



# Confidence-based Ensembles of End-to-End Speech Recognition Models

Igor Gitman<sup>1</sup>, Vitaly Lavrukhin<sup>1</sup>, Aleksandr Laptev<sup>1,2</sup>, Boris Ginsburg<sup>1</sup>

<sup>1</sup>NVIDIA, USA

<sup>2</sup>ITMO University, Russia

igitman@nvidia.com, vlavrukhin@nvidia.com, alaptev@nvidia.com, bginsburg@nvidia.com

## Abstract

The number of end-to-end speech recognition models grows every year. These models are often adapted to new domains or languages resulting in a proliferation of expert systems that achieve great results on target data, while generally showing inferior performance outside of their domain of expertise. We explore combination of such experts via confidence-based ensembles: ensembles of models where only the output of the most-confident model is used. We assume that models' target data is not available except for a small validation set. We demonstrate effectiveness of our approach with two applications. First, we show that a confidence-based ensemble of 5 monolingual models outperforms a system where model selection is performed via a dedicated language identification block. Second, we demonstrate that it is possible to combine base and adapted models to achieve strong results on both original and target data. We validate all our results on multiple datasets and model architectures.

**Index Terms:** ensembles, confidence, end-to-end, speech recognition, language identification, accent adaptation

## 1. Introduction

Deep neural network (DNN) models are commonly combined in an ensemble to improve predictions' accuracy. Most widely used types of ensembles include bagging [1], boosting [2] and stacking [3]. A less popular ensembling approach is to only use an output of a single model that is deemed best for the current input. A typical way to pick the "best" output is to select a model with the highest confidence score, which provides an estimate of how likely the output is to be correct. We refer to such systems as *confidence-based ensembles* (see Figure 1).

There exists a number of applications of confidence-based ensembles to automatic speech recognition (ASR) and classification. Metze et al. [4] used confidence scores of multiple monolingual Hidden Markov Model (HMM)-based ASR models to improve language identification (LID) performance. This idea has been later applied to build a hybrid DNN-HMM multilingual ASR model where a combination of a dedicated LID block and a confidence-based model selection was used [5, 6]. Wang et al. [7] extended this further to include additional acoustic and language model scores and trained a neural network to make the final model selection. Confidence-based ensembles have also been applied to dialectal and accented speech recognition. Soto et al. [8] showed that a confidence-based hypothesis selection improves recognition results on Arabic dialects. Kukk and Alumäe [9] improved accented speech LID by directly building text classification models on top of the ASR output.

While there are many applications of confidence-based en-

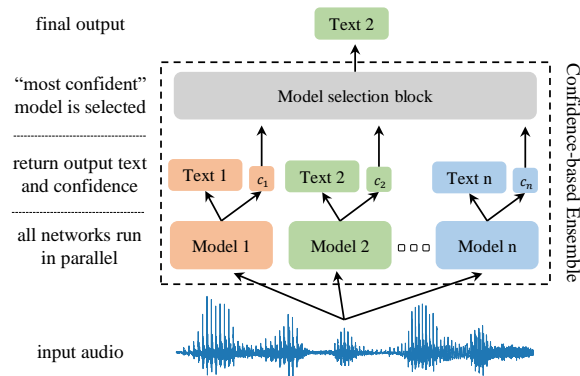


Figure 1: An illustration of a confidence-based ensemble applied to ASR. The same input audio is passed to all models. They run in parallel to produce an output text as well as an estimation of the output's correctness (confidence). A confidence of each model can generally consist of multiple values (e.g., acoustic and language model scores). Model selection block can be implemented as any mapping from confidence values to the index of the "most confident" model.

sembles of HMM-based ASR systems, there is no study on using confidence to combine multiple *end-to-end* neural ASR models. We fill this gap by building ensembles of connectionist temporal classification (CTC) [10] and recurrent neural network transducer (RNNT) [11] models (using Conformer [12] as an encoder). There are different ways how to define confidence measure for such models. A traditional approach is to use a maximum probability of the output tokens [13]. Laptev et al. [14] observed that using entropy-based measures helps to reduce networks' overconfidence. It is possible to train separate models for confidence estimation [15, 16], but such methods are not directly applicable to pretrained models, so we do not consider them in this paper. Instead, we focus on the collection of entropy-based measures proposed in [14] and systematically study their effect on the quality of confidence-based ensembles.

We use two applications to evaluate the effectiveness of our approach. First, we consider a task of multilingual ASR. The current trend in building multilingual end-to-end ASR models is to train a single network that can recognize multiple languages [17, 18, 19]. Another popular line of work is to build mixture-of-expert systems that use separate sub-networks specializing in different languages [20, 21]. However, such techniques require significant computational resources for training and might produce worse recognition than monolingual predictors [18]. An alternative approach is to combine monolingual

Table 1: Comparison of different confidence estimation methods for ensembles of Conformer models. All the numbers in the table represent an average per-dataset accuracy on the validation subsets. “LID 5 languages” is a combination of MCV, MLS and VoxPopuli data. For the “LID 5 languages” and “SLR83” columns we report mean  $\pm$  standard deviation of results across all datasets. “untuned max-prob” confidence is defined as a product of probabilities of emitted tokens (including blanks). “default confidence” corresponds to Rényi entropy with linear normalization, mean aggregation (excluding blanks),  $T = 1.0$ ,  $\alpha = 0.25$ . “tuned confidence” is the best confidence measure, which is generally different for each dataset.

Confidence	Conformer-Transducer			Conformer-CTC		
	LID 5 languages	CORAAL	SLR83	LID 5 languages	CORAAL	SLR83
untuned max-prob	98.18 $\pm$ 0.77	88.00	75.32 $\pm$ 4.35	98.76 $\pm$ 0.82	84.89	57.84 $\pm$ 4.53
default confidence	99.13 $\pm$ 0.44	94.13	81.85 $\pm$ 4.73	99.37 $\pm$ 0.39	91.39	77.08 $\pm$ 5.23
tuned confidence	99.39 $\pm$ 0.32	95.31	93.60 $\pm$ 1.88	99.48 $\pm$ 0.37	93.31	90.37 $\pm$ 4.35

models into a multilingual system using a dedicated language identification (LID) block [6]. We show that confidence-based ensembles outperform dedicated LID models on long audio segments ( $> 5$  seconds) and can be combined with LID model’s scores to improve predictions on short segments.

We then show that it is possible to use confidence-based ensembles for the models sharing the same input language. We consider a task of accent and dialect adaptation. A typical solution to this problem is to finetune an ASR model on the target data. A common issue of this approach is that model’s performance on the original domain might significantly degrade, which is known as catastrophic forgetting [22]. While many techniques have been proposed to address this issue [23, 24], they often require access to the original training data, which might not be available. Majumdar et al. [25] proposes using limited training strategy and regularized adapter modules to reduce the degradation on the source domain without access to the original training data. In this paper we show that a simple confidence-based combination of original and finetuned models achieves significantly better results. Additionally, we demonstrate that the proposed ensemble can interpolate between better performance on source or target domain via a *runtime* change in the weights of the confidence scores.

A clear limitation of most ensemble systems is that required computation grows linearly with each added model. To address this problem, we show that it is possible to use an output of intermediate layers for confidence estimation. We demonstrate this technique through an ensemble of models finetuned with intermediate CTC loss [26] and show that using only 4 out of 18 layers is enough for accurate model selection. We apply this technique to both CTC and Transducer models and find that it generally improves the base models’ accuracy.

Summing up, our findings are as follows:

- Confidence-based ensembles of end-to-end ASR models have better recognition accuracy than a system with a dedicated LID block on long audio segments. Our method provides approximately 10% relative word-error-rate reduction (WERR) compared to the state-of-the-art LID systems on 3 public multilingual datasets.
- Our method can be used to significantly limit catastrophic forgetting during model adaptation. Compared to the constraint adaptation of [25], confidence-based ensembles show 10–50% WERR on the target data, while being 2% relatively better on the source domain.
- Confidence can be reliably estimated from the output of early layers (4 out of 18) if the model was trained with intermediate CTC loss [26]. This reduces the runtime cost of adding a new model to the ensemble by 4.5 times.

## 2. Method

An illustration of confidence-based ensembles is presented in Figure 1. This general structure allows for any vector of confidence values to be produced by each model and any non-linear mapping of the confidence vectors to the model indices to be used as a “model selection block”. In our experiments we adopt a simpler pipeline. We only use a single confidence score for each model and train a logistic regression (LR) on a few audio samples to predict matching model index from the generated confidence values.

### 2.1. Model selection

Formally, our model selection block is defined as follows. Let’s denote the set of models in the ensemble as  $\mathcal{M} = \{m_k\}_{k=1}^M$  and the set of evaluation datasets as  $\mathcal{D} = \{d_i\}_{i=1}^D$ . We assume that there is a surjective mapping  $L : \mathcal{D} \rightarrow \mathcal{M}$  that defines which model is considered “correct” for which dataset. We select a small set of  $N$  utterances<sup>1</sup> of the *training* subset of each dataset  $d_i$  denoted as  $\{t_i^j\}_{j=1}^N$ . Finally, let’s define  $c_k(t_i^j)$  as the confidence score of the model  $k$  computed on the sample  $j$  from the dataset  $i$ . Then we train LR :  $\mathcal{R}^M \rightarrow \mathcal{M}$  to map  $x_i^j = [c_1(t_i^j), c_2(t_i^j), \dots, c_M(t_i^j)] \rightarrow y_i^j = L(d_i)$ .

Recall that logistic regression is a linear classifier that only needs to train  $n + 1$  coefficients for  $n$  models. E.g., in the case of a 2-model ensemble, LR produces the following model selection rule:  $ac_1 + bc_2 > c$ . While a more powerful non-linear functions might achieve better results, using a simple LR allows us to train model selection block with only a few audio samples and helps to avoid overfitting.

### 2.2. Evaluation metrics

To evaluate the accuracy of the ASR models we use a traditional word-error-rate (WER) metric. To evaluate the accuracy of the model selection within ensemble we use the *average per-dataset accuracy* ( $A_{avg}$ ) metric:

$$A_{avg} = \frac{1}{D} \sum_{i=1}^D \frac{1}{D_i} \sum_{j=1}^{D_i} [\text{LR}(c_1(v_i^j), \dots, c_M(v_i^j)) = L(d_i)]$$

where  $v_i^j, j = 1..D_i$  are the *validation/test* utterances from dataset  $d_i$  and  $D_i$  is the size of the validation/test set.  $[X = Y]$  is an indicator function that equals 1 if  $X = Y$  and 0 otherwise. Put simply, an average per-dataset accuracy is a mean of the model prediction accuracies across all datasets. We use

<sup>1</sup>We use  $N = 100$  in our experiments, but for datasets with a lot of variability it might be better to use larger training data.

this metric instead of a regular multi-class prediction accuracy in order to account for label imbalances present in some dataset combinations. E.g., multilingual datasets can be very imbalanced, but we assume that the performance on each language is of equal importance.

### 3. Experiments

For all experiments we followed the same evaluation setup. We fit LR on 100 training utterances and tuned LR’s hyperparameters (regularization strength and class weights in case of imbalanced datasets) maximizing the  $A_{\text{avg}}$  metric on the validation subsets of all datasets. In the following tables we report the results on the test subsets unless noted otherwise. All experiments were performed using NeMo toolkit<sup>2</sup> [27].

#### 3.1. Datasets

##### 3.1.1. Multilingual ASR.

To evaluate the performance of our method we used 3 public multilingual ASR datasets: Mozilla Common Voice (MCV) [28], Multilingual LibriSpeech (MLS) [29] and VoxPopuli [30]. For all datasets we used standard train, validation and test splits. We refer the reader to the original papers for additional information about the used datasets. We built ensembles that supports five languages: English, Italian, Spanish, German and French. For all languages we used “Conformer Large” models from NVIDIA NGC catalog.

##### 3.1.2. Accent and dialect adaptation.

We tested our method on 2 public datasets: Corpus of Regional African American Language (CORAAL) [31] and Open-source Multi-speaker Corpora of the English Accents in the British Isles (SLR83) [32]. We manually split CORAAL into train, validation and test splits containing 18, 3 and 22 hours respectively. The training, validation and test sets contained audio from different speakers. For SLR83 we re-used data splits from [25]. To be able to have direct comparison with [25] we used Conformer Large models trained on LibriSpeech dataset [33] as the base for model adaptation on SLR83 datasets. For CORAAL, we started finetuning from Conformer Large trained on 24325 hours of English speech from NVIDIA NGC catalog.

#### 3.2. Impact of the confidence measure

In this section we study the effect of choosing a confidence estimation method on the ensemble’s quality. We evaluated Gibbs, Tsallis [34] and Rényi [35] entropies (with both linear and exponential normalization) proposed in [14] as well as the maximum probability as confidence measures. To aggregate confidence information across time steps we tested minimum, maximum, mean and product aggregation functions. Both RNNT and CTC models use a special *blank* symbol to align input to the output. We test both including and ignoring blanks in the confidence aggregation. For all measures we tuned the softmax’s temperature  $T \in [0.01, 0.05, 0.1, 0.25, 0.5, 0.75, 1.0, 2.0, 5.0, 10.0]$ . For the entropy-based measures we also ran a grid-search over the parameter  $\alpha \in [0.1, 0.2, 0.25, 0.33, 0.5, 1.0]$ . Overall, the search space spans 2960 combinations of confidence measures.

While the best performing measure for each dataset is typically different, we found that using *Rényi entropy with linear normalization, mean aggregation*,  $T = 1.0, \alpha = 0.25$  and

<sup>2</sup><https://github.com/NVIDIA/NeMo>

Table 2: *LID accuracy and WER on VoxPopuli, MLS and MCV datasets. All numbers show an average of the results across all 5 languages. Confidence-based ensembles outperform all baseline LID systems.*

LID model	VoxPopuli	MLS	MCV
<i>LID accuracy</i>			
ECAPA-TDNN-CE [36]	96.29	98.41	94.42
XLS-R [37]	97.63	98.69	97.47
confidence-based	<b>98.76</b>	<b>99.69</b>	<b>98.82</b>
<i>WER</i>			
oracle	9.44	5.89	6.05
ECAPA-TDNN-CE [36]	11.26	6.90	9.44
XLS-R [37]	10.53	6.78	7.23
confidence-based	<b>9.57</b>	<b>6.04</b>	<b>6.24</b>

Table 3: *Model selection accuracy with different audio duration on VoxPopuli dataset. Only utterances with more than 15 seconds of audio were used to have the same evaluation set for all durations. “Combination” column shows that using a combination of confidence and LID scores can improve the results for short segments.*

Audio duration	Confidence	XLS-R	Combination
3 sec	90.63	93.45	96.00
5 sec	95.50	95.45	97.56
10 sec	98.05	97.07	98.86
15 sec	99.02	97.56	99.46
full audio	99.57	97.67	99.57

*blank symbols excluded* tends to perform well across all evaluated tasks<sup>3</sup>. Table 1 shows the comparison of using the product of probabilities of the output tokens (untuned max-prob), Rényi entropy with the aforementioned parameters (default confidence) and the best confidence measure for each dataset (tuned confidence). The proposed default confidence always outperforms maximum probability and can be further improved if confidence hyperparameters are tuned for each dataset.

Due to the space constraints all subsequent results are reported with default confidence and only for Transducer models.

#### 3.3. Multilingual ASR

In this section we evaluate the performance of the confidence-based ensembles for the multilingual ASR task. We compared our method with several state-of-the-art LID models [36, 37]. Since these models were trained to recognize much larger set of languages, we constrained them to only select the highest probability index out of the 5 evaluated languages. Table 2 shows LID accuracy and WER of the resulting systems on VoxPopuli, MLS and MCV datasets. Table 3 shows how the LID accuracy changes for different audio duration. Our experiments demonstrate that while confidence-based ensembles outperform dedicated LID models on the long audio segments ( $> 5$  seconds), they underperform the best baseline for shorter durations. Last column of Table 3 shows that using a combination of confidence and LID scores as an input to logistic regression obtains significantly better predictions for the short segments.

<sup>3</sup>This measure might not be well-calibrated as it is tuned to optimize the model-selection accuracy.

Table 4: Comparison of the WER for different model adaptation techniques on the SLR83 dataset. The LibriSpeech test-other WER is averaged across all speakers. We abbreviate each speaker id using the first two letters of speaker’s dialect and the first letter of gender (e.g., “Ir-M” stands for “Irish-English Male”). Our “base” model is the same as in [25], but the “finetuned” model is generally better on the target data as we do not constraint training.

Model	LS other	Ir-M	Mi-F	Mi-M	No-F	No-M	Sc-F	Sc-M	So-F	So-M	We-F	We-M
base	5.12	20.69	9.61	11.25	11.11	10.18	12.26	11.94	9.70	10.22	8.51	11.46
finetuned	7.85 ± 2.28	9.17	7.51	8.36	7.60	6.87	8.32	6.35	4.14	3.59	4.67	6.49
constrained [25]	5.40 ± 0.05	15.86	8.40	9.43	9.33	8.54	10.00	8.68	7.73	7.90	6.64	9.70
ensemble	<b>5.26 ± 0.11</b>	<b>9.17</b>	<b>7.51</b>	<b>8.36</b>	<b>7.63</b>	<b>6.78</b>	<b>8.40</b>	<b>6.82</b>	<b>4.64</b>	<b>3.98</b>	<b>4.79</b>	<b>6.70</b>

Table 5: WER after model adaptation on CORAAL dataset. We trained ensemble to select “base” model for both LibriSpeech and VoxPopuli datasets and “finetuned” model for CORAAL dataset. The last two rows demonstrate capability of ensemble to prioritize original vs target domain via a runtime adjustment of logistic regression probability threshold.

Model	LS other	VoxPopuli	CORAAL
base	3.73	6.37	31.23
finetuned	5.25	10.51	8.29
ensemble default	3.82	6.57	8.72
tuned for base	3.76	6.39	9.77
tuned for target	4.31	8.85	8.33

### 3.4. Accent and dialect adaptation

For all finetuning experiments in this section we ran training for 100 epochs on a single 16GB NVIDIA V100 GPU re-using training configuration of the base model, except for the batch size and learning rate parameters. Batch size was selected to fully utilize GPU memory and learning rate was tuned over 8 log-uniform grid points in  $[10^{-7}, 10^{-3}]$ . Table 4 shows the comparison of the confidence-based ensembles with the constraint adaptation approach of [25] on the SLR83 dataset. Our method achieves significantly better accuracy on both original and target domains. Note that for one dialect confidence-based ensemble outperforms the finetuned model. This is possible because the base model can be “incorrectly” selected on utterances where it has lower WER than the finetuned model.

Table 5 shows the results of the confidence-based ensemble after adaptation on the CORAAL dataset. The last two rows of the table demonstrate an ability of the ensemble to trade-off between performance on the target and the original domain via a runtime change of the logistic regression threshold. Based on the specific requirements users can adjust the performance of the model with a simple configuration change.

### 3.5. Decreasing runtime cost

A clear limitation of model ensembles is that runtime cost increases linearly with ensemble size. To partially overcome this, we propose to use outputs of the intermediate layers for model selection<sup>4</sup>. Table 6 shows that it is possible to get a high-quality confidence estimate using intermediate encoder layers. We finetuned all models for 100 epochs on the combination of MCV, MLS and VoxPopuli datasets using intermediate CTC loss [26]

<sup>4</sup>This approach requires finetuning models and thus can only be used if the training data is available.

Table 6: Model selection accuracy using confidence estimation from different layers. All models have been finetuned with intermediate CTC loss [26].

Confidence from	CORAAL	VoxPopuli	MLS
Layer 4	91.58	98.67	99.80
Layer 9	93.32	98.08	99.75
Layer 18 (final)	94.13	98.76	99.69

applied to different layers and used the output of those layers to compute confidence scores. Following the original paper, we added the new loss with coefficient of 0.3. We re-used training configuration of the base models except the learning rate which was 10 times smaller. Unlike the original paper we applied this technique to both CTC and Transducer models and it generally increased models’ accuracy because of the extra regularization. Surprisingly, even though WER of intermediate layers is significantly worse, using intermediate confidence has only a minor reduction in the model-selection accuracy.

## 4. Conclusion

In this paper we showed that confidence-based ensembles of end-to-end ASR models can improve state-of-the-art results for several speech recognition problems. Our method achieves close to 10% WERR on 3 multilingual ASR datasets compared to systems using a dedicated LID block. It also improves results on accent and dialect adaptation tasks by 10–50% WERR.

However, there are several limitations that need to be acknowledged. Confidence-based ensembles are not well suited for latency-critical applications as they require a few seconds of audio to select the most confident model. The runtime cost grows linearly with each added model, which limits the practically useful ensemble size. Finally, given enough compute and data, it is likely possible to build specialized models that would outperform confidence-based ensembles on most tasks.

Taking these limitations into account, we think that confidence-based ensembles can be useful in a wide range of applications beyond what is covered in our experiments. We consider building confidence-based ensembles as a general technique to combine multiple black-box expert models into a single system that achieves competitive results on all target domains. The models can be combined with almost no additional training and without the need to share either data or model weights. We think that these properties can enable new applications of speech technology for users who don’t have access to large compute clusters or big training datasets.

## 5. References

- [1] D. W. Opitz and R. F. Maclin, "An empirical evaluation of bagging and boosting for artificial neural networks," in *Proceedings of International Conference on Neural Networks (ICNN'97)*, vol. 3. IEEE, 1997, pp. 1401–1405.
- [2] H. Schwenk and Y. Bengio, "Boosting neural networks," *Neural computation*, vol. 12, no. 8, pp. 1869–1887, 2000.
- [3] L. Deng, D. Yu, and J. Platt, "Scalable stacking and learning for building deep architectures," in *2012 IEEE International conference on Acoustics, speech and signal processing (ICASSP)*. IEEE, 2012, pp. 2133–2136.
- [4] F. Metze, T. Kemp, T. Schaaf, T. Schultz, and H. Soltau, "Confidence measure based language identification," in *ICASSP*, 2000.
- [5] H. Lin, J.-t. Huang, F. Beaufays, B. Strope, and Y.-h. Sung, "Recognition of multilingual speech in mobile applications," in *ICASSP*, 2012.
- [6] J. Gonzalez-Dominguez, D. Eustis, I. Lopez-Moreno, A. Senior, F. Beaufays, and P. J. Moreno, "A real-time end-to-end multilingual speech recognition architecture," *IEEE Journal of selected topics in signal processing*, vol. 9, no. 4, pp. 749–759, 2014.
- [7] S. Wang, L. Wan, Y. Yu, and I. L. Moreno, "Signal combination for language identification," *arXiv:1910.09687*, 2019.
- [8] V. Soto, O. Siohan, M. Elfeky, and P. Moreno, "Selection and combination of hypotheses for dialectal speech recognition," in *ICASSP*, 2016.
- [9] K. Kukk and T. Alumäe, "Improving language identification of accented speech," in *Interspeech*, 2022.
- [10] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *ICML*, 2006.
- [11] A. Graves, "Sequence transduction with recurrent neural networks," in *ICML*, 2012.
- [12] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu *et al.*, "Conformer: Convolution-augmented transformer for speech recognition," in *Interspeech*, 2020.
- [13] D. Hendrycks and K. Gimpel, "A baseline for detecting misclassified and out-of-distribution examples in neural networks," in *ICLR*, 2017.
- [14] A. Laptev and B. Ginsburg, "Fast entropy-based methods of word-level confidence estimation for end-to-end automatic speech recognition," in *Spoken Language Technology Workshop (SLT)*, 2023.
- [15] A. Woodward, C. Bonnín, I. Masuda, D. Varas, E. Bou-Balust, and J. C. Riveiro, "Confidence measures in encoder-decoder models for speech recognition," in *Interspeech*, 2020, pp. 611–615.
- [16] Q. Li, D. Qiu, Y. Zhang, B. Li, Y. He, P. C. Woodland, L. Cao, and T. Strohmaier, "Confidence estimation for attention-based sequence-to-sequence models for speech recognition," in *ICASSP*, 2021.
- [17] S. Punjabi, H. Arsikere, Z. Raesy, C. Chandak, N. Bhave, A. Bansal, M. Müller, S. Murillo, A. Rastrow, A. Stolcke *et al.*, "Joint ASR and language identification using RNN-T: An efficient approach to dynamic language switching," in *ICASSP*, 2021.
- [18] B. Li, T. N. Sainath, R. Pang, S.-y. Chang, Q. Xu, T. Strohmaier, V. Chen, Q. Liang, H. Liu, Y. He *et al.*, "A language agnostic multilingual streaming on-device asr system," in *Interspeech*, 2022.
- [19] L. Zhou, J. Li, E. Sun, and S. Liu, "A configurable multilingual model is all you need to recognize all languages," in *ICASSP*, 2022.
- [20] N. Gaur, B. Farris, P. Haghani, I. Leal, P. J. Moreno, M. Prasad, B. Ramabhadran, and Y. Zhu, "Mixture of informed experts for multilingual speech recognition," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6234–6238.
- [21] Y. Kwon and S.-W. Chung, "MoLE : Mixture of language experts for multi-lingual automatic speech recognition," in *ICASSP*, 2023.
- [22] R. M. French, "Catastrophic forgetting in connectionist networks," *Trends in cognitive sciences*, vol. 3, no. 4, pp. 128–135, 1999.
- [23] Y. Zhao, C. Ni, C.-C. Leung, S. Joty, E. S. Chng, and B. Ma, "A unified speaker adaptation approach for ASR," in *EMNLP*, 2021.
- [24] D. Hwang, A. Misra, Z. Huo, N. Siddhartha, S. Garg, D. Qiu, K. C. Sim, T. Strohmaier, F. Beaufays, and Y. He, "Large-scale ASR domain adaptation using self-and semi-supervised learning," in *ICASSP*, 2022.
- [25] S. Majumdar, S. Acharya, V. Lavrukhin, and B. Ginsburg, "Damage control during domain adaptation for transducer based automatic speech recognition," in *Spoken Language Technology Workshop (SLT)*. IEEE, 2023, pp. 130–135.
- [26] J. Lee and S. Watanabe, "Intermediate loss regularization for ctc-based speech recognition," in *ICASSP*, 2021.
- [27] O. Kuchaiev, J. Li, H. Nguyen, O. Hrinchuk, R. Leary, B. Ginsburg, S. Kriman, S. Beliaev, V. Lavrukhin, J. Cook *et al.*, "NeMo: a toolkit for building AI applications using neural modules," in *Systems for ML Workshop, NeurIPS*, 2019.
- [28] R. Ardila, M. Branson, K. Davis, M. Kohler, J. Meyer, M. Henretty, R. Morais, L. Saunders, F. Tyers, and G. Weber, "Common voice: A massively-multilingual speech corpus," in *Proceedings of the 12th Language Resources and Evaluation Conference*, 2020, pp. 4218–4222.
- [29] V. Pratap, Q. Xu, A. Sriram, G. Synnaeve, and R. Collobert, "MLS: A large-scale multilingual dataset for speech research," in *Interspeech*, 2020.
- [30] C. Wang, M. Riviere, A. Lee, A. Wu, C. Talnikar, D. Haziza, M. Williamson, J. Pino, and E. Dupoux, "Voxpopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation," in *Proc. of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 2021, pp. 993–1003.
- [31] T. Kendall and C. Farrington, "The corpus of regional African American language. The Corpus of Regional African American Language. Version 2021.07. Eugene, OR," *The Online Resources for African American Language Project*, 2021.
- [32] I. Demirsahin, O. Kjartansson, A. Gutkin, and C. Rivera, "Open-source multi-speaker corpora of the English accents in the British Isles," in *Proc. of the Twelfth Language Resources and Evaluation Conference*, 2020, pp. 6532–6541.
- [33] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an ASR corpus based on public domain audio books," in *ICASSP*. IEEE, 2015.
- [34] C. Tsallis, "Possible generalization of Boltzmann-Gibbs statistics," *Journal of statistical physics*, 1988.
- [35] A. Rényi *et al.*, "On measures of entropy and information," in *Proceedings of the fourth Berkeley symposium on mathematical statistics and probability*, 1961.
- [36] B. Desplanques, J. Thienpondt, and K. Demuynck, "ECAPA-TDNN: Emphasized Channel Attention, Propagation and Aggregation in TDNN Based Speaker Verification," in *Interspeech*, 2020.
- [37] A. Babu, C. Wang, A. Tjandra, K. Lakhotia, Q. Xu, N. Goyal, K. Singh, P. von Platen, Y. Saraf, J. Pino, A. Baevski, A. Conneau, and M. Auli, "XLS-R: self-supervised cross-lingual speech representation learning at scale," in *Interspeech*, 2022.