

Unlocking Foundation Models for Privacy-Enhancing Speech Understanding: An Early Study on Low Resource Speech Training Leveraging Label-guided Synthetic Speech Content

Tiantian Feng, Digbalay Bose, Xuan Shi, Shrikanth Narayanan

¹Signal Analysis and Interpretation Laboratory, University of Southern California, USA

tiantiaf@usc.edu

Abstract

Automatic Speech Understanding (ASU) leverages the power of deep learning models for accurate interpretation of human speech, leading to a wide range of speech applications that enrich the human experience. However, training a robust ASU model requires the curation of a large number of speech samples, creating risks for privacy breaches. In this work, we investigate using foundation models to assist privacy-enhancing speech computing. Unlike conventional works focusing primarily on data perturbation or distributed algorithms, our work studies the possibilities of using pre-trained generative models to synthesize speech content as training data with just label guidance. We show that zero-shot learning with training label-guided synthetic speech content remains a challenging task. On the other hand, our results demonstrate that the model trained with synthetic speech samples provides an effective initialization point for low-resource ASU training. This result reveals the potential to enhance privacy by reducing user data collection but using label-guided synthetic speech content.

Index Terms: speech emotion recognition, spoken language understanding, privacy, synthetic data, foundation model

1. Introduction

¹Speech provides a natural way for us to express ourselves, enabling the effortless exchange of thoughts, emotions, and ideas. Automatic speech understanding (ASU) typically involves interpreting and comprehending human speech, bringing significant benefits in connecting individuals from diverse backgrounds and assisting individuals with speech impairments. Modern ASU systems typically leverage the power of deep learning [1] models for accurate and robust interpretation of human speech, leading to a wide range of speech applications that improve productivity, enhance the quality of life, and enrich the human experience. For example, popular virtual assistants, like Amazon Alexa, Apple Siri, and Google Assistant, equipped with advanced ASU models, empower numerous novel human experiences that substantially increase user satisfaction.

The success of such an ASU system crucially depends on the performance and robustness of the deployed deep learning models. One critical factor in ensuring the reliability of these deep learning models for ASU relies primarily on collecting qualitative training data. Typically, these datasets include large-scale speech samples from various speakers, languages, and acoustic conditions, leading models to learn diverse variations in speech patterns, intonations, and other environmental factors that may affect speech understanding. However, acquiring speech data frequently raises significant concerns regarding

data privacy. In addition to relevant information for developing the ASU system, speech data often carry sensitive information about a person, such as individual demographic attributes (e.g., age, gender), states (e.g., health), or biometric fingerprints (e.g., voice print) [2]. Therefore, it is essential for ML practitioners to investigate privacy-enhancing approaches for training ASU models.

Data transformation is widely used to prevent privacy inferences in training the ASU model. For example, researchers applied perturbations to suppress private information from the speech data while minimally interfering with the ASU task [3, 4, 5, 6, 7]. Recently, Federated Learning (FL) [8] has become an emerging privacy-enhancing learning algorithm that allows clients to train a model collaboratively without sharing their data. For instance, prior works explored the ability of FL on diverse ASU tasks, including keyword spotting [9], ASR [10], and emotion recognition [11]. Although numerous efforts have been made in privacy-enhancing research, researchers demonstrated that the above methods could still leak private information in every possible way [12, 13, 14, 15].

Recent advances in foundation models have empowered many evolving technologies, notably generative AI (e.g., ChatGPT and DALL-E-2², etc.), enabling the creation of high-fidelity content in formats of images, audio, and natural language based on user-input requirements or prompts. These advances in generative AI also present opportunities for privacy-enhancing computing, as high-quality generated content poses minimum privacy leaks while offering possibilities to serve as training data. For example, the prior study by [16] demonstrated that training with pure synthetic images could perform better than training with authentic images under zero-shot learning settings. Moreover, researchers studying large language models (LLMs) have broadly adopted synthetic text data as self-instructions for fine-tuning [17]. Finally, [18] also report that using synthetic speech samples in speech recognition training can improve ASR performance. However, these synthetic speech data are still generated from an existing text corpus.

Inspired by [19], in this paper, we provide an early investigation of using label-guided synthetic speech data for ASU training. Unlike previous studies that aim to protect privacy using data transformation or distributed algorithms, **our method focuses on unlocking the possibilities of using synthetic content as training data to enhance privacy, leading to the reduced need to collect user data or potential in local few-shot training.** Specifically, we perform a two-stage speech synthesis approach combining the LLMs with the text-to-speech (TTS) model. We begin with prompting the LLMs with the label-related information to generate spoken text data which

¹This work is supported by USC Amazon Center.

²<https://openai.com/>

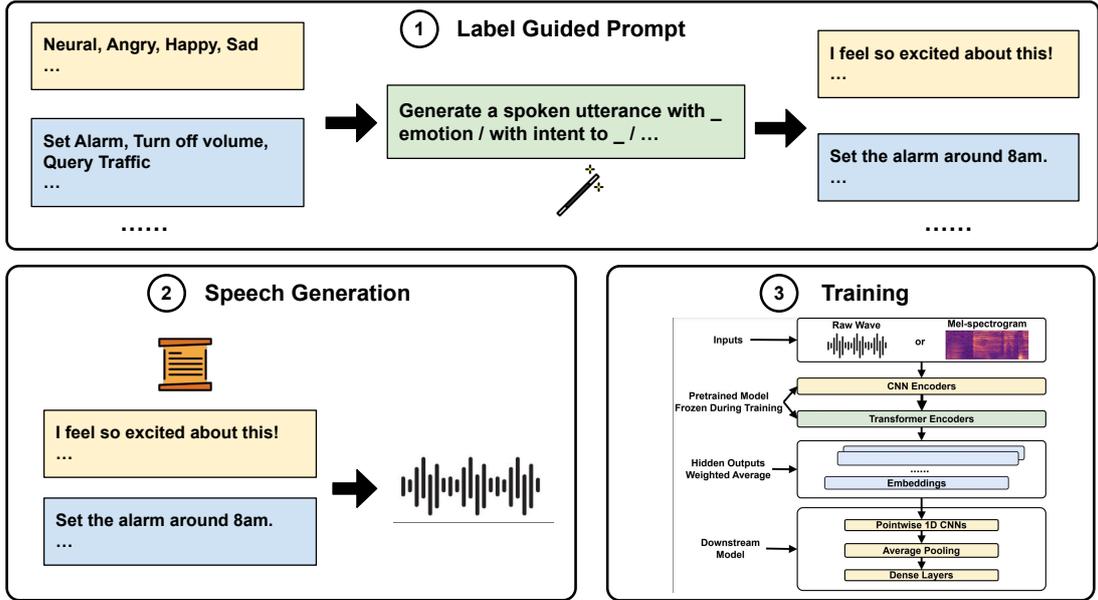


Figure 1: ASU training framework used for this work leveraging synthetic speech data. The framework begins with label-guided prompts that generate spoken utterances, followed by a text-to-speech model that creates synthetic speech data. The synthetic speech data is used to train the end-to-end ASU model. We use the pre-trained WavLM Base+ as the backbone for fine-tuning.

is then transformed into speech data using the text-to-speech (TTS) model. Our proposed method can synthesize speech data without the constraints of an existing text corpus. Specifically, our study aims to answer two fundamental and critical research questions related to privacy-enhancing speech understanding:

- Training ASU models with only synthetic data provide the strongest guarantees to user privacy as no actual user data is exposed. However, there is constrained knowledge about the training utilizing label-guided synthetic speech. **As a result, are label-guided synthetic speech content ready for training ASU models compared to real speech data?**
- Recent privacy laws, like the EU’s General Data Protection Regulation (GDPR) [20], have prevented the unauthorized collection or misuse of user data, leading to limited data for model training. It is more prevalent to train models with low-resource data to comply with privacy regulations. **Therefore, are label-guided synthetic speech content ready for low-resource speech training for ASU?**

2. ASU Tasks

This work focuses on two popular ASU tasks: speech emotion recognition (SER) and spoken language understanding (SLU).

Speech Emotion Recognition [21] aims to automatically classify expressed emotions (e.g., neutral, sad, and happy) from speech signals, where the categorical emotion labels are typically acquired from human annotators. Understanding emotions from human conversations has broader applications in virtual assistants, education, and healthcare.

Spoken Language Understanding [22] is the task that involves processing speech utterances for domain classification, intent detection, and slot filling. There is a steadily growing interest in deploying robust and efficient SLU systems on mobile and edge applications. Conventional SLU systems typically transcribe the speech signal into spoken format text using ASR models, followed by natural language understanding (NLU) models to recognize the domain and intent. In recent years, researchers

have focused on developing end-to-end recipes that directly map speech signals to classify SLU related tasks [23]. In this study, we focus on using end-to-end models for SLU tasks.

3. Two-stage Speech Content Synthesis

In this work, we propose a two-stage speech synthesis framework as demonstrated in Figure 1.

Spoken Text Generation We first apply label information to generate spoken text using LLMs. As shown in Figure 1, the process involves generating spoken utterances that convey particular emotions with prompt messages: Generate a spoken utterance with _ emotion, where _ represents categorical emotion labels, such as neutral, happy, and sad. When we aim to generate spoken utterances with intent, we adopt the following prompt message: Generate a spoken utterance with intent to .. In this study, we use the FLAN-T5 [24] as the language model that consists of 11.3B parameters.

Text-to-Speech The second stage in speech synthesis involves text-to-speech module. We use the prompted output from spoken text generation as the input to the text-to-speech model to directly synthesize speech data. Specifically, we utilize the recently released SpeechT5 model [25] for text-to-speech. We augment the speaker information by sampling the x-vector [26] from The CMU Arctic dataset [27].

4. Modeling Approach

Pre-trained Speech Backbone Our downstream modeling approach utilizes the widely adopted pre-trained speech models, WavLM, as the backbone. WavLM is a self-supervised model that optimizes multiple training objectives, including masked speech prediction, masked speech denoising, and frame prediction. This model is effective in a wide range of ASU tasks [28]. The WavLM base encoder includes 12 encoding layers with approximately 90M parameters.

End-to-End Downstream Modeling Our modeling approach

Table 1: Summary of dataset statistics used in this work.

Datasets	Unique Speakers	Classes	Total Utterances
IEMOCAP	10	4	5,531
MSP-Improv	12	4	7,798
SLURP	177	46	72,277

draws inspiration from [29], which highlights that combining the hidden outputs from all encoder layers provides substantially higher performance than relying on the last hidden output for the downstream speech task. In addition, [29] also shows that fine-tuning pre-trained speech models with the backbone encoder frozen provides simple but competitive results compared to unfreezing the encoder. In summary, our end-to-end modeling framework starts with weighted averaging to combine the hidden outputs from all encoder layers, where the weights are parameterized. Our downstream model consists of two 1D pointwise convolutional layers with a kernel size of 1 and a filter size of 256. The convolutional layers are connected with the ReLU activation functions. A global average pooling is applied to the convolutional layer output, leading to an output vector of size 256. Following this, the output vector is fed into two fully connected layers for the classification tasks in ASU.

Parameter-Efficient Fine-tuning In addition to relying on training downstream models, we decide to incorporate LoRA (Low-rank Adaptation) [30] in the fine-tuning stage. LoRA adapts the model updates with low-rank matrices during the training phase, bringing benefits to lower inference latency and ease of optimization. Previous work has demonstrated the effectiveness of applying LoRA in fine-tuning ASU tasks [31].

5. Datasets

Table 1 displays data statistics for the three datasets included in this work. We use IEMOCAP and MSP-Improv datasets for SER experiments and the SLURP dataset for SLU training.

Speech Emotion Recognition IEMOCAP [32] used in this work contains multi-modal (motion, audio, and video) recordings of acted human interactions from ten subjects, with half males and half females. Moreover, MSP-Improv [33] corpus is developed with the target of investigating naturalistic emotions that were elicited from improvised situations. Similar to the IEMOCAP dataset, the corpus comprises audio and visual data collected from 12 individuals, with an equal number of subjects in both genders. Due to imbalanced label distribution in both datasets, we decided to keep the four most frequently presented emotions for all the datasets, as recommended in [29].

Spoken Language Understanding SLURP [34] is the spoken language dataset collected for designing in-home personal robot assistants. The speech data were recorded from over 100 participants reading the collected text prompts carefully crafted by Mechanical Turk (AMT) workers. An example prompt with the intent to ask for time is: How would you ask for the time? Each recording belongs to one unique intent from 46 intent classes. The complete dataset contains a total of 72,277 utterances. We studied the intent classification using the SLURP dataset.

6. Results

6.1. Experiment Details

Data Split: We apply 5-fold and 6-fold evaluation on IEMOCAP and MSP-Improv datasets, where each session is regarded as a unique test fold. During each training fold, one

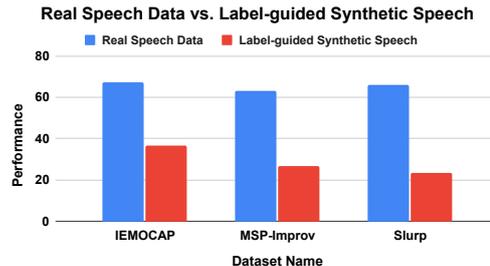


Figure 2: Baseline fine-tuning performance between real speech data and synthetic speech data.

data session is used for validation, while the rest are for training. In the SLURP study, we perform the experiment using the standard splits for training, validation, and testing.

Speech Generation: We generate 1,000 spoken text of each emotion class belonging to neutral, sad, happy, and angry, leading to 4,000 synthetic utterances for SER training. On the other hand, we generate 100 spoken texts with each unique intent associated with the SLURP dataset, resulting in 4,600 synthetic speech samples for training SLU. We control the maximum output token size as 32 and model temperature as 1.0.

End-to-end Training: We set the batch size as 64 in baseline experiments, including fine-tuning both natural and synthetic data. Specifically, we set the learning rate as 0.0005 and the maximum training epoch as 30 for SER training. We use a learning rate of 0.005 and a maximum training epoch of 50 for SLU training. We set the maximum audio duration as 6 seconds and 3 seconds for SER and SLU datasets, respectively. All experiments are implemented using PyTorch. The experiments are conducted on a high-performance computing server with A40 GPUs. We use the checkpoints of each pre-trained model from HuggingFace [35].

6.2. Examples of Spoken Text by Prompt

We provide examples of spoken text generated by prompting the LLM to showcase the quality of text input to the TTS module.

SER Examples We provide one example of each emotion:

- Neutral: "Okay, can you finish a glass of wine, please."
- Happy: "We had so much fun in Florida."
- Sad: "All I want to do is cry."
- Angry: "Why can't you just admit that you broke my car!"

SLU Examples We provide one example with each intent to set alarms, query contact, and mute the volume:

- Set Alarms: "You want the alarm to go off in 3 hours."
- Query contact: "I'd like to get a list of contacts who are in the San Francisco area."
- Mute the volume: "Turn your volume down."

6.3. Are label-guided synthetic speech content ready for training ASU models?

This subsection compares the training performance between real speech data and label-guided synthetic speech data. The results are presented in Figure 2. We use the unweighted average recall (UAR) and Macro-F1 to assess the performance of SER and SLU tasks, respectively. The results demonstrate that zero-shot performance with label-guided synthetic speech data is notably worse than training with real speech data in both SER and SLU tasks. This finding aligns with a prior study [36] that reported a substantial decline in performance when training exclu-

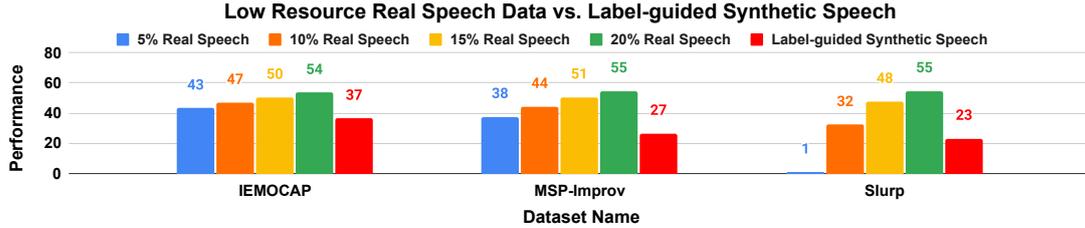


Figure 3: Fine-tuning performance between low-resource real speech data and synthetic speech data.

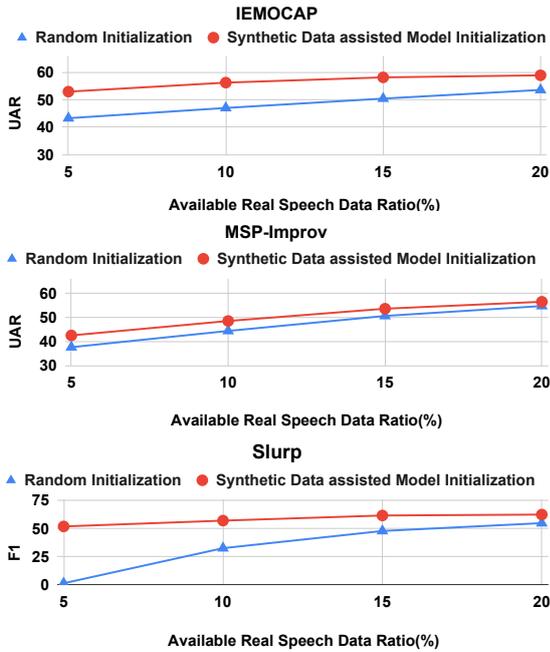


Figure 4: Fine-tuning performance between random initialization and synthetic data assisted model initialization in low-resource speech training. The x-axis represents the available real speech data ratio presented for low-resource training.

sively with synthetic speech data in ASR tasks. Consequently, it implies that even achieving the same level of performance with real speech training is still challenging using zero-shot learning with label-guided synthetic speech content.

6.4. Are label-guided synthetic speech content ready for low-resource speech training for ASU?

Here, we further investigate the role of label-guided synthetic speech content in low-resource speech training. We reduce the learning rate in training SER models to 0.0001 to prevent overfitting from fewer training samples.

Zero-shot Learning with Synthetic Speech still underperforms Low Resource Speech Training: We first compare the zero-shot performance with label-guided synthetic speech content and regular training with low-resource speech content. The comparison in Figure 3 reveals that, in most cases, zero-shot performance with label-guided synthetic speech content fails to provide competitive performance to low-resource speech training. The only exception to this trend is observed when utilizing 5% real speech with SLURP data for the SLU task. This finding further strengthens the implication that zero-shot learning with label-guided synthetic speech content training is challenging.

Low Resource Speech Training Leveraging Label-guided

Synthetic Speech Can Substantially Boost the Performance:

We further investigate low-resource speech training leveraging the label-guided synthetic speech content. Inspired by [19], we employ the model trained with label-guided synthetic speech samples as the initialization point for low-resource speech training. Figure 4 illustrates the performance differences between random initialization and synthetic data assisted initialization. The results indicate that, with an equal amount of real training speech data, synthetic data assisted initialization yields a substantial performance improvement compared to random initialization. Moreover, synthetic data assisted initialization leads to comparable ASU performance between training using 20% and full data (e.g., 20% data and full data yield 62.32% and 66.01% F1 scores on SLURP data, respectively). These findings highlight that although it is challenging to train the ASU model using label-guided synthetic data in lower-resource settings, the synthetic data assisted initialization proves to be effective, significantly enhancing the low-resource training performance.

7. Implications For Enhancing Privacy

Privacy-Enhancing Speech Training Leveraging Label-guided Synthetic Speech content is Feasible. It is promising that under the extremely low-resource condition (5%), synthetic data assisted model initialization provides competitive model performance. This finding suggests the possibility of collecting significantly fewer private data, therefore reducing privacy concerns and risks associated with data leaks or misuse.

Personalized/Local Learning to Enhance Privacy The current results also imply the possibility of few-shot learning on the user side without sharing user data. Our future studies would focus on personalized/local learning utilizing synthetic data assisted model initialization to minimize privacy risks.

Improving Text Generation Could Further Enhance the Data-free Zero-Shot Learning We conjecture that improving the current text prompt would substantially boost the zero-shot learning performance. As presented in Fig 2, the current label-guided text generation may not produce qualitative texts (e.g. low perplexity) for training. This can be related to the constraints of the chosen LLM. These findings emphasize the need for exploring more powerful LLMs such as GPT-4. For example, the current LLM failed to understand the prompt to generate the spoken utterance with the intent to tell a joke, and an incorrect example output is: "Hey man, it's raining again."

Enhance Text-to-speech Quality To Improve Data-free Zero-Shot Learning Apart from the quality of the text generation, the TTS output also plays a critical role in the zero-shot learning performance with synthetic speech. The existing TTS module lacks the capability to convey a wide range of emotions and prosodic styles comparable to human speech which is essential for generating emotional speech samples. Our future work aims to explore more advanced TTS models.

8. References

- [1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [2] T. Feng, R. Hebban, N. Mehlman, X. Shi, A. Kommineni, and S. Narayanan, "A review of speech-centric trustworthy machine learning: Privacy, safety, and fairness," *APSIPA Transactions on Signal and Information Processing*, vol. 12, no. 3, 2023.
- [3] C. Pierre, A. Larcher, and D. Jouvét, "Are disentangled representations all you need to build speaker anonymization systems?" in *Proc. Interspeech 2022*, 2022, pp. 2793–2797.
- [4] P.-G. Noé, M. Mohammadamini, D. Matrouf, T. Parcollet, A. Nautsch, and J.-F. Bonastre, "Adversarial Disentanglement of Speaker Representation for Attribute-Driven Privacy Preservation," in *Proc. Interspeech 2021*, 2021, pp. 1902–1906.
- [5] T. Feng, H. Hashemi, M. Annavaram, and S. S. Narayanan, "Enhancing privacy through domain adaptive noise injection for speech emotion recognition," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 7702–7706.
- [6] N. Tomashenko, X. Wang, E. Vincent, J. Patino, B. M. L. Srivastava, P.-G. Noé, A. Nautsch, N. Evans, J. Yamagishi, B. O'Brien et al., "The voiceprivacy 2020 challenge: Results and findings," *Computer Speech & Language*, vol. 74, p. 101362, 2022.
- [7] N. Tomashenko, X. Wang, X. Miao, H. Nourtel, P. Champion, M. Todisco, E. Vincent, N. Evans, J. Yamagishi, and J. F. Bonastre, "The voiceprivacy 2022 challenge evaluation plan," *arXiv preprint arXiv:2203.12468*, 2022.
- [8] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," *arXiv preprint arXiv:1610.05492*, 2016.
- [9] T. Zhang, T. Feng, S. Alam, S. Lee, M. Zhang, S. S. Narayanan, and S. Avestimehr, "Fedaudio: A federated learning benchmark for audio tasks," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [10] D. Dimitriadis, M. H. Garcia, D. M. Diaz, A. Manoel, and R. Sim, "Flute: A scalable, extensible framework for high-performance federated learning simulations," *arXiv preprint arXiv:2203.13789*, 2022.
- [11] T. Feng and S. Narayanan, "Semi-FedSER: Semi-supervised Learning for Speech Emotion Recognition On Federated Learning using Multiview Pseudo-Labeling," in *Proc. Interspeech 2022*, 2022, pp. 5050–5054.
- [12] M. Nasr, R. Shokri, and A. Houmansadr, "Comprehensive privacy analysis of deep learning: Passive and active white-box inference attacks against centralized and federated learning," in *2019 IEEE Symposium on Security and Privacy*. IEEE, 2019, pp. 739–753.
- [13] L. Zhu, Z. Liu, and S. Han, "Deep leakage from gradients," *Advances in neural information processing systems*, vol. 32, 2019.
- [14] W.-C. Tseng, W.-T. Kao, and H. yi Lee, "Membership Inference Attacks Against Self-supervised Speech Models," in *Proc. Interspeech 2022*, 2022, pp. 5040–5044.
- [15] T. Feng, H. Hashemi, R. Hebban, M. Annavaram, and S. S. Narayanan, "Attribute inference attack of speech emotion recognition in federated learning settings," *arXiv preprint arXiv:2112.13416*, 2021.
- [16] R. He, S. Sun, X. Yu, C. Xue, W. Zhang, P. Torr, S. Bai, and X. Qi, "Is synthetic data from generative models ready for image recognition?" *arXiv preprint arXiv:2210.07574*, 2022.
- [17] Y. Wang, Y. Kordi, S. Mishra, A. Liu, N. A. Smith, D. Khashabi, and H. Hajishirzi, "Self-instruct: Aligning language model with self generated instructions," *arXiv preprint arXiv:2212.10560*, 2022.
- [18] T.-Y. Hu, M. Armandpour, A. Shrivastava, J.-H. R. Chang, H. Koppula, and O. Tuzel, "Synt++: Utilizing imperfect synthetic data to improve speech recognition," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 7682–7686.
- [19] T. Zhang, T. Feng, S. Alam, M. Zhang, S. S. Narayanan, and S. Avestimehr, "Gpt-fl: Generative pre-trained model-assisted federated learning," *arXiv preprint arXiv:2306.02210*, 2023.
- [20] P. Voigt and A. Von dem Bussche, "The eu general data protection regulation (gdpr)," *A Practical Guide, 1st Ed., Cham: Springer International Publishing*, vol. 10, no. 3152676, pp. 10–5555, 2017.
- [21] C. M. Lee and S. S. Narayanan, "Toward detecting emotions in spoken dialogs," *IEEE transactions on speech and audio processing*, vol. 13, no. 2, pp. 293–303, 2005.
- [22] G. Tur and R. De Mori, *Spoken language understanding: Systems for extracting semantic information from speech*. John Wiley & Sons, 2011.
- [23] D. Serdyuk, Y. Wang, C. Fuegen, A. Kumar, B. Liu, and Y. Bengio, "Towards end-to-end spoken language understanding," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5754–5758.
- [24] H. W. Chung, L. Hou, S. Longpre, B. Zoph, Y. Tay, W. Fedus, E. Li, X. Wang, M. Dehghani, S. Brahma et al., "Scaling instruction-finetuned language models," *arXiv preprint arXiv:2210.11416*, 2022.
- [25] J. Ao, R. Wang, L. Zhou, C. Wang, S. Ren, Y. Wu, S. Liu, T. Ko, Q. Li, Y. Zhang et al., "Speech5: Unified-modal encoder-decoder pre-training for spoken language processing," *arXiv preprint arXiv:2110.07205*, 2021.
- [26] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust dnn embeddings for speaker recognition," in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 5329–5333.
- [27] J. Kominek and A. W. Black, "The cmu arctic speech databases," in *Fifth ISCA workshop on speech synthesis*, 2004.
- [28] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao et al., "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [29] L. Pepino, P. Riera, and L. Ferrer, "Emotion Recognition from Speech Using wav2vec 2.0 Embeddings," in *Proc. Interspeech 2021*, 2021, pp. 3400–3404.
- [30] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "Lora: Low-rank adaptation of large language models," *arXiv preprint arXiv:2106.09685*, 2021.
- [31] Y. Li, A. Mehrish, R. Bhardwaj, N. Majumder, B. Cheng, S. Zhao, A. Zadeh, R. Mihalcea, and S. Poria, "Evaluating parameter-efficient transfer learning approaches on sure benchmark for speech understanding," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [32] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan, "IEMOCAP: Interactive emotional dyadic motion capture database," *Language resources and evaluation*, vol. 42, no. 4, pp. 335–359, 2008.
- [33] C. Busso, S. Parthasarathy, A. Burmania, M. AbdelWahab, N. Sadoughi, and E. M. Provost, "Msp-improv: An acted corpus of dyadic interactions to study emotion perception," *IEEE Transactions on Affective Computing*, vol. 8, no. 1, pp. 67–80, 2016.
- [34] E. Bastianelli, A. Vanzo, P. Swietojanski, and V. Rieser, "Slurp: A spoken language understanding resource package," *arXiv preprint arXiv:2011.13205*, 2020.
- [35] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. L. Scao, S. Gugger, M. Drame, Q. Lhoest, and A. M. Rush, "Transformers: State-of-the-art natural language processing," in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, Oct. 2020, pp. 38–45.
- [36] J. Li, R. Gadde, B. Ginsburg, and V. Lavrukhin, "Training neural speech recognition systems with synthetic speech augmentation," *arXiv preprint arXiv:1811.00707*, 2018.