



SGEM: Test-Time Adaptation for Automatic Speech Recognition via Sequential-Level Generalized Entropy Minimization

Changhun Kim¹ Joonhyung Park¹ Hajin Shim¹ Eunho Yang^{1,2}

¹Korea Advanced Institute of Science and Technology (KAIST), South Korea

²AITRICS, South Korea

{changhun.kim, deepjoon, shimazing, eunhoy}@kaist.ac.kr

Abstract

Automatic speech recognition (ASR) models are frequently exposed to data distribution shifts in many real-world scenarios, leading to erroneous predictions. To tackle this issue, an existing test-time adaptation (TTA) method has recently been proposed to adapt the pre-trained ASR model on unlabeled test instances without source data. Despite decent performance gain, this work relies solely on naive greedy decoding and performs adaptation across timesteps at a frame level, which may not be optimal given the sequential nature of the model output. Motivated by this, we propose a novel TTA framework, dubbed SGEM, for general ASR models. To treat the sequential output, SGEM first exploits beam search to explore candidate output logits and selects the most plausible one. Then, it utilizes generalized entropy minimization and negative sampling as unsupervised objectives to adapt the model. SGEM achieves state-of-the-art performance for three mainstream ASR models under various domain shifts.¹

Index Terms: automatic speech recognition, test-time adaptation, beam search, entropy minimization, negative sampling

1. Introduction

While deep neural networks have achieved remarkable progress in a broad range of areas, such as computer vision [1, 2], natural language processing [3, 4], and speech processing [5, 6], these models are known to be susceptible to data distribution shifts [7, 8]. This so-called *domain shift problem* readily occurs in recent automatic speech recognition (ASR) models, which imposes challenges in deploying ASR models to real-world applications. For example, utterances of unseen speakers/words not exposed during training or utterances with accidental background noise can be given at test time.

To tackle this domain shift problem for ASR models, many prior works have been suggested, including data augmentation [9], feature alignment [10], domain adversarial learning [11, 12], and knowledge distillation [13]. These works mainly mitigate the domain shift problem under the unsupervised domain adaptation (UDA) setting, where the source models are adapted to unlabeled target domains. However, the UDA setting has several impractical assumptions in real-world scenarios. First, it assumes that the source data is accessible, which might be unavailable due to privacy/storage issues. Second, a pile of target data has to be collected *in advance*. This is also unrealistic as it requires substantial resources. Even worse, it restricts the generalization capacity of the model only to the pre-collected target data, although the target distribution can change arbitrarily at test time. Meanwhile, several speaker adaptation

methods [14, 15, 16] have demonstrated satisfactory adaptation performance on variation in speakers during inference. However, their restricted focus on speaker changes and reliance on prior knowledge of test-time speakers impose limitations in effectively addressing arbitrary domain shifts during test time.

Inspired by these limitations, a method called SUTA [17] was first proposed to address an arbitrary domain shift problem under a more realistic setting - test-time adaptation (TTA) - for ASR models as in other domains [18, 19, 20, 21, 22, 23, 24, 25, 26, 27]. Given an off-the-shelf ASR model pre-trained on the source domain, TTA methods aim to adapt the model on-the-fly using unlabeled instances from the target domain in test time without access to source data. Inheriting the ideas of the TTA approaches in the computer vision domain [19, 21, 22], SUTA shows decent performance in a single-utterance TTA setting for the CTC-based ASR model [6].

However, directly adopting this approach to advanced ASR models [28, 29] could not be optimal as SUTA was developed with the CTC-based models [30] in mind. This is because, unlike the CTC-based models, which generate each output token independently in a greedy manner, advanced ASR models are designed to work in an autoregressive manner or typically utilize beam search decoding with an external language model during the test phase. This indicates that output logits acquired by greedy decoding may not adequately capture the output distribution, and naively adapting with these logits at a frame level as in [17] can cause undesirable behavior at the sequential level for general ASR models.

In this paper, we propose a Sequential-Level Generalized Entropy Minimization (SGEM) framework towards an effective TTA for general ASR models. To this end, SGEM first explores candidate output logits and selects the most plausible one using beam search to leverage the sequential nature of the output. Then, SGEM leverages generalized entropy minimization loss and negative sampling loss as auxiliary unsupervised objectives to adapt the model at a sequential level. We validate SGEM for three representative ASR models on various datasets with different distribution shifts, including unseen speakers/words and severe background noise, and demonstrate that SGEM achieves state-of-the-art performance in most settings. To the best of our knowledge, this is the first work suggesting the TTA method for general ASR models.

2. Proposed Method: SGEM

This section introduces SGEM, an effective TTA framework for general ASR models. To this end, we describe the test-time adaptation setup for ASR models in Section 2.1 and illustrate the core strategies of SGEM from Section 2.2 to Section 2.5. Figure 1 depicts the overall pipeline of the proposed method.

¹Our code is available at <https://github.com/drump/SGEM>.

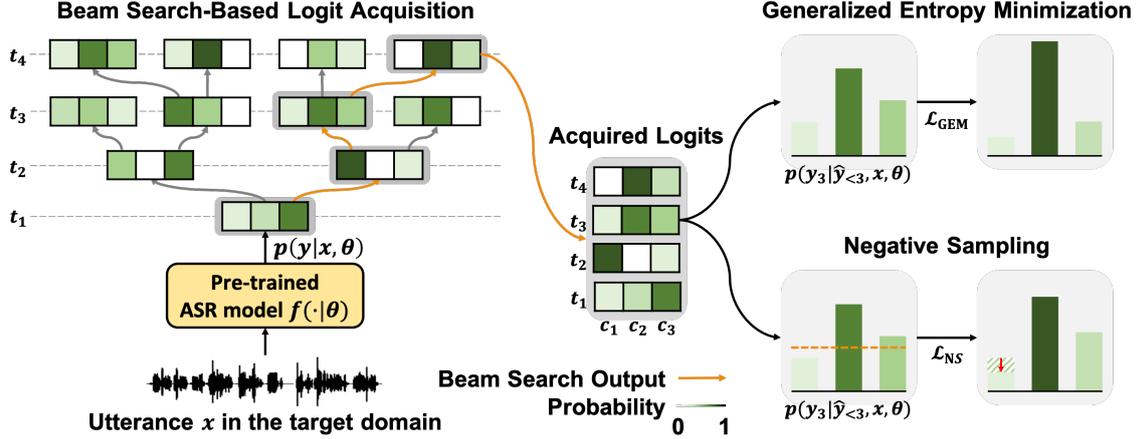


Figure 1: The overall pipeline of the SGEM framework. Given a single utterance x in the target domain, we first acquire logits based on the most plausible beam search output. Then, we utilize generalized entropy minimization and negative sampling objectives with x itself to adapt the ASR model $f(\cdot|\theta)$ pre-trained on the source domain.

2.1. Test-Time Adaptation Setup for ASR Models

We briefly explain general ASR models before formulating test-time adaptation (TTA). Let $f(\cdot|\theta)$ be an ASR model trained on the labeled source domain $\mathcal{D}_s = \{(x_i^s, y_i^s)\}_i$ in pairs of speech and text, which takes a raw waveform x and returns logits $f(x|\theta) \in \mathbb{R}^{L \times C}$ for each timestep. Here, L is the number of timesteps, C is the number of vocabulary classes, and θ is the model parameter. It models the log joint probability $\log p(\mathbf{y}|x, \theta)$ of a candidate transcript $\mathbf{y} = (y_1, \dots, y_L)$, in an *autoregressive manner* as follows:

$$\begin{aligned} \log p(\mathbf{y}|x, \theta) &:= \log p_{\text{AM}}(\mathbf{y}|x, \theta) + \lambda_{\text{LM}} \log p_{\text{LM}}(\mathbf{y}) + Z \\ &= \sum_{i=1}^L \log p_{\text{AM}}(y_i|y_{<i}, x, \theta) + \lambda_{\text{LM}} \log p_{\text{LM}}(y_i|y_{<i}) + Z, \end{aligned}$$

where $y_i \in \{1, \dots, C\}$, $p_{\text{AM}}(\mathbf{y}|x, \theta)$ is the joint probability given by the model output $f(x|\theta)$, $p_{\text{LM}}(\mathbf{y})$ is the joint probability of an autoregressive language model (LM), λ_{LM} is a hyperparameter to control the effect of the LM, and Z is a normalizing constant. LM aims to boost the ASR model to generate more faithful sentences. ASR decoding strategies approximate the optimal solution $\mathbf{y}^* = \arg \max_{\mathbf{y}} \log p(\mathbf{y}|x, \theta)$.

TTA methods for an ASR model $f(\cdot|\theta)$ aim to adapt the model to the unlabeled target speech domain $\mathcal{D}_t = \{x_i^t\}_i$ without access to \mathcal{D}_s . Specifically, we consider a *single-utterance TTA setting* where we fine-tune the ASR model $f(\cdot|\theta)$ for each utterance $x_i^t \in \mathcal{D}_t$ to get more precise output logits $\log p(\mathbf{y}|x_i^t, \theta)$ with unsupervised objectives using only x_i^t itself. This single-utterance TTA setting is considerably pragmatic regarding low latency without presuming that the test instances are independent and identically distributed [17, 24, 27].

2.2. Beam Search-Based Logit Acquisition

An existing TTA method for ASR models [17] exploits the greedy decoding strategy without leveraging the external LM (i.e., $\lambda_{\text{LM}} = 0$) to get output logits for all timesteps. Also, it utilizes them with unsupervised objectives such as entropy minimization [31] and minimum class confusion [32] as if they are appropriate logits to adapt the ASR model. However, naively using greedy decoding is proven defective [33] and can mislead the model to be adapted on the wrong labels. Furthermore, this frame-level greedy adaptation might be sub-optimal on the se-

quential output since it only considers the joint probability of a sequence myopically over timesteps.

To this end, we exploit a novel logit acquisition strategy based on more systematic beam search decoding. Given a beam width B , we find the most plausible output sequence $\hat{\mathbf{y}} = (\hat{y}_1, \dots, \hat{y}_L)$, which approximates \mathbf{y}^* using beam search [33]. Note that we do not hold the logits of beam candidates in this step to reduce memory consumption. Instead, the estimated sequence $\hat{\mathbf{y}}$ is passed to the model again to acquire the i -th logit $\mathbf{o}_i = (o_{i1}, \dots, o_{iC}) \in \mathbb{R}^C$ for all $i \in \{1, \dots, L\}$ where $o_{ij} = \log p(y_i = j|\hat{y}_{<i}, x, \theta)$. Our intuition behind the beam search-based logit acquisition is that *considering the logits obtained from beam search for adaptation naturally aligned with how ASR models decode sentences, i.e., our approach can tailor the model to adapt toward the actual yet accurate sentences generated by the ASR models.*

2.3. Generalized Entropy Minimization

While entropy minimization achieves decent performance on domain adaptation tasks [19, 21, 22, 34] by reducing the uncertainty of predictions, and extracting domain-invariant features on the target domain, we can further improve this objective by adopting its generalized version, Rényi entropy, and searching for more effective hyperparameter. For a discrete random variable X , which takes values in $\{1, \dots, C\}$, Rényi entropy $H_\alpha(X)$ of order α with $\alpha \in (0, 1) \cup (1, \infty)$ is defined as $H_\alpha(X) = \frac{1}{1-\alpha} \log \left(\sum_{j=1}^C \mathbb{P}(X=j)^\alpha \right)$. When $\alpha \rightarrow 1$ and $\alpha \rightarrow \infty$, $H_\alpha(X)$ becomes Shannon entropy and cross-entropy with a pseudo-label $\arg \max_j \mathbb{P}(X=j)$, respectively. For a single-utterance TTA setting, we hypothesize that there exists an optimal $\alpha \in (1, \infty)$ and define the generalized entropy loss as follows:

$$\mathcal{L}_{\text{GEM}} = \frac{1}{L} \sum_{i=1}^L \frac{1}{1-\alpha} \log \left(\sum_{j=1}^C p_{ij}^\alpha \right), \quad (1)$$

where $p_{ij} = \frac{\exp(o_{ij}/T)}{\sum_{j'=1}^C \exp(o_{ij'}/T)}$ and T is a temperature hyperparameter for preventing *vanishing gradient*. As the blank token dominates all timesteps, we ignore timesteps with the highest probability of blank token among all classes to alleviate the class imbalance problem as in [17].

Table 1: Comparison of TTA performance, measured as word error rate (%), for three mainstream ASR models across 12 datasets with various types of domain shifts. The results are obtained using greedy decoding for inference.

	Dataset	CH	TD	CV	VA	AC	AA	BA	CM	MU	NB	SD	TP	Avg.
CTC-based model	Unadapted	31.2	13.2	36.9	14.5	28.1	40.9	66.6	49.8	50.4	119.2	19.2	26.2	41.4
	SUTA	25.0	12.0	31.4	11.8	17.7	31.3	55.2	39.4	39.7	113.0	15.0	17.8	34.1
	SGEM	24.7	12.0	31.1	11.6	17.3	30.7	53.1	38.5	38.6	110.5	14.8	17.5	33.4
Conformer	Unadapted	28.7	15.1	36.8	17.4	18.8	44.8	74.3	45.7	56.0	122.1	20.8	36.9	43.1
	SUTA	25.2	13.4	32.4	14.7	14.5	39.8	73.3	38.4	48.7	125.5	16.4	28.8	39.3
	SGEM	24.5	13.3	31.6	14.6	14.4	38.5	70.4	38.7	48.5	120.9	16.8	28.9	38.4
Transducer	Unadapted	11.8	7.2	12.9	6.5	14.1	20.4	31.0	29.7	31.3	74.6	12.7	16.2	22.4
	SUTA	10.3	6.8	12.1	5.5	12.0	18.5	28.3	26.7	28.7	74.6	11.7	14.7	20.8
	SGEM	9.9	6.6	12.0	5.2	11.6	18.0	27.5	26.0	28.0	76.5	11.5	14.3	20.6

2.4. Negative Sampling

We further exploit negative sampling loss, originally adopted for semi-supervised learning (SSL) in [35]. Negative sampling loss penalizes the probabilities of low-confident classes, and Chen et al. [35] have shown that adding it can further boost the performance of existing SSL algorithms. It can be derived from the standard cross-entropy loss as follows. Given L labeled samples $\{(x_i, y_i)\}_{i=1}^L$, the standard cross-entropy loss is defined as $\mathcal{L}_{CE} = -\frac{1}{L} \sum_{i=1}^L \sum_{j=1}^C \mathbb{1}_{[j=y_i]} \log p_{ij}$. Note that $\sum_{j=1}^C \mathbb{1}_{[j=y_i]} \log p_{ij} = \log(1 - \sum_{j \neq y_i} p_{ij})$. Since we do not know the ground truth label y_i for each x_i in the unlabeled target domain, we approximate \mathcal{L}_{CE} with the negative sampling loss \mathcal{L}_{NS} defined as follows:

$$\mathcal{L}_{NS} = -\frac{1}{L} \sum_{i=1}^L \log \left(1 - \sum_{j=1}^C \mathbb{1}_{[p'_{ij} < \tau]} p_{ij} \right), \quad (2)$$

where $p_{ij} = \frac{\exp(o_{ij}/T)}{\sum_{j'=1}^C \exp(o_{ij'}/T)}$, $p'_{ij} = \frac{\exp(o_{ij})}{\sum_{j'=1}^C \exp(o_{ij'})}$ with a temperature hyperparameter T for avoiding vanishing gradient, and $\mathbb{1}$ is an indicator function. j -th class of x_i is considered as a negative class when the probability p'_{ij} is less than a threshold τ . Without modification, Equation (2) can be interpreted in a single-utterance TTA setting as penalizing probabilities of negative classes at every timestep for a sequential output of length L .

2.5. Overall Framework

Our final unsupervised objective is the weighted sum of the generalized entropy loss and the negative sampling loss as follows:

$$\mathcal{L} = \mathcal{L}_{GEM} + \lambda_{NS} \mathcal{L}_{NS}, \quad (3)$$

where λ_{NS} is negative sampling weight for balancing two losses. For each utterance, we adapt the model for N iterations in an *episodic manner* where we newly reset the model to the pre-trained one to preserve the knowledge from the source domain.

3. Experiments

3.1. Experimental Setup

Source ASR Models To verify the efficacy of SGEM, we evaluate it on three mainstream ASR architectures: the CTC-based model [30], Conformer [36], and Transducer [37]. More specifically, for the CTC-based model, we use wav2vec 2.0 [6]² trained on the LibriSpeech dataset [38]. For Conformer, we exploit Conformer-CTC [36]³ trained on the LibriSpeech

dataset. For Transducer, we adopt Conformer-Transducer [28]⁴ trained on a composite NeMO ASRSET dataset, including the LibriSpeech dataset. We utilize the external 4-gram language model⁵ for the CTC-based model and Conformer.

Datasets We assess the performance of SGEM on multiple datasets under various domain shift settings. To test SGEM under unseen speakers/words, we use the test set of four datasets: CHiME-3 (CH) [39], TED-LIUM 2 (TD) [40], Common Voice (CV) [41], and Valentini (VA) [42]. In addition, we validate SGEM under accident background noise by injecting the following eight types of noises to each utterance of in-domain LibriSpeech test-other dataset [38]: air conditioner (AC), airport announcement (AA), babble (BA), copy machine (CM), munching (MU), neighbors (NB), shutting door (SD), and typing (TP) with SNR = 10dB. For each type of noise, we randomly select one noise sample from the MS-SNSD noise test set [43]. We also evaluate SGEM on L2-Arctic [44], non-native English speech corpora, to verify SGEM under extreme pronunciation/accent shifts. Specifically, we randomly select one speaker for each first language.

Implementation Details Since the TTA setting has no validation set, we optimize hyperparameters on the CH dataset for each model and apply them to the other datasets. The best settings are as follows. For all models, we use AdamW optimizer [45] and cosine annealing learning rate scheduler with η_i and η_f for initial and final learning rates, respectively, and set $(N, T, \tau) = (10, 2.5, 0.4/C)$ with vocabulary size C . We only train feature extractors for the CTC-based model and encoders for the others. Furthermore, we set $(\eta_i, \eta_f, B, \lambda_{LM}, \alpha, \lambda_{NS}) = (4 \cdot 10^{-5}, 2 \cdot 10^{-5}, 5, 0.3, 1.5, 1)$ for the CTC-based model, $(4 \cdot 10^{-5}, 2 \cdot 10^{-5}, 5, 0.3, 1.25, 2)$ for Conformer, and $(4 \cdot 10^{-6}, 2 \cdot 10^{-6}, 3, 0, 1.25, 0.5)$ for Transducer. All experiments are conducted on Nvidia TITAN Xp and GeForce RTX 3090. Adaptation takes about 0.771 seconds for a 1-second utterance averaged over three models.

3.2. Main Results

We compare the TTA performance of three mainstream ASR models, including the CTC-based model, Conformer, and Transducer, across 12 datasets with various domain shifts. Table 1 presents the word error rate (WER) of ASR model outputs generated by the greedy search decoding method, following the evaluation protocol used in the previous study [17]. Additionally, Table 2 showcases the TTA performance for the CTC-based models using beam search decoding with external LM. For both decoding methods, the ASR models with SGEM

²<https://huggingface.co/facebook/wav2vec2-base-960h>

³https://catalog.ngc.nvidia.com/orgs/nvidia/teams/nemo/models/stt.en.conformer_ctc_small_ls

⁴https://catalog.ngc.nvidia.com/orgs/nvidia/teams/nemo/models/stt.en.conformer_transducer_small

⁵<https://huggingface.co/patrickvonplaten/wav2vec2-base-100h-with-lm>

Table 2: Comparison of TTA performance, measured as word error rate (%), for the CTC-based model across 12 datasets with various domain shifts. The results are obtained using beam search decoding with an external language model for inference.

Dataset	CH	TD	CV	VA	AC	AA	BA	CM	MU	NB	SD	TP	Avg.
Unadapted	29.5	12.2	36.9	13.0	26.1	38.6	58.9	48.9	49.0	91.6	17.4	23.7	37.2
SUTA	24.1	11.6	31.5	11.4	16.8	30.3	53.2	38.1	38.6	107.9	14.1	16.9	32.9
SGEM	24.1	11.7	31.1	11.1	16.5	29.8	51.6	37.7	37.7	106.7	14.0	16.9	32.4

consistently enhance the recognition accuracy of target utterances with an average word error rate reduction of 15.6%, except for two cases on NB, where the performance without adaptation is the best when using beam search decoding. Furthermore, SGEM outperforms SUTA in terms of average WER across all 12 datasets for each of the three model architectures (**CTC-based model**: (greedy) 34.1% \rightarrow **33.4%**, (beam search) 32.9% \rightarrow **32.4%** / **Conformer**: 39.3% \rightarrow **38.4%** / **Transducer**: 20.8% \rightarrow **20.6%**). This indicates the superiority of our unsupervised objectives as well as the logit acquisition method for adapting sequential language outputs regardless of the decoding strategy.

3.3. Non-Native English Speech Corpora

To show the usability of SGEM at various domain shifts, we further analyze SGEM on six different non-native English speech corpora. The result is summarized in Table 3. As shown in Table 3, SGEM achieves the best results for all corpora, outperforming the baseline. This implies the adaptability of SGEM under extreme pronunciation/accent shifts, demonstrating its versatility in practical situations with severe speaker shifts, such as globally used online ASR systems.

Table 3: Comparison of TTA performance, measured as word error rate (%), for the CTC-based model on six non-native English speech corpora.

Setting	Unadapted	SUTA	SGEM
Arabic	32.5	27.1	26.5
Mandarin	28.5	23.3	23.1
Hindi	15.7	12.5	12.3
Korean	23.3	19.7	19.5
Spanish	35.7	29.8	29.3
Vietnamese	18.5	15.7	15.4
Average	25.7	21.4	21.0

3.4. Data Deficient Condition

It is commonly known that TTA methods fail under the data deficient condition where the number of test instances is limited [27, 46]. This still holds in the single-utterance TTA setting for ASR models, where the length of utterance is short, so the number of output tokens is insufficient. To validate SGEM on this harsh condition, we split the CH dataset according to utterance length and evaluate SGEM with the CTC-based model on each split. As shown in Figure 2, SGEM performs best in every length interval. In addition, it is worth noting that SGEM significantly outperforms the baseline for extremely short utterances of less than 2 seconds, showing the superiority of our method in real-world scenarios where short utterances are prevalent and negligible latency is required.

3.5. Ablation Study

To validate the core components of SGEM: beam search-based logit acquisition (**BS**, Section 2.2), generalized entropy minimization (**GEM**, Section 2.3), and negative sampling (**NS**, Section 2.4), we conduct an ablation study for three mainstream ASR models on CH dataset. As shown in Table 4, both general-

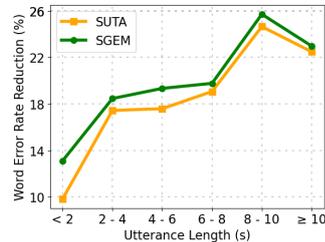


Figure 2: Comparison of TTA performance, measured as relative word error rate reduction (%), on different utterance lengths of CH dataset for the CTC-based model.

ized entropy minimization and negative sampling alone achieve remarkable performance gains for every model, indicating the efficacy of each component. Meanwhile, substituting greedy search for beam search even with small beam width (for all models) and without external LM (for Transducer) consistently boosts the performance in all cases, showing the effectiveness of beam search-based logit acquisition and implying that additional performance improvement can be expected using larger beam sizes or language model if resources are allowed.

Table 4: Ablation study of SGEM.

BS	GEM	NS	CTC	Conformer	Transducer
x	x	x	31.2	28.7	11.8
x	✓	x	24.9	24.7	10.0
x	x	✓	25.2	25.0	10.1
x	✓	✓	24.8	24.7	10.0
✓	x	✓	24.8	24.7	10.1
✓	✓	✓	24.7	24.5	9.9

4. Conclusion

We have suggested SGEM, an effective single-utterance TTA framework for general ASR models. SGEM exploits beam search-based logit acquisition and utilizes generalized entropy minimization and negative sampling objectives to adapt the model at a sequential level. SGEM has achieved state-of-the-art performance for the three mainstream ASR models across 12 datasets with various domain shifts, including utterances of unseen speakers/words during training and utterances with severe background noise. We have also verified SGEM under harsh conditions, such as non-native English utterances with severe pronunciation/accent shifts and the data deficient condition, and the efficacy of each component of SGEM through an ablation study. SGEM sheds light on the careful design of speech-specific components when devising test-time adaptation methods for ASR models.

5. Acknowledgement

This work was supported by National Research Foundation of Korea (NRF) grants (No.2018R1A5A1059921, No.2019R1C1C1009192) and Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2019-0-00075, Artificial Intelligence Graduate School Program(KAIST)).

6. References

- [1] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, 2016.
- [2] L. B. Alexey Dosovitskiy and A. K. et al., "An image is worth 16x16 words: Transformers for image recognition at scale," in *ICLR*, 2021.
- [3] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *NeurIPS*, 2017.
- [4] N. R. Tom B. Brown, Benjamin Mann and M. S. et al., "Language models are few-shot learners," in *NeurIPS*, 2020.
- [5] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerrv-Ryan *et al.*, "Natural tts synthesis by conditioning wavenet on mel spectrogram predictions," in *ICASSP*, 2018.
- [6] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *NeurIPS*, 2020.
- [7] E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell, "Adversarial discriminative domain adaptation," in *CVPR*, 2017.
- [8] Y. Ganin and V. Lempitsky, "Unsupervised domain adaptation by backpropagation," in *ICML*, 2015.
- [9] W.-N. Hsu, Y. Zhang, and J. Glass, "Unsupervised domain adaptation for robust speech recognition via variational autoencoder-based data augmentation," in *ASRU*, 2017.
- [10] W. Hou, J. Wang, X. Tan, T. Qin, and T. Shinozaki, "Cross-domain speech recognition with unsupervised character-level distribution matching," in *INTERSPEECH*, 2021.
- [11] S. Sun, B. Zhang, L. Xie, and Y. Zhang, "An unsupervised deep domain adaptation approach for robust speech recognition," in *Neurocomputing*, 2017.
- [12] S. Sun, C.-F. Yeh, M.-Y. Hwang, M. Ostendorf, and L. Xie, "Domain adversarial training for accented speech recognition," in *ICASSP*, 2018.
- [13] V. Manohar, P. Ghahremani, D. Povey, and S. Khudanpur, "A teacher-student learning approach for unsupervised domain adaptation of sequence-trained asr models," in *SLT*, 2018.
- [14] F. Seide, G. Li, X. Chen, and D. Yu, "Feature engineering in context-dependent deep neural networks for conversational speech transcription," in *ASRU*, 2011.
- [15] K. Yao, D. Yu, F. Seide, H. Su, L. Deng, and Y. Gong, "Adaptation of context-dependent deep neural networks for automatic speech recognition," in *SLT*, 2012.
- [16] D. Yu, K. Yao, H. Su, G. Li, and F. Seide, "Kl-divergence regularized deep neural network adaptation for improved large vocabulary speech recognition," in *ICASSP*, 2013.
- [17] G.-T. Lin, S.-W. Li, and H.-y. Lee, "Listen, adapt, better wer: Source-free single-utterance test-time adaptation for automatic speech recognition," in *INTERSPEECH*, 2022.
- [18] Y. Sun, X. Wang, Z. Liu, J. Miller, A. Efros, and M. Hardt, "Test-time training with self-supervision for generalization under distribution shifts," in *ICML*, 2020.
- [19] D. Wang, E. Shelhamer, S. Liu, B. Olshausen, and T. Darrell, "TENT: Fully test-time adaptation by entropy minimization," in *ICLR*, 2021.
- [20] Y. Liu, P. Kothari, B. van Delft, B. Bellot-Gurlet, T. Mordan, and A. Alahi, "TTT++: When does self-supervised test-time training fail or thrive?" in *NeurIPS*, 2021.
- [21] F. Fleuret *et al.*, "Test time adaptation through perturbation robustness," in *NeurIPS DistShift Workshop*, 2021.
- [22] M. Zhang, S. Levine, and C. Finn, "MEMO: Test time robustness via adaptation and augmentation," in *NeurIPS*, 2022.
- [23] D. Chen, D. Wang, T. Darrell, and S. Ebrahimi, "Contrastive test-time adaptation," in *CVPR*, 2022.
- [24] Q. Wang, O. Fink, L. Van Gool, and D. Dai, "Continual test-time domain adaptation," in *CVPR*, 2022.
- [25] A. Bartler, A. Bühler, F. Wiewel, M. Döbler, and B. Yang, "MT3: Meta test-time training for self-supervised test-time adaption," in *AISTATS*, 2022.
- [26] S. Goyal, M. Sun, A. Raghunathan, and Z. Kolter, "Test-time adaptation via conjugate pseudo-labels," in *NeurIPS*, 2022.
- [27] S. Niu, J. Wu, Y. Zhang, Z. Wen, Y. Chen, P. Zhao, and M. Tan, "Towards stable test-time adaptation in dynamic wild world," in *ICLR*, 2023.
- [28] M. Burchi and V. Vielzeuf, "Efficient conformer: Progressive downsampling and grouped attention for automatic speech recognition," in *ASRU*, 2021.
- [29] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *ICML*, 2022.
- [30] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *ICML*, 2006.
- [31] Y. Grandvalet and Y. Bengio, "Semi-supervised learning by entropy minimization," in *NIPS*, 2004.
- [32] Y. Jin, X. Wang, M. Long, and J. Wang, "Minimum class confusion for versatile domain adaptation," in *ECCV*, 2020.
- [33] M. Freitag and Y. Al-Onaizan, "Beam search strategies for neural machine translation," in *ACL*, 2017.
- [34] S. Khurana, N. Moritz, T. Hori, and J. Le Roux, "Unsupervised domain adaptation for speech recognition via uncertainty driven self-training," in *ICASSP*, 2021.
- [35] J. Chen, V. Shah, and A. Kyriillidis, "Negative sampling in semi-supervised learning," in *ICML*, 2020.
- [36] A. Gulati and Q. et al., "Conformer: Convolution-augmented transformer for speech recognition," in *INTERSPEECH*, 2020.
- [37] A. Graves, "Sequence transduction with recurrent neural networks," in *ICML*, 2012.
- [38] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *ICASSP*, 2015.
- [39] J. Barker, R. Marxer, E. Vincent, and S. Watanabe, "The third 'chime' speech separation and recognition challenge: Dataset, task and baselines," in *ASRU*, 2015.
- [40] A. Rousseau, P. Deléglise, Y. Esteve *et al.*, "Enhancing the ted-lium corpus with selected data for language modeling and more ted talks," in *LREC*, 2014.
- [41] R. Ardila, M. Branson, K. Davis, M. Henretty, M. Kohler, J. Meyer, R. Morais, L. Saunders, F. M. Tyers, and G. Weber, "Common voice: A massively-multilingual speech corpus," in *LREC*, 2020.
- [42] C. Valentini-Botinhao *et al.*, "Noisy speech database for training speech enhancement algorithms and tts models," *University of Edinburgh. School of Informatics. Centre for Speech Technology Research (CSTR)*, 2017.
- [43] C. K. Reddy, E. Beyrami, J. Pool, R. Cutler, S. Srinivasan, and J. Gehrke, "A scalable noisy speech dataset and online subjective test framework," in *INTERSPEECH*, 2019.
- [44] G. Zhao, S. Sonsaat, A. Silpachai, I. Lucic, E. Chukharev-Hudilainen, J. Levis, and R. Gutierrez-Osuna, "L2-arctic: A non-native english speech corpus," in *INTERSPEECH*, 2018.
- [45] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *ICLR*, 2017.
- [46] J. Gao, J. Zhang, X. Liu, T. Darrell, E. Shelhamer, and D. Wang, "Back to the source: Diffusion-driven test-time adaptation," in *ICML DynN Workshop*, 2022.