



# MSDA: Combining Pseudo-labeling and Self-Supervision for Unsupervised Domain Adaptation in ASR

*Dimitrios Damianos<sup>1,2</sup>, Georgios Paraskevopoulos<sup>2</sup>, Alexandros Potamianos<sup>1</sup>*

<sup>1</sup>Speech and Language Processing Group, National Technical University of Athens, Greece

<sup>2</sup>Institute for Language and Speech Processing, Athena Research Center, Greece

d.damianos@athenarc.gr, g.paraskevopoulos@athenarc.gr, potam@central.ntua.gr

## Abstract

In this work, we investigate the Meta PL unsupervised domain adaptation framework for Automatic Speech Recognition (ASR). We introduce a Multi-Stage Domain Adaptation pipeline (MSDA), a sample-efficient, two-stage adaptation approach that integrates self-supervised learning with semi-supervised techniques. MSDA is designed to enhance the robustness and generalization of ASR models, making them more adaptable to diverse conditions. It is particularly effective for low-resource languages like Greek and in weakly supervised scenarios where labeled data is scarce or noisy. Through extensive experiments, we demonstrate that Meta PL can be applied effectively to ASR tasks, achieving state-of-the-art results, significantly outperforming state-of-the-art methods, and providing more robust solutions for unsupervised domain adaptation in ASR. Our ablations highlight the necessity of utilizing a cascading approach when combining self-supervision with self-training.

**Index Terms:** Unsupervised Domain Adaptation, Automatic Speech Recognition, Self-supervision, Semi-supervision, Pseudo-labeling, Low-resource Languages

## 1. Introduction

Automatic Speech Recognition (ASR) has transformed human-machine interaction, driving technologies such as virtual assistants, automated transcription services, and voice-controlled devices. However, ASR systems often face challenges in maintaining accuracy when exposed to data that differ from their training environment. Factors such as background noise, variations in recording equipment, diverse speaking styles, and regional accents can significantly degrade performance. To address these issues, domain adaptation techniques are crucial. They help close the gap between training and real-world deployment conditions, enhancing the system's robustness in diverse and unfamiliar scenarios.

Unsupervised Domain Adaptation (UDA) methods are a popular choice to mitigate this problem, as they rely on unlabeled data to achieve adaptation to the target domain, which are easily accessible [1, 2]. In the context of ASR, UDA has been employed to improve robustness across various recording conditions, such as environmental noise and reverberation. In addition, UDA techniques have been applied to cross-lingual and multilingual adaptation tasks, enhancing performance in low-resource languages [3] and different dialects [4].

A prevalent approach to Unsupervised Domain Adaptation (UDA) is teacher-student learning or pseudo-labeling [5, 6], which reframes the unsupervised task as a supervised one. In this framework, a teacher model trained on the source domain generates pseudo-labels for the target domain, which are then

used by a student model for supervised training. However, the teacher's limited exposure to the target domain often results in erroneous pseudo-labels, leading to the accumulation of errors. To address this, a Confidence Estimation Module (CEM) [7, 8] filters out unreliable pseudo-labels by estimating the confidence of each prediction, thereby improving training stability and label quality. Additionally, Noisy Student Training (NST) [9] enhances domain adaptation by training the student model on heavily augmented target data, such as time masking or speed perturbation, encouraging robustness against domain shifts and improving generalization across domains. Several works propose dynamically updating the teacher model based on the student's performance to enhance UDA. KAIZEN [10] updates the teacher every  $\Delta$  steps using the Exponential Moving Average (EMA) of the student's parameters, ensuring that the teacher remains a stable yet adaptive guide during training. Meta Pseudo Labels [11] introduces a feedback loop where the teacher is optimized based on the student's source domain performance, resulting in more accurate pseudo-labels and improved student generalization across domains. Progressive Pseudo-Labeling [12] refines pseudo-labels through an iterative process, where a student trained with hard labels from beam search decoding and a language model subsequently serves as a teacher, progressively improving Word Error Rate (WER). Similarly, kNN-CTC [13] enhances pre-trained CTC-based ASR models by retrieving fine-grained CTC pseudo-labels from a kNN datastore. Its "skip-blank" strategy effectively addresses alignment challenges in audio-transcript pairs, leading to more accurate predictions and improved UDA performance.

Another widely used approach for domain adaptation is self-supervision, which was first introduced for Natural Language Processing (NLP) tasks [14, 15]. This strategy has proven to be an effective and straightforward technique for domain adaptation. The core concept behind self-supervision involves using self-supervised loss during training to adapt large pre-trained models to new domains. Continual Pre-Training (CPT) [15] has been extensively explored in ASR, involves further pre-training of large pre-trained models on unlabeled data to enhance performance on target domain tasks. Hsu et al. [16] emphasize the importance of incorporating in-domain unlabeled data during continued pre-training (CPT) to enhance domain adaptation performance, while CASTLE [17] demonstrates that integrating CPT with a two-stage pseudo-labeling strategy can substantially reduce error rates in the target domain. Another approach is to integrate self-supervision into a multitask objective for model training, ensuring both domains contribute to the self-supervised task. In UDALM [14], BERT [18] is pre-trained on both source and target domains, using the pre-training loss on the target domain during fine-tuning to maintain the adaptation. Similarly, [19] applies self-




Figure 1: *The Multi-Stage Domain Adaptation (MSDA). At Stage 1, we use self-supervision to adapt our pre-trained model to both source  $(x_s, y_s)$  and target  $(x_t)$  domain, using the loss objective described in Eq. (1). At Stage 2, we apply an enhanced version of Meta PL to further improve the overall adaptation. In this stage, the teacher provides pseudo-labels  $y_t^*$  for student training, and itself is trained using the objective described in Eq. (2).*

supervision to ASR domain adaptation, ensuring both domains contribute to the self-supervised task to avoid mode collapse.

Our main contributions are: (a) we demonstrate that the Meta Pseudo Labels framework [11], originally designed for image recognition tasks, is a viable adaptation method for ASR. (b) We introduce Multi-Stage Domain Adaptation (MSDA), a novel, sample efficient, two-stage domain adaptation framework that integrates self-supervised techniques from [19] with semi-supervised strategies from [11], built on the Wav2Vec 2.0 architecture [20], achieving state-of-the-art adaptation results. (c) Through extensive experimentation, we demonstrate that the cascading framework is the most effective method to combine self-supervised learning with the Meta PL framework. Our results highlight that MSDA significantly enhances model performance in target domains, demonstrates enhanced sample efficiency, and achieves state-of-the-art results, particularly for low-resource languages and weakly supervised data.

## 2. Methodology

Fig. 1 illustrates our proposed approach, which builds on and extends the methodologies presented in [19] and [11]. This approach, termed Multi-Stage Domain Adaptation (MSDA), consists of two sequential adaptation stages: an initial self-supervised stage followed by a semi-supervised stage.

**Stage 1:** In this stage, we follow the method proposed in [19] to develop a teacher model trained on both domains in a self-supervised manner, using the following objective:

$$L_{M2DS2} = L_{CTC}(x_s, y_s) + \alpha L_s(x_s) + \beta L_s(x_t) \quad (1)$$

Here,  $L_{CTC}(x_s, y_s)$  is the Connectionist Temporal Classification (CTC) loss [21] applied to source domain samples  $(x_s, y_s)$ , ensuring accurate alignment between the input speech and corresponding transcriptions. The terms  $L_s(x_s)$  and  $L_s(x_t)$  represent contrastive self-supervised losses for the source and target speech data, respectively, following the approach in [20]. These losses promote the learning of robust, domain-invariant feature representations. The coefficients  $\alpha$  and  $\beta$  control the contribution of the self-supervised objectives.

**Stage 2:** In this stage, we leverage the teacher model trained in Stage 1 to generate pseudo-labels for the target domain. These pseudo-labels are then used to train a student model with the CTC objective. Simultaneously, the teacher model is refined using the following objective:

$$L = L_{feedback}(x_s, y_s) + \gamma L_{CTC}(x_s, y_s) + \delta L_d(x_t) \quad (2)$$

In this formulation,  $L_{feedback}$  represents the student’s CTC loss on source domain data, providing feedback that guides the teacher to produce higher-quality pseudo-labels. The term

$L_{CTC}$  corresponds to the teacher’s own CTC loss in the source domain, ensuring continued alignment with the labeled data.  $L_d(x_t)$  is a diversity loss, as defined in [20], which encourages output diversity to prevent overfitting. The parameters  $\gamma$  and  $\delta$  are discount coefficients that balance the contribution of each term.

To effectively integrate self-supervised learning with Meta PL’s feedback, we adopt a cascading framework. In Stage 1, we apply SpecAugment [22] to both domains’ audio inputs to train a more robust teacher. In Stage 2, however, SpecAugment is only used in student’s inputs. We avoid applying it to the teacher’s inputs in this stage, as the added noise and augmentation lead to the generation of low-quality pseudo-labels.

We use Mixed Multi-Domain Self-Supervision (M2DS2) training at the first stage because, as described in [19], multi-domain self-supervision enables domain adaptation while preventing potential mode collapse. This approach facilitates the generation of higher-quality pseudo-labels in the second stage. Furthermore, incorporating  $L_d$  into the target domain is essential to prevent mode collapse, while  $L_{CTC}$  ensures that the teacher model remains adaptable to the source domain despite adjustments made by  $L_{feedback}$ .

This two-stage framework not only leverages the strengths of both self-supervised and semi-supervised learning but also enhances domain generalization by maintaining source knowledge while adapting to the target domain. The integration of SpecAugment and diversity loss further improves model robustness, leading to a more stable and reliable adaptation process.

## 3. Experimental Setup

### 3.1. Pre-trained model:

For our base model, we utilize XLSR-53 [23], a state-of-the-art pre-trained speech model developed on the Wav2Vec 2.0 [20] architecture. XLSR-53 stands out for its extensive training on a diverse and extensive corpus, comprising 56,000 hours of speech data across 53 languages. This diverse training set equips XLSR-53 with a robust and comprehensive understanding of various linguistic and acoustic features, making it highly versatile for cross-lingual and multi-lingual applications. We employed the openly available version of XLSR-53<sup>1</sup>, which offers a pre-trained implementation ready for fine-tuning and adaptation tasks.

### 3.2. Dataset:

In our experiments we used the following datasets:

<sup>1</sup><https://huggingface.co/facebook/wav2vec2-large-xlsr-53>

Adaptation Setting	FT	CPT	M2DS2	FT-MP	M2DS2-MP	CASTLE	MSDA
LG → CV	59.57	66.43	51.31	61.76	55.21	<u>51.20</u>	<b>50.18</b>
LG → HP	62.13	67.51	<u>60.09</u>	61.01	61.98	<b>59.27</b>	60.98
CV → LG	69.55	71.12	63.40	65.89	<u>63.33</u>	64.71	<b>62.51</b>
CV → HP	70.72	73.83	68.70	66.64	<u>63.18</u>	64.51	<b>60.08</b>
HP→CV	55.90	59.68	52.95	54.50	50.83	<b>46.13</b>	<u>47.88</u>
HP→LG	48.65	52.63	46.37	47.03	<u>46.05</u>	48.06	<b>45.12</b>

Table 1: MSDA performance compared to our selected baselines, using greedy decoding, evaluated with the WER (%) metric.  $A \rightarrow B$  indicates adaptation from  $A$  (source) to  $B$  (target).

**Logotypografia:** Logotypografia (LG) [24] is one of the first corpora for Large Vocabulary Continuous Speech Recognition in Greek. The dataset contains 33,136 newscast utterances, or 72 hours of speech.

**Common Voice:** Common Voice (CV) [25] is a multilingual, crowd-sourced dataset developed by Mozilla. For our experiments, we utilize version 9.0 of CV, specifically the Greek corpus, which comprises 12 hours of training speech data.

**HParl:** HParl (HP) [19] consists of parliamentary recordings from sessions of the Hellenic Parliament, spanning the years 2018 to 2022, which contains 99 hours of training audio.

**Greek Podcast Corpus:** The Greek Podcast Corpus (GPC) [26] consists of 3,124 hours of weakly-supervised audio sourced from Greek podcasts, spanning 16 diverse categories such as *True Crime* and *Comedy*. Transcriptions were automatically generated using the WhisperX pipeline [27], resulting in pseudo-labeled audio data. For our experiments, we used the GPC-20 subset, which includes 20 hours of audio per category, focusing specifically on the *True Crime*, *Business*, *Comedy*, and *Education* categories.

### 3.3. Baselines

We evaluate the effectiveness of MSDA using the Word Error Rate (WER) metric. All models are tested on the target domain test set. Our approach is compared with the following six baselines:

**Finetuning (FT):** We perform supervised finetuning of XLSR-53 (CTC) using only the source domain data and evaluate on the target domain test set.

**Continual Pre-Training (CPT):** We perform a pretraining phase using the diversity and contrastive losses from Eq. (1) on the target domain train set to create adapted versions of XLSR. These versions are then finetuned on the source domain using the CTC objective and evaluated on the target domain.

**M2DS2:** We train XLSR-53 using the mixed objective of Eq. (1) on both domains and evaluate on the target test data.

**Finetuned-Meta PL (FT-MP):** We employ the Meta PL framework [11], using the model fine-tuned on the source domain as the teacher, and evaluate on the target domain.

**M2DS2-Meta PL:** We employ the Meta PL framework [11], using the M2DS2-trained model on both domains as the teacher, and evaluate on the target domain.

**CASTLE:** Finally, we utilize the CASTLE [17] framework on our chosen datasets, using the publicly available code<sup>2</sup> provided by the authors, and evaluate on the target test set.

### 3.4. Experimental Settings

For the experiments involving the LG, CV and HP datasets, we apply Eq. (1) and (2) with  $\alpha = 0.01$ ,  $\beta = 0.02$ , and  $\gamma = \delta = 0.0001$ . CASTLE [17] consists of an initial CPT

phase followed by two pseudo-labeling phases, where a 4-gram language model is used to generate higher-quality pseudo-labels via beam search. Since MSDA focuses on the acoustic model, we adapted CASTLE for a fair comparison by directly generating pseudo-labels from the acoustic model using greedy decoding, without any language model assistance. For our weakly-supervised data, the parameters in Eq. (1) and (2) are set to  $\alpha = 0.01$ ,  $\beta = 0.01$ , and  $\gamma = \delta = 0.001$ . For Stage 1, we followed the training setup described in [19]. For Stage 2, we used the AdamW optimizer [28] with a learning rate of  $3 \times 10^{-4}$ , no learning rate scheduler, and a maximum of 30 epochs. All experiments were conducted on a single NVIDIA RTX 3090 GPU (24GB) using mixed precision training.

## 4. Results

In Table 1, we compare the performance of MSDA with selected baselines across six adaptation scenarios involving the LG, CV and HP datasets. Meta PL consistently exhibits strong adaptation capabilities, with FT-MP and M2DS2-MP achieving notable WER reductions of 1–4% and 1–2%, respectively, compared to the FT and M2DS2 baselines. These results highlight the effectiveness of Meta PL in ASR adaptation. MSDA consistently outperforms all baselines, delivering significant WER improvements. In particular, MSDA achieves WER reductions of 2–10% compared to FT and outperforms M2DS2 by 1–10%. Notably, MSDA demonstrates robustness in challenging scenarios where M2DS2 struggles, such as the  $HP \rightarrow LG$  adaptation. Moreover, MSDA surpasses the state-of-the-art CASTLE framework in four out of six adaptation scenarios, achieving absolute WER improvements ranging from 1 to 4%.

Adaptation Setting	FT	M2DS2	FT-MP	M2DS2 MP	MSDA
TC→Bus	48.8	45.32	43.47	42.35	<b>37.98</b>
TC→Ed	62.7	57.16	58.26	56.27	<b>50.3</b>
TC→Com	61.9	68.77	57.56	57.1	<b>49.21</b>
Ed→TC	37.28	43.23	<b>33.3</b>	42.07	39.89
Ed→Bus	50.93	56.01	48.25	55.17	<b>47.03</b>
Ed→Com	52.54	71.14	<b>50.7</b>	58.06	56.16

Table 2: MSDA performance on weakly-supervised data, using WER. We used GPC’s *TrueCrime* (TC), *Education* (Ed), *Business* (Bus) and *Comedy* (Com) categories as different domains.

A similar trend is observed in Table 2, where we assess MSDA’s performance on weakly-supervised data across six adaptation scenarios, comparing it with the FT, M2DS2, FT-MP, and M2DS2-MP baselines. Once again, Meta PL shows consistent adaptation, with FT-MP and M2DS2-MP achieving WER reductions of 1–4% and 1–2% compared to FT and M2DS2. MSDA continues to outperform our baselines in most scenarios, achieving significant WER improvements. Specifically, MSDA reduces WER by 3–12% compared to FT, and outperforms M2DS2 by 7–18%, even in cases where M2DS2 fails to

<sup>2</sup><https://github.com/zhu-han/CASTLE/tree/CASTLE-base>




Figure 2: WER comparison between different setting of  $\gamma$  (blue) and  $\delta$  (red), FT (yellow) and M2DS2 (green) baselines, evaluated in the  $TC \rightarrow Ed$  adaptation setting.

adapt, such as  $Ed \rightarrow TC$ .

Fig. 2 illustrates the effect of  $\gamma$  and  $\delta$  on MSDA performance for *TrueCrime-Education* adaptation. In two separate experiments, one coefficient is fixed at 0.001 while the other varies from 1 to  $10^{-5}$ . MSDA struggles when  $\gamma, \delta > 10^{-3}$  due to source overfitting and strong  $L_d$  influence, which hinder adaptation. For  $\gamma, \delta < 10^{-4}$ , performance drops from mode collapse as  $L_{CTC}$  and  $L_d$  lose effectiveness. The best results are achieved when  $\gamma, \delta$  are between  $10^{-4}$  and  $10^{-3}$ .

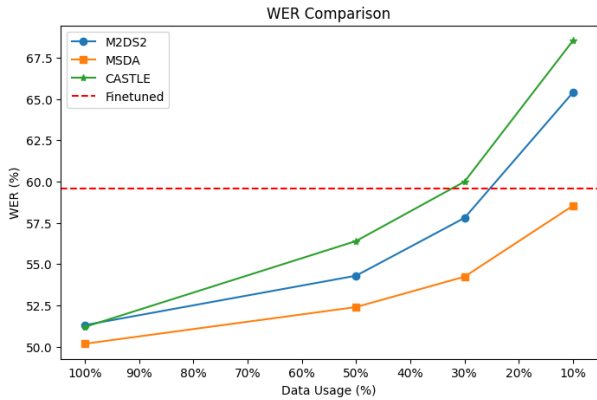


Figure 3: Performance comparison between M2DS2, CASTLE and MSDA for the  $LG \rightarrow CV$  setting, when reducing the amount of available target samples to 50%, 25%, and 10% of the original dataset. FT performance is indicated with the red line.

Finally, Fig. 3 demonstrates the sample efficiency of MSDA in the  $LG \rightarrow CV$  scenario, compared with M2DS2 and CASTLE, when using 50%, 25%, and 10% of the available target data (corresponding to 6, 3, and 1.2 hours of audio, respectively). The results indicate that MSDA consistently outperforms M2DS2 across all cases, achieving noticeable adaptation even with just 10% of the available target data. This highlights MSDA’s ability to facilitate adaptation even in scenarios where data collection is challenging or impractical.

## 5. Meta PL & Self-Supervision

We explored different ways to integrate self-supervised learning with the Meta PL framework in a non-cascading manner. We experimented with combining self-supervised objectives with student feedback to improve teacher performance, as described

in Eq. (3).

$$L_T = L_{feedback}(x_s, y_s) + L_s(x_s) + L_s(x_t) \quad (3)$$

In a second approach, we applied self-supervision solely to the target domain for the student model during the pseudo-labeling phase, using Eq. (4).

$$L_S = L_{CTC}(x_t, y_t^*) + L_s(x_t) \quad (4)$$

Despite these efforts, both approaches failed to achieve successful adaptation, as shown in Table 3, where we experimented with both FT and M2DS2-trained teacher models on the GPC  $TC \rightarrow Ed$  source-target adaptation setting.

Method	FT		M2DS2	
	Teacher	Student	Teacher	Student
MSDA	62.7	<b>58.18</b>	57.16	<b>50.3</b>
teacher feedback+self super	62.7	97	57.16	95.83
student CTC+self super	62.7	72	57.16	67.1

Table 3: Alternatives explored during integration on the  $TC \rightarrow Ed$  source-target adaptation setting. The “Teacher” column contains teacher’s original WER on target domain, while the “Student” column contains student’s WER after adaptation

We found that applying Eq. (3) to the teacher model led to the generation of low-quality pseudo-labels, which impaired the student model’s ability to adapt effectively to the target domain. Likewise, using Eq. (4) on the student model produced misleading feedback for the teacher, ultimately reducing the overall quality of the pseudo-labels.

## 6. Conclusions & Future Work

In this work, we explored the application of the Meta PL framework [11] in ASR tasks and investigated its integration with self-supervision objectives. We found that Meta PL is an effective adaptation method, providing a straightforward and easily implementable solution. However, our proposed approach, Multi-Stage Domain Adaptation (MSDA), significantly outperforms Meta PL, M2DS2, and CASTLE, as highlighted in Tables 1 and 2. Furthermore, MSDA is a sample-efficient framework, as shown in Fig. 3, making it particularly suitable for scenarios where data collection is limited or infeasible. We also explored alternative ways to integrate self-supervision within the Meta PL framework and concluded that cascading methods show promising results. In future work, we plan to investigate the effectiveness of MSDA in additional adaptation settings, such as accent adaptation and cross-lingual adaptation. Furthermore, we aim to evaluate its performance on a broader range of tasks, including image recognition and natural language processing (NLP), to further assess its versatility and robustness across diverse domains.

## 7. Acknowledgements

This work is co-funded by the European Union’s Digital Europe Programme and by the European Union through the National Strategic Reference Framework (NSRF) 2021–2027, under the Operational Programme “Competitiveness” Grant No 101083565 and Grant No 6002637 respectively, for the project SmartAttica-AtHeNAI DIH (the Attica region - Greek Innovation hub for Artificial Intelligence in Energy and Environment, Supply chain and mobility, Culture and Tourism).

## 8. References

- [1] M. Long et al., “Learning transferable features with deep adaptation networks,” in *Proc. ICML*. PMLR, 2015, pp. 97–105.
- [2] Y. Ganin et al., “Domain-adversarial training of neural networks,” *Journal of machine learning research*, vol. 17, no. 59, pp. 1–35, 2016.
- [3] C. S. Anoop et al., “Unsupervised domain adaptation schemes for building asr in low-resource languages,” in *Proc. ASRU*. IEEE, 2021, pp. 342–349.
- [4] T. Asami et al., “Domain adaptation of dnn acoustic models using knowledge distillation,” in *Proc. ICASSP*. IEEE, 2017, pp. 5185–5189.
- [5] H. Scudder, “Probability of error of some adaptive pattern-recognition machines,” *IEEE Transactions on Information Theory*, vol. 11, no. 3, pp. 363–371, 1965.
- [6] E. Riloff and J. Wiebe, “Learning extraction patterns for subjective expressions,” in *Proc. EMNLP*, 2003, pp. 105–112.
- [7] D. Hwang et al., “Large-scale asr domain adaptation using self- and semi-supervised learning,” in *Proc. ICASSP*. IEEE, 2022, pp. 6627–6631.
- [8] D. Hwang, K. C. Sim, Z. Huo, and T. Strohmaier, “Pseudo label is better than human label,” in *Proc. Interspeech*. ISCA, 2022, pp. 1421–1425.
- [9] D. S. Park et al., “Improved noisy student training for automatic speech recognition,” in *Proc. Interspeech*. ISCA, 2020, pp. 2817–2821.
- [10] V. Manohar et al., “Kaizen: Continuously improving teacher using exponential moving average for semi-supervised speech recognition,” in *Proc. ASRU*. IEEE, 2021, pp. 518–525.
- [11] H. Pham et al., “Meta pseudo labels,” in *Proc. CVPR*. IEEE/CVF, 2021, pp. 11 557–11 568.
- [12] R. Ahmad, M. U. Farooq, and T. Hain, “Progressive unsupervised domain adaptation for asr using ensemble models and multi-stage training,” in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 11 466–11 470.
- [13] J. Zhou, S. Zhao, Y. Liu, W. Zeng, Y. Chen, and Y. Qin, “knn-ctc: Enhancing asr via retrieval of ctc pseudo labels,” in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 11 006–11 010.
- [14] C. Karouzos et al., “Udalm: Unsupervised domain adaptation through language modeling,” in *Proc. NAACL: HLT*, 2021, pp. 2579–2590.
- [15] S. Gururangan et al., “Don’t stop pretraining: Adapt language models to domains and tasks,” in *Proc. 58th Annual Meeting of ACL*, 2020, pp. 8342–8360.
- [16] W. Hsu et al., “Robust wav2vec 2.0: Analyzing domain shift in self-supervised pre-training,” in *Proc. Interspeech*. ISCA, 2021, pp. 721–725.
- [17] H. Zhu et al., “Boosting cross-domain speech recognition with self-supervision,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 32, pp. 471–485, 2024.
- [18] J. D. Kenton et al., “Bert: Pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of NAACL-HLT*, vol. 1, no. 2, 2019.
- [19] G. Paraskevopoulos et al., “Sample-efficient unsupervised domain adaptation of speech recognition systems: A case study for modern greek,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 32, pp. 286–299, 2024.
- [20] A. Baeviski, Y. Zhou, A. Mohamed, and M. Auli, “wav2vec 2.0: A framework for self-supervised learning of speech representations,” in *Proc. NeurIPS*, vol. 33, 2020, pp. 12 449–12 460.
- [21] A. Graves and N. Jaitly, “Towards end-to-end speech recognition with recurrent neural networks,” in *Proc. ICML*. PMLR, 2014, pp. 1764–1772.
- [22] D. Park et al., “SpecAugment: A simple data augmentation method for automatic speech recognition,” in *Proc. Interspeech*. ISCA, 2019, pp. 2613–2617.
- [23] A. Conneau et al., “Unsupervised cross-lingual representation learning for speech recognition,” in *Proc. Interspeech*. ISCA, 2021, pp. 2426–2430.
- [24] V. Digilakis et al., “Large vocabulary continuous speech recognition in greek: corpus and an automatic dictation system,” in *Proc. Eurospeech*, 2003, pp. 1565–1568.
- [25] R. Ardila et al., “Common voice: A massively-multilingual speech corpus,” in *Proc. LREC*, 2020, pp. 4218–4222.
- [26] G. Paraskevopoulos et al., “The greek podcast corpus: Competitive speech models for low-resourced languages with weakly supervised data,” in *Proc. Interspeech*. ISCA, 2024, pp. 3969–3973.
- [27] M. Bain et al., “Whisperx: Time-accurate speech transcription of long-form audio,” in *Proc. Interspeech*. ISCA, 2023, pp. 4489–4493.
- [28] I. Loshchilov and F. Hutter, “Decoupled weight decay regularization,” *arXiv e-prints*, pp. arXiv–1711, 2017.