$  from  $[20, 50]$  as the length of masked continuous speech frames.
3: for  $i = 0; i < l; i ++$  do
4:   Draw a random number  $r$  from  $U(0, 1)$ ;
5:   if  $r < 0.15$  then ▷ Mask the continuous speech frames from index  $i$  to  $i + n - 1$ 
6:      $m_{\mathbf{f}}[i : i + n] = 1$ ;
7:     for  $j = 0; j < n$  and  $i + j < l; j ++$  do
8:       Draw a random number  $t$  from  $U(0, 1)$ ;
9:       if  $t < 0.8$  then
10:         $\tilde{\mathbf{f}}[i + j] = \mathbf{0}$ ;
11:       else if  $t < 0.9$  then
12:        Replace  $\tilde{\mathbf{f}}[i + j]$  with a random speech frame in  $\mathbf{f}$ ;
13:       end if
14:     end for
15:      $i = i + n - 1$ 
16:   end if
17: end for
```

A Implementation details of the tokenizer

We describe how we convert each I_i into input embeddings \mathbf{x}_i . First, we split the sequence I_i into list of tokens using a BBPE algorithm (Radford et al., 2019b) and convert each token into its index according to the dictionary of our tokenizer. Then, we pass the token indexes to the pre-trained token embedding layer of RoBERTa model to get the token embedding of each token. Finally, we sum up the token embedding, the absolute positional embedding and the segment embedding ($\mathbf{e}_{t,0}$ or $\mathbf{e}_{t,1}$) to get the input text embedding of every token in I_i .

B Method to Mask Speech Tokens

We report our method to mask speech tokens in Algorithm 1. We note that masked speech tokens are set to 0 at 80% of the time, a random token 10%, and an un-altered 10%. In our experiments, the maximum length of speech features f_i is 99 since the longest slice of our speech input is 10 seconds. We estimate the expectation of the number of masked frames of our method and the original MAM method proposed by Liu et al. (2020a) by simulating both masking steps for 1,000,000 times and calculating the average number of masked tokens. Simulation results show that our method masks approximately 57% of tokens in the se-

quence, while the original MAM method masks around 15%.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
We include the limitations of our work after Section 6.
- A2. Did you discuss any potential risks of your work?
There are no potential risks of this work.
- A3. Do the abstract and introduction summarize the paper’s main claims?
Left blank.
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Left blank.

- B1. Did you cite the creators of artifacts you used?
No response.
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
No response.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
No response.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
No response.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
No response.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
No response.

C Did you run computational experiments?

Sections 4 and 5

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Section 4.2

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
Section 4.2
- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?
Section 4
- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?
Section 4
- D** **Did you use human annotators (e.g., crowdworkers) or research with human participants?**
Section 4.6.
- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?
Due to limited space, we do not report the full text of instructions given to participants.
- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?
We crowdsourced the recordings through phone calls the Appen platform (<https://appen.com/>), and obtained the transcriptions via a commercial ASR system, and annotated speech-text pairs with a schema similar to MultiWoz.
- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?
We explain that the data is used for research purposes.
- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?
determined exempt
- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?
Since we use the Appen platform to do annotation, we do not know the basic demographic and geographic characteristics of the annotator population.