



# Addressing Cold Start Problem for End-to-end Automatic Speech Scoring

Jungbae Park\*, Seungtaek Choi

Riiid AI Research

jbpark0614@gmail.com, seungtaek.choi@riiid.co

## Abstract

Integrating automatic speech scoring/assessment systems has become a critical aspect of second-language speaking education. With self-supervised learning advancements, end-to-end speech scoring approaches have exhibited promising results. However, this study highlights the significant decrease in the performance of speech scoring systems in new question contexts, thereby identifying this as a cold start problem in terms of items. With the finding of cold-start phenomena, this paper seeks to alleviate the problem by following methods: 1) prompt embeddings, 2) question context embeddings using BERT or CLIP models, and 3) choice of the pretrained acoustic model. Experiments are conducted on TOEIC speaking test datasets collected from English-as-a-second-language (ESL) learners rated by professional TOEIC speaking evaluators. The results demonstrate that the proposed framework not only exhibits robustness in a cold-start environment but also outperforms the baselines for known content.

**Index Terms:** automatic speech scoring, cold-start problem, multi-modal system

## 1. Introduction

With the rise of globalization and the accelerated adoption of online education following the COVID-19 pandemic, there is a growing demand for automatic speech assessment/scoring (ASA) systems that can help English-as-a-second-language (ESL) learners enhance their speaking proficiency. In the application of ASA, the performance is quite crucial for ESL learners since the predicted score of ASA directly impacts the learners' decision on the subsequent learning curricula and affects user segmentation of the intelligent tutoring systems (ITS).

Computer-assisted speech-scoring systems can be classified into two technical approaches: cascade systems, which rely on automatic speech recognition (ASR) followed by acoustic or linguistic analysis, and end-to-end scoring systems that are not dependent on any other posterior logic. Conventionally, cascade methods [1] were the general methods, while the end-to-end approaches had failed due to scarce gold labels from domain experts in speech scoring systems. The cascade methods utilize ASR models with the following score modules like GoP (goodness of pronunciation) [2], relating the test speech to an ASR model trained on native speech. Nevertheless, conversely, cascade methods also have limitations for their performances since the transcript text from ASR models for ESL learners can be distrusted due to their immature pronunciation [3, 4]. In addition, the score modules are needed to be constructed separately for each criterion. However, on the other hand, recently raised

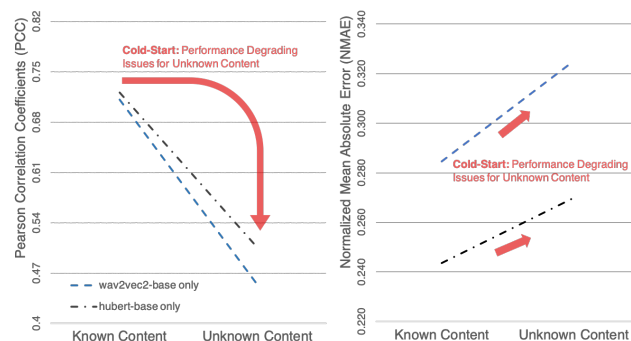


Figure 1: Cold start problem in end-to-end automatic speech scoring. The performances of ASA models directly finetuned from self-supervised acoustic models (Wav2Vec 2.0 and HuBERT) are dramatically decreased for unseen questions.

self-supervised pretrained acoustic models like contrastive predictive coding (CPC), Wav2Vec, Hubert, Data2Vec [5, 6, 7, 8, 9] successfully boost the accuracy of end-to-end scoring models [10] and make ease of ASA without other following logic like cascade systems. Because the pretraining methods do not rely on the score labels from evaluating experts for each test exam, it may relieve the sparsity problem of labels per speaker and makes scoring models robust for unseen speakers.

However, unfortunately, the previously investigated evaluations do not guarantee the consistency for unseen or newly added content in speech scoring systems. Since the spoken contexts from ESL learners should be changed according to the given contexts from each question, the ASA models also should be robust for each question context. However, as shown in Fig. 1, we find that the performances of end-to-end speech scoring models can be dramatically degraded for unseen content. The ‘item-wise cold-start problem’ is usually raised when the question content are added or updated in speech scoring systems. While the score labels are generally costly to be collected from experts, this cold-start problem cannot be quickly resolved and makes the system hard for updating content.

From our initial discovery of the cold-start issue, in this study, we suggest an evaluation strategy for verifying the performance of the speech scoring systems with the unseen content. After that, to be robust for the cold-start problem, we introduce methodologies addressing the cold-start problem with empirical experiments, including adding prompt embeddings, question content embeddings, and selecting the pretrained acoustic models. To the best of our knowledge, this is the first work on defining cold-start evaluation in speech scoring systems and introducing methods to address the item-wise cold-start problem.

\*Work done while affiliated with Riiid; now at Bucketplace Inc.

## 2. Related Work

The cold-start problem generally concerns the issue when the computational system cannot draw any inferences for users or items about which it has not yet gathered sufficient information [11, 12]. This issue arises when the system encounters a new user or item for which it has little or no prior knowledge. The cold-start problem is a common challenge in various domains, including recommendation systems, healthcare, advertising, and visual/textual/auditory recognition tasks [13, 14, 15, 16]. However, the specific perspectives and approaches to addressing this challenge may differ depending on the context and application of each domain.

In previous studies for ASA systems, the evaluation process follows general user splits (see Fig. 2 left) [10, 17, 18]. This evaluation is efficient since the characteristic of speakers, like tone, intonation from their nationality, and personalized traits, may affect the systems. Since its effectiveness, user-based splits also can be found in the evaluation of speaker-verification [19] and automatic speech recognition tasks. Compared with the previous works, in this study, the cold-start problem relates to the challenge of accurately evaluating the speaking abilities of new items (content) that the system has not previously assessed. In most ASA systems, machine learning algorithms may grade the speaking performance of test-takers based on various linguistic features and criteria extracted from the content. The item-wise cold-start problem can be particularly challenging in ASA systems, especially when the system needs to understand the context of each content, such as in the TOEIC speaking test. In such cases, accurate evaluation of ASA systems requires not only assessing pronunciation but also understanding the context of the spoken content.

## 3. Approach for Cold Start Problem

### 3.1. Split Methods for Cold Start in Speech Scoring

While the conventional user split-based evaluation process shows the generality of speakers' traits (see Fig. 2 left), as shown in Fig. 1, this evaluation strategy does not guarantee the efficiency and generality of unseen prompt content. As mentioned before, however, the item(content)-wise evaluation is also crucial for cold-start problems in the speech scoring system, where the prompt contexts are provided for each question to test-takers like TOEIC speaking. For assessment of the generality of newly added content of systems, we introduce item(content)-wise splits as shown in the right-side of Fig. 2. Since the consecutive testers on the same testing types are few in datasets, in this study, we split test sets from data (blue area in Fig. 2), which is isolated from training sets (yellow area in Fig. 2) with no intersection of prompt content and testers.

### 3.2. Model Architecture

The overview of the model architecture is shown in Fig. 3. The framework is composed as follows. First, the pretrained acoustic model creates acoustic features from spoken answers. Second, prompt contextualized vectors are extracted from content metadata. Lastly, a multi-modal fusion layer is applied across all embeddings and criterion tokens.

### 3.3. Choice of Pretrained Acoustic Model

While [10] shows self-supervised pretrained acoustic models enhance the performance of ASA system in the user-split setup, as shown in Fig. 1, it is not sufficient in unknown content.

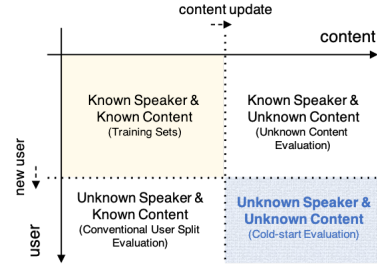


Figure 2: Dataset split strategy for ASA systems. While conventional user-splits only regards unknown speakers, content-wise splits should be considered in the context of content update.

However, if the acoustic model is pretrained to have linguistic understanding like Whisper [20], we find that the model outperforms in cold-start environments. In this study, we compare pretrained acoustic models, Wav2Vec 2.0 [7], HuBERT [8], Data2Vec [9] with Whisper, pretrained with automatic speech recognition data, uploaded in Huggingface hub<sup>1</sup>. Lastly, a linear acoustic adapter is added, and the final acoustic features are generated in the shape of  $N \times L_a \times H$ , if  $N$  is the batch size,  $L_a$  is the length of acoustic features, and  $H$  is hidden size.

### 3.4. Prompt Embedding

On English-speaking tests, there are various types of prompts to measure the skills of testers to contain different criteria and standpoints. TOEIC speaking exam has five types of prompts in a test set (see Table. 1). For understanding prompt contexts, several prompt embeddings are provided as followings.

#### 3.4.1. Question-type Embedding

The question-type encoder takes one hot encoded part and question number and is trained from scratch. The final question-type embedding is calculated by  $I_q = I_p + I_n$ , where  $I_p$  is part embedding and  $I_n$  is question number embedding.

#### 3.4.2. Transcript Embedding

When transcript text is given like part 1 in TOEIC speaking, character or phonetic level representations are essential for measuring pronunciation or intonations. Therefore, similar to [10, 18], the transcript embedding for phonetic-level representations is applied. Likewise, a character tokenizer is applied to embed character or phonetic levels for pronunciation assessment. During tokenization, to be matched with the feature vector shape, zero-padding or truncation is applied. For the embedding layer, the long short-term memory (LSTM) layer [21] is applied. The module is trained from random initialization. The final embedding vector is the shape of  $N \times L_t \times H$ , where  $L_t$  is a hyperparameter of the maximum length of transcript embeddings. If the content has no transcript text (part 2-5), the exact shape of zero vectors is added instead.

#### 3.4.3. Question Context Embedding

Unlike transcript text is given mainly for pronunciation measure, the semantic understanding of context embedding would be essential rather than character or phonetic level. We introduce pretrained models, BERT [22], or CLIP [23] to embed

<sup>1</sup><https://huggingface.co/models>

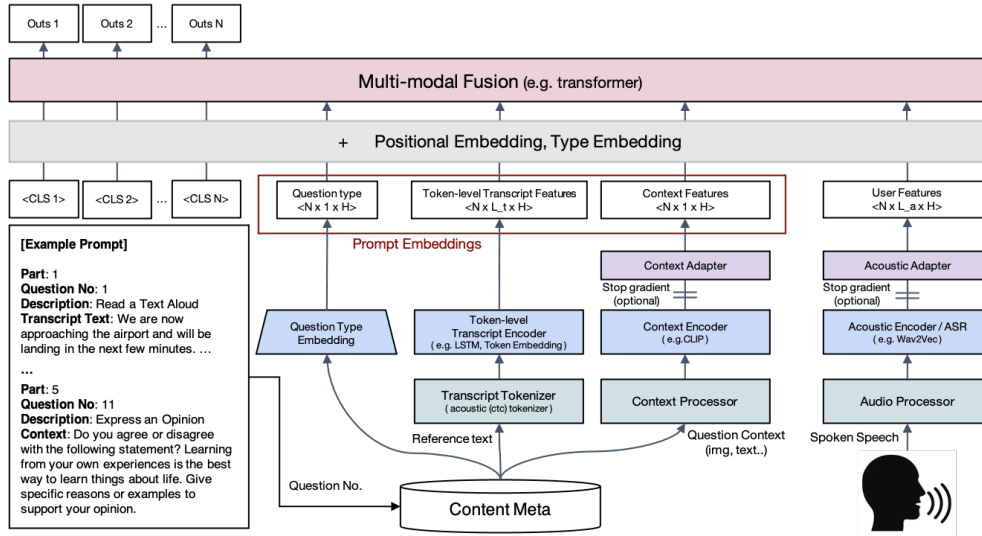


Figure 3: Proposed end-to-end speech scoring framework with the acoustic model, prompt embedding, and multi-modal fusion layer.

Table 1: Statistics of the collected dataset with prompt context information provided for 1,879 testees of the TOEIC speaking exam.

Part no.	Question no.	Prompt description	#Prompt and type	Criteria* with score statistics (criterion: mean $\pm$ standard deviation)	Score range per criterion	#Response
1	1~2	Read a text aloud	A transcript text	[ p: 1.96 $\pm$ 0.71, is: 1.97 $\pm$ 0.72 ]	~3.00	3,758
2	3~4	Describe a picture	An image	[ p: 2.00 $\pm$ 0.73, is: 1.93 $\pm$ 0.72, gv: 1.92 $\pm$ 0.69 ]	~3.00	3,758
3	5~7	Respond to questions	A question with a passage	[ p: 1.82 $\pm$ 0.89, is: 1.74 $\pm$ 0.87, gv: 1.70 $\pm$ 0.87 ]	~3.00	5,637
4	8~9	Respond to questions, using information provided	2 passages or tables with a question	[ p: 1.73 $\pm$ 0.97, is: 1.68 $\pm$ 0.94, gv: 1.64 $\pm$ 0.92, c: 1.66 $\pm$ 0.98, rc: 1.76 $\pm$ 1.06 ]	~3.00	3,758
5(a)	10	Express an opinion	Same passages of part 4, but a different question	[ p: 1.80 $\pm$ 0.91, is: 1.75 $\pm$ 0.90, gv: 1.73 $\pm$ 0.89, c: 1.77 $\pm$ 0.94, rc: 1.84 $\pm$ 1.00 ]	~3.00	1,879
5(b)	11	Express an opinion	A question	[ p: 2.51 $\pm$ 1.18, is: 2.46 $\pm$ 1.17, gv: 2.48 $\pm$ 1.17, c: 2.57 $\pm$ 1.25, rc: 2.66 $\pm$ 1.29 ]	~5.00	1,879

\*: 'p': pronunciation, 'is': intonation and stress, 'gv': grammar and vocabulary, 'c': cohesion, 'rc': relevance and completeness.

prompt question contexts. In TOEIC Speaking, the question contexts can be either textual passages (part 3-5) or images (part 2). Before embedding, all textual contexts are concatenated in advance to extract a representing feature vector. For BERT, we do not utilize the image feature. On the other hand, for CLIP, we utilize both textual and image features. Lastly, a linear context adapter is added, and the shape of question context embedding vectors is  $N \times 1 \times H$ . If all contexts are provided in transcript text, we simply append zero vectors.

### 3.5. Criterion Token Embedding

Since the English test requires measuring several criteria from the speeches of learners, so score labels are provided separately for each criterion (see Table 1), and the predictions for each criterion should be individual. Likewise, the training objectives also should be defined respectively. Similar to [18], we add start-tokens parameters like the other prompt embedding in a shape of  $N \times 1 \times H$  for each criterion. Unless all criteria are not used, the masking is applied to eliminate the unintended bias for other parts of the questions.

### 3.6. Multi-modal Fusion Layer with Additory Embeddings

To fusion the embeddings from acoustic, question prompt, and criterion token, we use a bi-directional transformer encoder [22, 24]. With these embeddings, trainable positional encoding ( $P_i$ )

and modality type encoding ( $T_i$ ) are added and inputs to the transformer. Finally, predicted scores of each criterion are from each linear output layer after the transformer layers.

## 4. Experiment and Result

### 4.1. Dataset and Experiment Setup

We privately collect data from 1,879 subjects and 20,669 vocal response samples. Each subject takes one of 12 TOEIC speaking test sets and must answer each prompt question, having the same preparation time and time limits. After then, professional raters for English tests assess each response according to official scoring criteria. The scoring criteria for each part are different, as shown in Table 1. These assessed scores are used for labels for finetuning and evaluations for the following experiments. We split datasets regarding users and items to evaluate our methods and baselines, as denoted in section 3.1. First, we isolate 3 test sets for item-wise cold start evaluations, leaving 9 test sets for training and conventional user-split tests. Then we randomly split train and test users in the ratio of 8:2.

For training, we employ the mean square error (MSE) loss with a fixed mini-batch size of 4 across all experiments. To optimize our models, we utilize the RAdam optimizer [25], which is a rectified version of Adam [26]. RAdam reduces the initial variance of gradients, mitigating initial convergence issues and offering robust training with reduced sensitivity to learning rate

Table 2: Performance (PCC) comparisons according to the type of the acoustic model and prompt embeddings.

Acoustic Model	Freeze?	Question-type Embedding	Transcript Embedding	Question Context Embedding	Test Avg. PCC in Known Content	Test Avg. PCC in Unknown Content
wav2vec2-base-960h	✓	✗	✗	✗	0.7134	0.5265
wav2vec2-base-960h	✓	✓	✗	✗	0.7394	0.5949
wav2vec2-base-960h	✓	✓	✓	BERT	0.7417	0.6041
wav2vec2-base-960h	✓	✓	✓	CLIP	<b>0.7576</b>	<b>0.6054</b>
hubert-base-ls960	✓	✗	✗	✗	0.7315	0.5710
hubert-base-ls960	✓	✓	✗	✗	0.7501	0.6201
hubert-base-ls960	✓	✓	✓	CLIP	<b>0.7513</b>	<b>0.6254</b>
data2vec-audio-base-960h	✓	✗	✗	✗	0.7144	0.5968
data2vec-audio-base-960h	✓	✓	✗	✗	0.7559	<b>0.6489</b>
data2vec-audio-base-960h	✓	✓	✓	CLIP	<b>0.7741</b>	0.6428
whisper-base	✓	✗	✗	✗	0.8004	0.7235
whisper-base	✓	✓	✗	✗	<b>0.8052</b>	0.6870
whisper-base	✓	✓	✓	BERT	0.7996	<b>0.7240</b>
whisper-base	✓	✓	✓	CLIP	<b>0.8052</b>	0.7108

Table 3: Performance (PCC) comparison as scaling acoustic model size, tested with CLIP context embeddings.

	whisper -tiny	whisper-tiny w/ CLIP	whisper -base	whisper-base w/ CLIP	whisper -small	whisper-small w/ CLIP	whisper -medium	whisper-medium w/ CLIP	whisper -large	whisper-large w/ CLIP	mean vanilla	mean whisper w/CLIP
Avg. PCC (known content)	0.7917	<b>0.8026</b>	0.8004	<b>0.8052</b>	0.8084	<b>0.8115</b>	0.8138	<b>0.8158</b>	0.8184	<b>0.8191</b>	0.8065	<b>0.8108</b>
Avg. PCC (unknown content)	0.6403	<b>0.6786</b>	<b>0.7235</b>	0.7108	0.6948	<b>0.7228</b>	0.7063	<b>0.7185</b>	0.7195	<b>0.7231</b>	0.6969	<b>0.7108</b>

and warm-up scheduling choices. Regarding hyperparameters, we set  $H$  to 512 and align  $L_a$  and  $L_t$  with the backbone model size or the maximum audio sample length of 60 seconds.

#### 4.2. Performance Comparison Result

To compare methods without parameter scales, we first experiment only with base size (74M parameters) pretrained acoustic models and investigate performance differences with prompt embeddings (see Table 2). From the base vanilla acoustic model, we gradually introduce each prompt embedding method and compare the performances with each other. For comparisons, we use the Pearson correlation coefficient (PCC), widely used in evaluation for speech scoring. Results show that the prompt embeddings are generally effective compared with the vanilla acoustic model, except for whisper-base. Furthermore, while performance is dramatically degraded for unknown content with vanilla Wav2Vec and HuBERT, prompt embeddings relieve this phenomenon. However, the results with Data2Vec indicate that context embedding also seems to help the performance increments in known content. On the other hand, the experiments with the whisper-base show that choosing the Whisper acoustic model can also be effective enough in known and unknown content. For Whisper, supervised multitask pretraining tasks, such as English transcription, speech translation, and phrase-level speech detection may help the learning of spoken contexts [20], so the performance is improved.

To analyze parameter scale effects and the exceptional cases in the Whisper-based model, we change the backbone acoustic model in various scales and compare vanilla methods with CLIP context embeddings (see Table 3). Overall, the ASA model performs better if the pretrained acoustic model is larger and

known to be better. Also, we find that except for Whisper-base, CLIP question context embedding increases performance as the other acoustic backbone cases in Table 2.

#### 4.3. Importance of Freezing the Acoustic Encoder

We investigate whether the performance ASA model can be potentially boosted when the acoustic encoder is also finetuned. However, unfreezing the backbone acoustic model has a detrimental effect on the overall performance (see Table 4). This phenomenon can be attributed to the crucial features of the audio encoder, which are utilized in decoding text for speech recognition tasks and play a significant role in the ASA model. Furthermore, the utilization of pretrained acoustic models in end-to-end ASA models may also leverage transcribed textual features, akin to cascade ASA methods. Additionally, it is noteworthy that the performance degradation is substantially more pronounced for unknown content compared to known content.

## 5. Conclusion

In this study, we raise the degrading performance problem in unknown content for ASA systems, item-wise cold start issue and propose an evaluation strategy for verifying the performance in cold start environments. Based on our evaluation, to be robust for the cold-start problem, we introduce potent methodologies addressing the issue: selecting the pretrained acoustic models and adding prompt and question content embeddings. We evaluate these approaches in conventional user splits with known content for models and cold start environments with collected TOEIC speaking data. Choosing a frozen Whisper-based acoustic encoder is the best choice for an acoustic encoder. Also, both question-type and question context embedding show effectiveness in both user-split and item-split setups. Especially these embeddings are more effective when the pretrained acoustic model only treats audio features like Wav2Vec or HuBERT. Based on our findings, our proposed approaches will be further extended with other context embedding and multi-modal fusions methodologies in future works for cold start problems in automatic speech scoring systems.

Table 4: PCC change after unfreezing the acoustic encoder

Acoustic Model	Test Avg. PCC in Known Content	Test Avg. PCC in Unknown Content
wav2vec2-base-960h	0.7576 → 0.5218	0.6054 → 0.4119
whisper-base	0.8004 → 0.7048	0.7235 → 0.5498
whisper-base w/ CLIP	0.7996 → 0.7110	0.7240 → 0.5710

## 6. References

- [1] P. Bamdev, M. S. Grover, Y. K. Singla, P. Vafaei, M. Hama, and R. R. Shah, "Automated speech scoring system under the lens: evaluating and interpreting the linguistic cues for language proficiency," *International Journal of Artificial Intelligence in Education*, pp. 1–36, 2022.
- [2] S. M. Witt and S. J. Young, "Phone-level pronunciation scoring and assessment for interactive language learning," *Speech communication*, vol. 30, no. 2-3, pp. 95–108, 2000.
- [3] M. S. Mirzaei, K. Meshgi, and T. Kawahara, "Exploiting automatic speech recognition errors to enhance partial and synchronized caption for facilitating second language listening," *Computer Speech & Language*, vol. 49, pp. 17–36, 2018.
- [4] S. Cheng, Z. Liu, L. Li, Z. Tang, D. Wang, and T. F. Zheng, "Asr-free pronunciation assessment," *arXiv preprint arXiv:2005.11902*, 2020.
- [5] A. v. d. Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.
- [6] S. Schneider, A. Baevski, R. Collobert, and M. Auli, "wav2vec: Unsupervised pre-training for speech recognition," *arXiv preprint arXiv:1904.05862*, 2019.
- [7] 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.
- [8] 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.
- [9] A. Baevski, W.-N. Hsu, Q. Xu, A. Babu, J. Gu, and M. Auli, "Data2vec: A general framework for self-supervised learning in speech, vision and language," in *International Conference on Machine Learning*. PMLR, 2022, pp. 1298–1312.
- [10] E. Kim, J.-J. Jeon, H. Seo, and H. Kim, "Automatic pronunciation assessment using self-supervised speech representation learning," *arXiv preprint arXiv:2204.03863*, 2022.
- [11] A. I. Schein, A. Popescul, L. H. Ungar, and D. M. Pennock, "Methods and metrics for cold-start recommendations," in *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*, 2002, pp. 253–260.
- [12] M. Saveski and A. Mantrach, "Item cold-start recommendations: learning local collective embeddings," in *Proceedings of the 8th ACM Conference on Recommender systems*, 2014, pp. 89–96.
- [13] Y. Tan, C. Yang, X. Wei, C. Chen, W. Liu, L. Li, J. Zhou, and X. Zheng, "Metacare++: Meta-learning with hierarchical subtyping for cold-start diagnosis prediction in healthcare data," in *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2022, pp. 449–459.
- [14] F. Pan, S. Li, X. Ao, P. Tang, and Q. He, "Warm up cold-start advertisements: Improving ctr predictions via learning to learn id embeddings," in *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2019, pp. 695–704.
- [15] H. Wu, F. Pitié, and G. Jones, "Cold start problem for automated live video comments," in *Proceedings of the Third Workshop on Multimodal Artificial Intelligence*, 2021, pp. 54–62.
- [16] X. Zhang and H. Zhao, "Cold-start recommendation based on speech personality traits," *Journal of Computational and Theoretical Nanoscience*, vol. 14, no. 3, pp. 1314–1323, 2017.
- [17] Y. K. Singla, A. Gupta, S. Bagga, C. Chen, B. Krishnamurthy, and R. R. Shah, "Speaker-conditioned hierarchical modeling for automated speech scoring," in *Proceedings of the 30th ACM international conference on information & knowledge management*, 2021, pp. 1681–1691.
- [18] Y. Gong, Z. Chen, I.-H. Chu, P. Chang, and J. Glass, "Transformer-based multi-aspect multi-granularity non-native english speaker pronunciation assessment," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 7262–7266.
- [19] G. Velius, "Variants of cepstrum based speaker identity verification," in *ICASSP-88, International Conference on Acoustics, Speech, and Signal Processing*. IEEE Computer Society, 1988, pp. 583–584.
- [20] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," *arXiv preprint arXiv:2212.04356*, 2022.
- [21] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [22] 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.
- [23] 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.
- [24] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [25] L. Liu, H. Jiang, P. He, W. Chen, X. Liu, J. Gao, and J. Han, "On the variance of the adaptive learning rate and beyond," *arXiv preprint arXiv:1908.03265*, 2019.
- [26] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.