

# DUAL: Discrete Spoken Unit Adaptive Learning for Textless Spoken Question Answering

Guan-Ting Lin<sup>1</sup>, Yung-Sung Chuang<sup>2</sup>, Ho-Lam Chung<sup>1</sup>, Shu-wen Yang<sup>1</sup>, Hsuan-Jui Chen<sup>1</sup>, Shuyan Dong<sup>3</sup>, Shang-Wen Li<sup>3</sup>, Abdelrahman Mohamed<sup>3</sup>, Hung-yi Lee<sup>1</sup>, Lin-shan Lee<sup>1</sup>

<sup>1</sup>National Taiwan University, Taiwan  
<sup>2</sup>Massachusetts Institute of Technology, USA  
<sup>3</sup>Meta AI, USA

r10942104@ntu.edu.tw, shangwel@fb.com, abdo@fb.com, hungyilee@ntu.edu.tw

## Abstract

Spoken Question Answering (SQA) is to find the answer from a spoken document given a question, which is crucial for personal assistants when replying to the queries from the users. Existing SQA methods all rely on Automatic Speech Recognition (ASR) transcripts. Not only does ASR need to be trained with massive annotated data that are time and cost-prohibitive to collect for low-resourced languages, but more importantly, very often the answers to the questions include name entities or out-of-vocabulary words that cannot be recognized correctly. Also, ASR aims to minimize recognition errors equally over all words, including many function words irrelevant to the SQA task. Therefore, SQA without ASR transcripts (textless) is always highly desired, although known to be very difficult.

This work proposes Discrete Spoken Unit Adaptive Learning (DUAL), leveraging unlabeled data for pre-training and fine-tuned by the SQA downstream task. The time intervals of spoken answers can be directly predicted from spoken documents. We also release a new SQA benchmark corpus, NMSQA, for data with more realistic scenarios. We empirically showed that DUAL yields results comparable to those obtained by cascading ASR and text QA model and robust to real-world data.<sup>1</sup>

**Index Terms:** Spoken Question Answering, Textless NLP, Self-Supervised Representation

## 1. Introduction

Spoken Question Answering (SQA) aims to find the answer from a spoken document given a question in either text or spoken form. SQA is crucial for personal assistants when replying to the questions from the user’s spoken queries. Unlike many spoken language understanding tasks such as speech translation or intent classification, in which the required understanding of semantics is primarily on utterance level, the SQA task requires sophisticated comprehension and reasoning over much longer audio content. In addition to understanding the question and comprehending the global information in the audio context, it also needs to catch the fine-grained information to precisely locate the answer span out of the long audio context. Thus, SQA is known to be a very challenging task.

The conventional SQA system consists of the cascade of an Automatic Speech Recognition (ASR) engine followed by a text-trained QA (TQA) model. However, speech recognition errors naturally cause catastrophic problems to the TQA. [1],

<sup>1</sup>Dataset is collected and released by NTU. The dataset and source code is available at: <https://github.com/DanielLin94144/DUAL-textless-SQA>

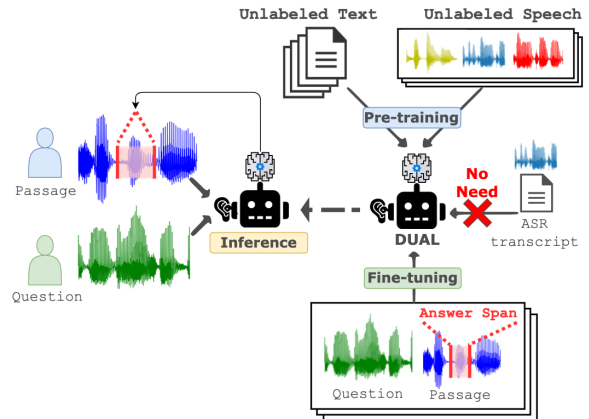


Figure 1: The proposed DUAL framework for textless (ASR transcript-free) SQA. All the passages, questions, and answers are in spoken form. The time intervals of the spoken answers can be extracted from the spoken passage without ASR transcripts.

and several works [2, 3, 4, 5] intended to alleviate such problems by knowledge distillation, including adapting the TQA model to be more robust against recognition errors. Some other efforts [6, 7] exploited paired speech and transcripts to construct a cross-modal speech and text pre-trained model with aligned semantics fine-tuned end-to-end, in which the speech recognition error problems can be mitigated to some degree, and SQA performance improved.

However, ASR errors remain to be a major problem for SQA tasks. The correct answers to the questions often include name entities or out-of-vocabulary (OOV) words that can never be recognized. The key information is thus inevitably lost when the audio signals are transformed into transcripts with errors, and there is no way to recover them in the following TQA stage. Also, the ASR engine was trained by minimizing the word error rate (WER), which was evaluated equally overall words including many function words irrelevant to the SQA task. So the cascade of two stages (ASR and TQA) individually optimized with two different criteria cannot perform as well as a single-stage global performance goal. So it is highly desired to capture the information directly, rather than from the ASR transcripts, and obtain overall performance not constrained by ASR accuracies.

Also, ASR engines have to be trained with vast quantities of human-annotated audio-transcript data, which are time-consuming and expensive to collect for the thousands of low-resourced languages over the world, when low and robust enough error rates are considered. Furthermore, there exist many languages without written form worldwide. All the above imply

technologies for ASR transcript-free (textless) SQA are highly desired although challenging.

In this work, we propose the first known textless (i.e., ASR transcript-free) SQA framework as in Figure 1. Inspired by the concept of *Textless NLP* [8, 9, 10, 11, 12, 13], which encodes speech signals into discrete units for modeling, and the *pre-trained language models transferability* of [14, 15, 16, 17, 18], we propose Discrete Unit Adaptive Learning (DUAL) for textless SQA. DUAL leverages pre-trained models to obtain quantized, length-condensed speech representations from audio signals and further adapts the pre-trained language model to achieve competitive SQA results without any ASR transcripts. The time span of the answer can be directly located from the audio context and played to the user, so the extracted answers do not suffer from speech recognition errors or out-of-vocabulary (OOV) problems because *NO* ASR is performed.

Furthermore, despite the increasing efforts to build SQA benchmark corpora [1, 19, 20, 21, 22, 23, 24], there is still a lack of natural and large-scale SQA datasets featuring real-world scenarios. For this purpose, we release a novel benchmark corpus, Natural Multi-speaker Spoken Question Answering (NMSQA). In this corpus, the test set was produced by human speakers, and the training and validation set were synthesized from Amazon Polly TTS service with industrial-grade quality. We also assign two different speakers to read the pairs of passage and question, examining whether our textless SQA system is speaker-independent.

The contributions of this paper are summarized below:

- We propose DUAL as the first known framework for textless SQA, not utilizing ASR transcripts and not suffering from ASR errors.
- We open-source the NMSQA dataset for SQA in real-world scenarios.
- DUAL achieved performance competitive to those obtained by cascading ASR and TQA, and significantly better when the word error rate exceeded 30 %.
- DUAL is more robust and retains the performance for the real-speaker testing set, which was not easily achievable for the cascade approach.

## 2. Method

### 2.1. Problem Formulation

The form of SQA dataset  $D$  is  $\{\mathbf{q}, \mathbf{p}, a\}$ , corresponding to the question  $\mathbf{q}$ , passage  $\mathbf{p}$ , and answer  $a$ , all in spoken form in this work. Our goal is to extract the starting and ending time  $(t_s, t_e)$ , denoted as the answer span  $a$ , from the spoken passage  $\mathbf{p}$  given the spoken question  $\mathbf{q}$ .

### 2.2. DUAL framework

The DUAL framework consists of the Speech Content Encoder (SCE) and Pre-trained Language Model (PLM) as in Figure 2 and introduced below.

#### 2.2.1. Speech Content Encoder (SCE)

The SCE transforms the question-passage audio waveform  $(\mathbf{q}, \mathbf{p})$  pair to sequences of discrete units  $(\mathbf{z}_q, \mathbf{z}_p)$ .

**Self-supervised Speech Representation:** A self-supervised speech pre-trained model can extract informative feature representations. We adopted the state-of-the-art self-supervised

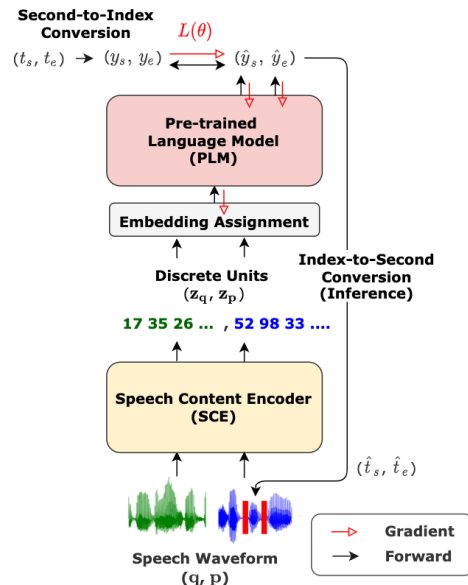


Figure 2: The overview of the DUAL framework.

speech pre-trained model HuBERT [25] for feature extraction<sup>2</sup>. HuBERT was trained by masked prediction objectives similar to BERT [27]. The prediction target was the K-means clustered index for speech signal processing features, e.g., Mel-frequency cepstral coefficients (MFCC) initially, and then learned latent representations after clustering in subsequent iterations. We utilized the HuBERT-Large model containing 24 transformer encoder layers pre-trained on LibriLight 60k hour dataset. HuBERT encoded the raw waveform into frame-level 1024 dimension features. Each frame was equivalent to 20 ms.

**Quantization:** The goal of quantization is to discretize the speech features so they can be fed into the pre-trained language model. K-means clustering was performed over the layer-wise representations of HuBERT-Large. We used LibriSpeech [28] 100-hour subset to train the K-means clustering model, and the number of clusters  $K$  is 64, 128, and 512. After clustering, the discrete units are represented by the clustering indices. The repetitive discrete units are merged to shorten the sequence length and remove the duration information, forming the dense discrete unit sequence of the question and passage  $(\mathbf{z}_q, \mathbf{z}_p)$ . We recorded the duration of duplication number of repetitions as  $\mathbf{c}_q$  and  $\mathbf{c}_p$  for  $\mathbf{z}_q$  and  $\mathbf{z}_p$ , so we can recover the frame-level indices to convert the answer span back to time interval at the inference stage.

#### 2.2.2. Pre-trained Language Model (PLM)

The learning model is a BERT-like transformer encoder model. The input was the discrete unit sequences of the spoken questions and passages  $(\mathbf{z}_q, \mathbf{z}_p)$ . Because SQA is a very challenging task to train from scratch, we leveraged the cross-disciplinary transferability of PLM [14, 15, 16, 17, 18] to help the SQA downstream task. Specifically, we used the weights of text PLM for network initialization, and randomly assigned the text pre-trained input embeddings for discrete units. The different random embedding assignments did not significantly affect the final performance. The input of PLM was the concatenated discrete unit sequences of the question and passage pair  $(\mathbf{z}_q, \mathbf{z}_p)$ , and the target was the start and end time  $(y_s, y_e)$  after the repetitions were reproduced.

<sup>2</sup>We used the open-source S3PRL [26] toolkit to extract the representations of the HuBERT-Large model.

Table 1: Word Error Rates on different datasets for the two open-sourced ASR models used in the baseline (cascade).

ASR	LS test-clean	NMSQA dev	NMSQA test
SB	3.1	15.6	61.7
W2v2	1.9	10.5	11.3

Because the length of a discrete unit sequence is much longer than the corresponding text, and the duration of a spoken passage itself is long, the standard maximal length of PLM (typically 512) is not enough in our case. As a result, we leveraged Longformer [29] to model the long  $(\mathbf{z}_q, \mathbf{z}_p)$ , which is a BERT-like model for long documents, pre-trained on the unlabeled long text documents and optimized for training efficiency by sparse attention mechanism, such as local and global attention, to support up to 4096 tokens.

### 2.2.3. Training Objective

The training objective is similar to the canonical QA fine-tuning in TQA. A randomly initialized linear layer is added on the top to predict the start and end time index. As the gradient flow shown in Figure 2,  $\theta$  represents the trainable weights of the model,  $\mathbf{c}_p = [\mathbf{c}_{p_1}, \mathbf{c}_{p_2}, \dots, \mathbf{c}_{p_n}]$  is the repetition of every discrete units  $\mathbf{z}_{p_i}$  in  $\mathbf{z}_p = [\mathbf{z}_{p_1}, \mathbf{z}_{p_2}, \dots, \mathbf{z}_{p_n}]$ .  $(t_s, t_e)$  is the ground truth start and end time in second, and  $(y_s, y_e)$  is the converted version on the index level. The overall training objective was to minimize the loss  $L(\theta)$  as the sum of the negative log probabilities of the true start and end indices on all the examples,

$$-\sum \log P(y_s | \mathbf{z}_q, \mathbf{z}_p; \theta) + \log P(y_e | \mathbf{z}_q, \mathbf{z}_p; \theta).$$

At the inference stage, we converted the predicted start and end indices  $(\hat{y}_s, \hat{y}_e)$  to the frame level by  $\mathbf{c}_p$  above, and finally transformed them to the time level  $(\hat{t}_s, \hat{t}_e)$ .

## 3. Experiments

### 3.1. Corpus Description

We developed and released here a new listening comprehension task named Natural Multi-speaker Spoken Question Answering (NMSQA). The train and dev set are the spoken version of the SQuAD v1.1 dataset, one of the largest QA datasets from Wikipedia paragraphs and human-written questions. We randomly split the SQuAD dev set into disjoint *SQuAD-dev-1* and *SQuAD-dev-2* for the NMSQA dev set and test set. The Amazon Polly Text-to-Speech service<sup>3</sup> was used for generating natural speech. We randomly assigned 12 TTS speakers and ensured that different speakers were used in producing each spoken document-question pair. Overall, there are 297.18 / 37.61 hours of audio for the train/dev set. Moreover, in order to have a realistic test set, 60 human speakers (30 male / 30 female) were requested to produce the *SQuAD-dev-2* naturally. This test set included 2.67 hours of audio. The answer intervals were annotated by Montreal Force Aligner [30].

### 3.2. Evaluation

Since the output target here is the temporal span of the spoken answer, following the evaluation metrics previously proposed by [1, 6], we adopted the Frame-level F1 score (FF1) and Audio Overlapping Score (AOS) to evaluate the performance based on the predicted time intervals. Higher FF1 and AOS scores imply more overlap between the predicted and ground truth spans.

<sup>3</sup><https://aws.amazon.com/tw/polly/>

Table 2: The performance of the proposed (DUAL) and baseline (cascade) approaches on the NMSQA dev and test sets. “Longformer<sup>†</sup>” indicates the Longformer model fine-tuned on clean text SQuAD-v1.1, while the normal “Longformer” was only pre-trained by unlabeled text data. The number after HuBERT (64, 128, 512) are numbers of clusters. “synth” and “human” represent the synthesized and human speech respectively.

Input	Model	dev (synth)		test (human)	
		FF1	AOS	FF1	AOS
<b>Baseline - Cascade (with ASR transcripts)</b>					
SB	Longformer <sup>†</sup>	56.7	49.7	17.3	15.3
W2v2	Longformer <sup>†</sup>	65.7	58.3	64.2	57.4
<b>Proposed - DUAL (without ASR transcripts)</b>					
HuBERT-64	Longformer	47.8	42.2	39.0	33.0
HuBERT-128	Longformer	54.2	48.5	<b>55.9</b>	<b>49.1</b>
HuBERT-512	Longformer	<b>55.0</b>	<b>49.6</b>	17.3	12.5

### 3.3. Baseline - Cascade (ASR plus TQA)

The conventional SQA approach cascading an ASR model and a TQA model was taken as the baseline to be compared here. Two open-sourced pre-trained ASR models were used in the first stage, one from Speechbrain [31]<sup>4</sup>, referred to as **SB**, the other the Wave2vec 2.0-large with self-training fine-tuning [32]<sup>5</sup>, referred to as **W2v2**. The Word Error Rates of them on different speech datasets are listed in Table 1. Both SB and W2v2 utilized LibriSpeech [28] 960-hour dataset as the supervised training data; however, we see in Table 1 W2v2 was much more robust than SB on the NMSQA test set, obviously because it leveraged 60k hrs of unlabeled data and the self-training procedure.

The TQA model, or the second stage of the baseline, is a Longformer-based model fine-tuned on SQuAD v1.1, denoted as Longformer<sup>†</sup> below. We used the online available model checkpoint<sup>6</sup> for TQA inference. The Longformer<sup>†</sup> obtained 91.5 F1 score and 85.1 EM (Exact Match) score on the text SQuAD v1.1 dataset. For the evaluation metrics used here, we adopted force alignment [30] to obtain the time intervals of the spoken answers in seconds.

### 3.4. Implementation Details of DUAL

We use the official Longformer-base model<sup>7</sup> as the PLM. The learning rate is searched in [3e-5, 5e-5, 7e-5, 1e-4], and we select models with the best performance on the validation set. The learning rate warmup step is 500, growing up linearly to the peak value and then linearly decaying to 0. All the DUAL experiments use 4 Tesla V100s with an overall 128 batch size for up to 5000 training steps. If the length of discrete units  $(\mathbf{z}_q, \mathbf{z}_p)$  input exceeds 4096, we truncate the passage  $\mathbf{z}_p$ .

## 4. Results

Encouraging results are reported here. Noting that the proposed (DUAL) approach achieves performance comparable to baselines (cascade), which require large ASR parallel data.

<sup>4</sup><https://huggingface.co/speechbrain/asr-crdnn-rnnlm-librispeech>

<sup>5</sup><https://huggingface.co/facebook/wav2vec2-large-960h-lv60-self>

<sup>6</sup><https://huggingface.co/valhalla/longformer-base-4096-finetuned-squadv1>

<sup>7</sup><https://huggingface.co/allenai/longformer-base-4096>

#### 4.1. “dev” set for synthesized speech

The experimental results are shown in Table 2. We first consider results for synthesized speech on the “dev” set, whose style is very similar to the training set, in the first column.

The top section of Table 2 is for the baseline (cascade) approach with ASR transcripts. We see ASR with W2v2 offered much better performance than that with SB (65.7 vs. 56.7 for FF1), obviously due to its lower recognition error in Table 1. On the other hand, for the proposed (DUAL) approach without ASR transcripts in the lower section of Table 2, we see DUAL achieved good FF1 scores (55.0, 54.2, 47.8 respectively for 512, 128, 64 clusters). Here, “HuBERT- $K$ ” denotes the input of DUAL is the clustering results from HuBERT-Large 22<sup>th</sup> layer representations with  $K$  clusters. These numbers are competitive to those in the top section for the baseline, verifying DUAL can learn semantics in the audio signals and find the answers almost as good as the case where ASR transcripts were available, though transcripts were not available at all. The relatively weak performance for 64 clusters suggested that the too small codebook size may lose important fine-grained content information. The situation can be improved significantly using larger codebook sizes (128 or 512 clusters).

#### 4.2. “test” set for human speakers

The experimental results on a more realistic scenario of human speech (“test”) are shown in the right column of Table 2. We observe that for the baseline (cascade) approach using SB for ASR on the top section, the performance dropped sharply due to the very high WER (61.7 in Table 2) on the human speech, while using the W2v2 for ASR model offered more robust results similar to the “dev” set. The result indicated that the performance of the baseline (cascade) approach relied heavily on the accuracy and robustness of ASR.

On the other hand, in the lower section of Table 2, we see the proposed (DUAL) approach could retain outstanding performance when  $K = 128$ , showing remarkable robustness of the approach for realistic human speech (55.9 of FF1 for  $K = 128$ ). However, the performance dropped drastically for  $K = 512$  (17.3 of FF1). The observation suggested that the cluster number played a crucial role in performance. We surmise that 128 clusters of quantization provide a smooth transformation from the machine-synthesized speech used in training to the human-speech test set. In contrast, 512 clusters may retain too many details that differentiate synthesized and human speech. This finding inspires more research to understand what makes textless SQA/NLP work or not.

## 5. Analysis and Discussion

**Ablation study for embedding assignment:** Table 3 shows the ablation study regarding how the discrete units should be assigned to the pre-trained embeddings, as shown in the middle of Figure 2. In this table on the top two rows (“Most frequent” and “Least frequent”), we randomly assigned the  $K$  (128) discrete units to the pre-trained embedding of the top- $K$  and the least- $K$  frequent vocabularies, where the vocabulary frequency was determined by Byte-Pair Encoding (BPE) on unlabeled text data. “Random” refers to randomly selecting pre-trained input embedding regardless of the frequency. “Re-init” denotes re-initializing the input embedding by a normal distribution. “Scratch” means the Longformer model was not pre-trained on the unlabeled text data. The results in Table 3 indicate that randomly assigning the pre-trained input embeddings for discrete units did not result in

Table 3: Ablation study on embedding assignment. All experiments used the HuBERT-128 setting. Performance was measured on the NMSQA dev set.

Embedding Assignment	FF1	AOS
Most frequent	54.2	48.5
Least frequent	46.9	41.7
Random	51.7	46.2
Re-init	8.9	7.2
Scratch (baseline)	6.1	4.9




Figure 3: Frame-level F1 (FF1) scores for DUAL and cascade approach (SB), evaluated on the small groups of full NMSQA dev set at different levels of ASR (SB) WER.

very different performance, although the “Most frequent” initialization offered the best results, which are those listed in Table 2. **Performance for Poor ASR Accuracy:** We compared the performance of the baseline cascade approach (SB for ASR which gave poor accuracy) and the proposed DUAL with HuBERT-128 for different levels of WER. Specifically, we bucketize the NMSQA dev set into subsets based on the WER (from 0% to 70%) obtained with ASR (SB). In Figure 3, we observe that for the baseline (cascade) approach, the FF1 score dropped significantly and continuously as the WER increased. This is the typical phenomenon of recognition error propagation. In contrast, the proposed (DUAL) attained very similar FF1 scores for different levels of WER, even when WER went up to 70%. Because there is no ASR in DUAL and no ASR transcripts were used, there was actually no correlation between WER and the FF1 score. The cascade approach outperformed DUAL when the WER was below 30%; but DUAL became much higher when WER exceeded 30%. Since the content of SQuAD (and thus NMSQA) is based on Wikipedia, it includes many words that are name entities, abbreviations, and OOV, which led to recognition errors. Many of these words are part of the answers. DUAL handle such scenario much better than cascade approaches.

## 6. Conclusion

We propose DUAL, the first textless (i.e., ASR transcript-free) SQA framework in this work. This framework only utilizes unlabeled speech and text data for pre-training and fine-tuning by the spoken questions, passages, and answer time intervals. DUAL directly predicts the answer time span without text supervision or acoustic word boundaries. Furthermore, we release NMSQA, a new natural, multi-speaker SQA benchmark corpus, which contains human speakers for the test set and large-scaled synthesized data for the training and development sets. The experiments showed that DUAL yields competitive results with the conventional cascade approach using ASR transcripts and is robust to real-world scenarios on NMSQA.

## 7. References

- [1] C.-H. Lee, S.-L. Wu, C.-L. Liu, and H. yi Lee, “Spoken squad: A study of mitigating the impact of speech recognition errors on listening comprehension,” in *INTERSPEECH*, 2018.
- [2] C.-H. Lee, Y.-N. Chen, and H.-Y. Lee, “Mitigating the impact of speech recognition errors on spoken question answering by adversarial domain adaptation,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 7300–7304.
- [3] C. You, N. Chen, and Y. Zou, “Knowledge distillation for improved accuracy in spoken question answering,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 7793–7797.
- [4] Y. Z. Chenyu You, Nuo Chen, “Mrd-net: Multi-modal residual knowledge distillation for spoken question answering,” in *IJCAI*, 2021.
- [5] D. Su and P. Fung, “Improving spoken question answering using contextualized word representation,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 8004–8008.
- [6] Y.-S. Chuang, C.-L. Liu, H. yi Lee, and L.-S. Lee, “Speechbert: An audio-and-text jointly learned language model for end-to-end spoken question answering,” in *INTERSPEECH*, 2020.
- [7] Y.-A. Chung, C. Zhu, and M. Zeng, “SPLAT: Speech-language joint pre-training for spoken language understanding,” in *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Online: Association for Computational Linguistics, Jun. 2021, pp. 1897–1907. [Online]. Available: <https://aclanthology.org/2021.naacl-main.152>
- [8] K. Lakhotia, E. Kharitonov, W.-N. Hsu, Y. Adi, A. Polyak, B. Bolte, T.-A. Nguyen, J. Copet, A. Baevski, A. Mohamed *et al.*, “Generative spoken language modeling from raw audio,” *arXiv preprint arXiv:2102.01192*, 2021.
- [9] A. Polyak, Y. Adi, J. Copet, E. Kharitonov, K. Lakhotia, W.-N. Hsu, A. Mohamed, and E. Dupoux, “Speech resynthesis from discrete disentangled self-supervised representations,” *arXiv preprint arXiv:2104.00355*, 2021.
- [10] E. Kharitonov, A. Lee, A. Polyak, Y. Adi, J. Copet, K. Lakhotia, T.-A. Nguyen, M. Rivière, A. Mohamed, E. Dupoux *et al.*, “Text-free prosody-aware generative spoken language modeling,” *arXiv preprint arXiv:2109.03264*, 2021.
- [11] F. Kreuk, A. Polyak, J. Copet, E. Kharitonov, T.-A. Nguyen, M. Rivière, W.-N. Hsu, A. Mohamed, E. Dupoux, and Y. Adi, “Textless speech emotion conversion using decomposed and discrete representations,” *arXiv preprint arXiv:2111.07402*, 2021.
- [12] A. Lee, P.-J. Chen, C. Wang, J. Gu, X. Ma, A. Polyak, Y. Adi, Q. He, Y. Tang, J. Pino *et al.*, “Direct speech-to-speech translation with discrete units,” *arXiv preprint arXiv:2107.05604*, 2021.
- [13] A. Lee, H. Gong, P.-A. Duquenne, H. Schwenk, P.-J. Chen, C. Wang, S. Popuri, J. Pino, J. Gu, and W.-N. Hsu, “Textless speech-to-speech translation on real data,” *arXiv preprint arXiv:2112.08352*, 2021.
- [14] I. Papadimitriou and D. Jurafsky, “Learning Music Helps You Read: Using transfer to study linguistic structure in language models,” in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Online: Association for Computational Linguistics, Nov. 2020, pp. 6829–6839. [Online]. Available: <https://aclanthology.org/2020.emnlp-main.554>
- [15] W.-T. Kao and H.-y. Lee, “Is BERT a cross-disciplinary knowledge learner? a surprising finding of pre-trained models’ transferability,” in *Findings of the Association for Computational Linguistics: EMNLP 2021*. Punta Cana, Dominican Republic: Association for Computational Linguistics, Nov. 2021, pp. 2195–2208. [Online]. Available: <https://aclanthology.org/2021.findings-emnlp.189>
- [16] C.-H. Chiang and H.-y. Lee, “On the transferability of pre-trained language models: A study from artificial datasets,” *arXiv preprint arXiv:2109.03537*, 2021.
- [17] K. Lu, A. Grover, P. Abbeel, and I. Mordatch, “Pretrained transformers as universal computation engines,” *arXiv preprint arXiv:2103.05247*, 2021.
- [18] R. Ri and Y. Tsuruoka, “Pretraining with artificial language: Studying transferable knowledge in language models,” in *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2022, pp. 7302–7315.
- [19] B.-H. Tseng, S. syun Shen, H. yi Lee, and L.-S. Lee, “Towards machine comprehension of spoken content: Initial toefl listening comprehension test by machine,” in *INTERSPEECH*, 2016.
- [20] P. Rajpurkar, J. Zhang, K. Lopyrev, and P. Liang, “SQuAD: 100,000+ questions for machine comprehension of text,” in *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Austin, Texas: Association for Computational Linguistics, Nov. 2016, pp. 2383–2392. [Online]. Available: <https://aclanthology.org/D16-1264>
- [21] C. You, N. Chen, F. Liu, D. Yang, and Y. Zou, “Towards data distillation for end-to-end spoken conversational question answering,” *ArXiv*, vol. abs/2010.08923, 2020.
- [22] C.-H. Lee, S.-M. Wang, H.-C. Chang, and H. yi Lee, “Odsqa: Open-domain spoken question answering dataset,” *2018 IEEE Spoken Language Technology Workshop (SLT)*, pp. 949–956, 2018.
- [23] A. Ravichander, S. Dalmia, M. Ryskina, F. Metzger, E. Hovy, and A. W. Black, “NoiseQA: Challenge Set Evaluation for User-Centric Question Answering,” in *Conference of the European Chapter of the Association for Computational Linguistics (EACL)*, Online, April 2021. [Online]. Available: <https://arxiv.org/abs/2102.08345>
- [24] F. Faisal, S. Keshava, A. Anastasopoulos *et al.*, “Sd-qa: Spoken dialectal question answering for the real world,” *arXiv preprint arXiv:2109.12072*, 2021.
- [25] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, “Hubert: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [26] S.-W. Yang, P.-H. Chi, Y.-S. Chuang, C.-I. J. Lai, K. Lakhotia, Y. Y. Lin, A. T. Liu, J. Shi, X. Chang, G.-T. Lin, T.-H. Huang, W.-C. Tseng, K. tik Lee, D.-R. Liu, Z. Huang, S. Dong, S.-W. Li, S. Watanabe, A. Mohamed, and H. yi Lee, “SUPERB: Speech Processing Universal PERFORMANCE Benchmark,” in *Proc. Interspeech 2021*, 2021, pp. 1194–1198.
- [27] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of NAACL-HLT*, 2019, pp. 4171–4186.
- [28] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “Librispeech: An asr corpus based on public domain audio books,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 5206–5210.
- [29] I. Beltagy, M. E. Peters, and A. Cohan, “Longformer: The long-document transformer,” *arXiv preprint arXiv:2004.05150*, 2020.
- [30] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, “Montreal forced aligner: Trainable text-speech alignment using kaldii,” in *INTERSPEECH*, 2017.
- [31] M. Ravanelli, T. Parcollet, P. Plantinga, A. Rouhe, S. Cornell, L. Lugosch, C. Subakan, N. Dawalatabad, A. Heba, J. Zhong, J.-C. Chou, S.-L. Yeh, S.-W. Fu, C.-F. Liao, E. Rastorgueva, F. Grondin, W. Aris, H. Na, Y. Gao, R. D. Mori, and Y. Bengio, “SpeechBrain: A general-purpose speech toolkit,” 2021, arXiv:2106.04624.
- [32] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 12 449–12 460.