

AN INVESTIGATION OF END-TO-END MODELS FOR ROBUST SPEECH RECOGNITION

Archiki Prasad¹, Preethi Jyothi², and Rajbabu Velmurugan¹

¹ Department of Electrical Engineering, Indian Institute of Technology Bombay

² Department of Computer Science & Engineering, Indian Institute of Technology Bombay

ABSTRACT

End-to-end models for robust automatic speech recognition (ASR) have not been sufficiently well-explored in prior work. With end-to-end models, one could choose to preprocess the input speech using speech enhancement techniques and train the model using enhanced speech. Another alternative is to pass the noisy speech as input and modify the model architecture to adapt to noisy speech. A systematic comparison of these two approaches for end-to-end robust ASR has not been attempted before. We address this gap and present a detailed comparison of speech enhancement-based techniques and three different model-based adaptation techniques covering data augmentation, multi-task learning, and adversarial learning for robust ASR. While adversarial learning is the best-performing technique on certain noise types, it comes at a cost of degrading clean speech WER. On other relatively stationary noise types, a new speech enhancement technique outperformed all the model-based adaptation techniques. This suggests that knowledge of the underlying noise type can meaningfully inform the choice of adaptation technique.

Index Terms— Robust ASR, Speech Enhancement, Multi-task and Adversarial Learning, Data Augmentation

1. INTRODUCTION

End-to-end (E2E) models, that directly convert a spoken utterance into a sequence of characters, are becoming an increasingly popular choice for ASR systems. They have been shown to outperform traditional cascaded ASR systems when large amounts of labeled speech are available. E2E ASR systems for low-resource scenarios have also emerged as an active new area of research. While E2E ASR systems for clean speech are growing rapidly in number, there have been relatively few investigations on the use of E2E models for noisy speech recognition [1, 2]. To the best of our knowledge, we are the first to provide a detailed comparison of speech enhancement-based techniques with a number of E2E model-based adaptation techniques for noisy speech across a diverse range of noise types. This comparison highlights the strengths and limitations of both types of approaches and offers prescriptions for which techniques are best suited for different noise types.

Prior work on robust ASR has predominantly used a two-pass approach to tackle the problem of robust ASR. The

input speech is first passed through a speech enhancement (SE) module and the enhanced speech is subsequently passed through a standard speech recognition system. We adopt this as one of our approaches as well and investigate the use of three different speech enhancement techniques in conjunction with an E2E ASR system. As opposed to modifying the input speech with front-end processing modules like speech enhancement, one could use the noisy speech as is and adapt the E2E model itself to handle the noisy input speech. We examine three model-based adaptation techniques:

1. **Data Augmentation-based Training (DAT):** Speech samples are augmented with varying noise types of varying signal to noise ratio (SNR) values and fed as input to an E2E model. Larger gradient updates are made in the lower layers of the E2E model compared to the higher layers. This technique first appeared in [2].
2. **Multi-task learning (MTL):** The E2E model is jointly trained with an auxiliary noise type classifier. MTL with a noise type classifier has not been previously explored for robust ASR and turns out to be quite effective as an adaptation technique.
3. **Adversarial Training (AvT):** Unlike MTL which drives the learned representations to be more noise-aware, in AvT, we train an E2E model with a gradient reversal layer in conjunction with the noise classifier to learn more noise-invariant representations.

2. RELATED WORK

We review some relevant supervised robust ASR techniques under noisy conditions and single-channel setting, and do not consider reverberation or multi-channel approaches. A summary of various deep learning techniques for robust ASR, datasets, and benchmarks are provided in [3]. In [4], a noise aware training (NAT) technique that uses a mean noise estimate (assuming stationarity) in the input has been proposed to give good results. This has been further improved in [5], by jointly training a source separation model for noise estimation and acoustic modeling. The work in [6] uses deep convolutional neural network (CNN) to achieve the best ASR results on Aurora-4 task [7].

One approach to improve the performance of E2E models is by using data augmentation along with fine tuning as in [8]. Another approach is to formulate this as a domain adaptation problem, and then use variational auto encoders (VAE) [9]. Here, the source domain is clean speech and target domain is noisy speech. Similarly, [1] uses penalty terms on an encoder-decoder type ASR model to achieve noise invariant representations. The work in [10] uses a jointly adversarial training framework, with a mask-based SE front-end along with an attention-based E2E ASR model.

Models that operate in time-domain, such as SE-GAN [11] and SE-VCAE [12] have been effective in SE and can be used as front-end in a two-pass approach. However, such models rely on large datasets for training and enhancement, but need not necessarily improve ASR. Deep Xi [13], a recent front-end SE has been shown to provide improved ASR when used with DeepSpeech2 [14] (DS2) as the back-end ASR. A more recent approach that operates on raw waveform for real-time speech enhancement is [15]. These two methods will be further explained in Section 4.1, and will be used in our analysis.

3. DATASETS AND E2E ASR SYSTEM

3.1. Implementation Details

In this work, we use DS2 as our main E2E ASR system. DS2 is trained using the Connectionist Temporal Classification (CTC) objective function and comprises two 2D convolutional layers, followed by five bidirectional LSTM layers with a final fully-connected (FC) softmax layer. The input features are derived from short-time Fourier transform (STFT) magnitude spectrograms and the baseline model is trained on 100 hours of clean speech from the Librispeech Dataset [16]. Additional training-specific details can be found in Section 2.2 of [17].¹ Henceforth, we refer to the five LSTM layers as LSTM₀, LSTM₁, LSTM₂, LSTM₃ and LSTM₄.

3.2. Data Description

We are specifically interested in applications that cater to certain noise types for which we do not have a lot of data. There are not many existing datasets that are designed to be low-resource.² Hence, we constructed our own custom dataset.

Our custom dataset consists of the following noise types: ‘Babble’, ‘Airport/Station’, ‘Car’, ‘Metro’, ‘Cafe’, ‘Traffic’, ‘AC/Vacuum’³. The noise samples are sampled at 16 kHz and

¹We use the DS2 implementation available at: <https://github.com/SeanNaren/deepspeech.pytorch>

²The DEMAND dataset [18] exists but it has very few samples per noise type.

³*Airport/Station* comprises background sounds containing announcements at both airports and stations, while *AC/Vacuum* comprises room sounds with an air conditioner or vacuum cleaner in the background. Within these noise types, the noise sub-types are equally distributed in train and test sets.

Method	Objective Scores			
	MOS-LQO	PESQ	STOI	eSTOI
Baseline	1.29	1.80	80.68	60.66
SE-VCAE	1.38	1.83	80.64	62.85
Deep Xi	1.98	2.50	86.89	74.52
DEMUCS	2.17	2.73	91.66	82.19

Table 1. Objective scores for various enhancement methods. Larger scores are better.

were collected from FreeSound [19]. For each noise type, we have 10 and 8 distinct samples in the train and test sets, respectively. The total duration of the noise train and test sets is close to 2 hours. For clean speech, we use the Librispeech corpus [16]. We train the DS2 models using the *train-clean-100* set and add simulated noise using the training samples of our noise dataset. Our development set is constructed using the *dev-clean* set of Librispeech with training noise samples from our noise dataset. During training, the SNR of the additive noise was varied from [0, 25] dB in steps of 5dB. For testing, we randomly picked 120 files from the *test-clean* set of Librispeech. To create noisy speech, for each utterance and each noise type from our noise database, a random section of the noise is added at SNR levels ranging from [0, 20] dB in increments of 5dB, resulting in a total of 4200 noisy speech test utterances. (This process of constructing noisy test samples was outlined in [13].)

4. APPROACHES

4.1. Front-End Speech Enhancement

We experiment with three state-of-the-art speech enhancement techniques detailed below.

SE-VCAE. Speech Enhancement Variance Constrained Autoencoder (SE-VCAE) [12] learns the distribution over latent features given noisy data and acts directly on the time-domain. It outperforms a popular generative modeling SE technique SE-GAN [11]. We finetune the pretrained SE-VCAE model on noisy speech samples from our dataset.

DeepXi. DeepXi was specifically proposed as a front-end to be used with DeepSpeech [13]. This enhancement technique acts on a noisy speech spectrogram and uses a priori SNR estimation to design a mask that is used to produce an estimate of a clean speech spectrogram. DS2 is then fine-tuned for 10 epochs on the enhanced examples from Librispeech while ensuring minimal loss in performance on clean speech data.

DEMUCS. [15] proposes an alternate encoder-decoder architecture for denoising speech samples. The model is trained on a low-resource noisy dataset along with reverberation, with data augmentation techniques to compensate for the limited data. We use their pretrained model and fine-tune it on our dataset for 20 epochs. Similar to DeepXi, DS2 is further finetuned on the denoised samples. Table 1 shows objective scores measuring the quality of enhanced speech using the

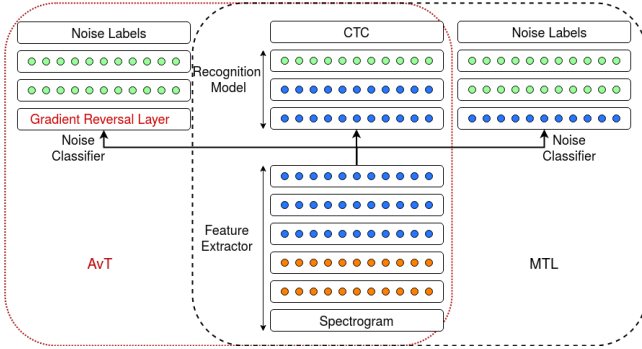


Fig. 1. Framework used in MTL (in black) and AvT (in red). 2D-convolutions, BiLSTM and FC layers are shown using orange, blue and green circles, respectively.

three SE techniques discussed here. DEMUCS clearly outperforms the other two techniques on all four metrics.

4.2. Data Augmentation-based Training (DAT)

For this technique, clean speech samples are augmented with noise with a probability of 0.5 and subsequently used to train DS2. The model was trained for 25 epochs with a batch size of 32 and a learning rate of 0.0001. In the end, the model that performed the best on the development set (with similar noise augmentation) was chosen. We will refer to this model as “Vanilla DAT”. To enable better transfer learning from clean to noisy speech, we incorporate the soft-freezing scheme proposed in [2]. The learning rate of the FC layer along with the last two LSTM layers is scaled down by factor of 0.5 (further discussed in Section 5). This training strategy has the effect of forcing the lower layers (that act as a feature extractor) to learn noise-invariant characteristics in the noisy speech. This model will henceforth be referred to as “Soft-Freeze DAT”.

4.3. Multi-Task Learning (MTL)

Figure 1 describes our MTL setup. The auxiliary classifier predicts noise type labels⁴ and uses representations from an intermediate LSTM layer as its input. This noise classifier comprises one bidirectional LSTM layer followed by two linear layers. The model is trained with a hybrid loss, $L_H = \lambda L_{CTC} + \eta(1 - \lambda)L_{CE}$, where L_{CTC} and L_{CE} are the CTC loss from DS2 and the cross entropy loss on the noise labels, respectively. η and λ are scaling factors and η is annealed by factor of 1.05 every epoch. In our experiments, we initialized the model using Soft-Freeze DAT⁵ and set $\lambda = 0.7$, $\eta = 10$.

⁴There are 8 noise labels: 7 for noise types + 1 for clean speech

⁵We observed that on starting with the baseline model, the initial 10-15 epochs behaved similar to DAT. Hence, Soft-Freeze DAT served as a good initialization and led to faster convergence.

4.4. Adversarial Training (AvT)

Contrary to MTL where the model jointly minimizes the CTC loss and the noise classification loss, adversarial training invokes the use of a gradient reversal layer (GRL) [20] before the auxiliary classifier as shown in Fig 1. This forces the representations before the GRL to be noise-invariant, thus making it hard for the noise classifier to distinguish between noise types⁶. AvT had to be carefully trained with setting different learning rates, λ_f , λ_r and λ_n corresponding to the feature extractor, recognition model and noise classifier, respectively. For our model, the base learning rate was set to 0.0008, $\lambda_f = 0.8$, $\lambda_r = 0.05$ and $\lambda_n = 1$. Similar to MTL, we initialized the model with Soft-Freeze DAT.

5. EXPERIMENTS AND RESULTS

Table 2 lists an exhaustive comparison of all previously mentioned techniques on seven different noise types and five different SNR values. We also report the WER on clean speech to observe degradation with each noise adaptation technique in place. We make a number of observations. Among the SE techniques, DEMUCS which performs best on objective SE scores (as shown in Table 1), also performs best in the ASR experiments. SE-VCAE suffers from distortion of content in the speech signal, which also reflects in its high degradation of clean speech WER. We observe that DeepXi is outperformed (in most noise conditions, and is otherwise matched) by techniques as simple as DAT. The performance mismatch between the reported numbers in DeepXi [13] and our numbers could be attributed to the difference in sizes of our noise datasets; their dataset was much larger in size compared to ours. Interestingly, even with the relatively smaller amounts of noise samples in our dataset, the DEMUCS technique is able to generalize well. On the relatively stationary noise types namely, ‘Car’, ‘Metro’ and ‘Traffic’ (Noise A), DEMUCS outperforms all techniques including all the model-adaptation techniques. For the relatively non-stationary noise types namely, ‘Babble’, ‘Airport/Station’, ‘Cafe’ and ‘AC/Vacuum’ (Noise B) DEMUCS and MTL are statistically very close. AvT yields the largest reductions in WER on Noise B samples (with few exceptions).

Another important distinction to make between the SE techniques and the ML-based techniques is that the SE techniques rely on high-quality pretrained models as a starting point. With our small noise dataset, training the SE model from scratch would not be an option. In contrast, our ML-based techniques are expressive enough to be able to learn from our limited noise datasets without any prior pretraining.

The overall takeaways from this investigation are the following: 1) Among the SE techniques, DEMUCS clearly outperforms the other two SE techniques by large margins. 2)

⁶A setup similar to AvT was explored in [1] using the Musan corpus[21]. Due to lack of noise labels, the classifier output 2 labels: clean and noisy.

Method	WER under SNR (dB)																			
	Babble					Airport/Station					AC/Vacuum					Cafe				
	0	5	10	15	20	0	5	10	15	20	0	5	10	15	20	0	5	10	15	20
Baseline	104.2	98.3	91.3	79.7	65.0	91.9	84.1	73.7	60.6	50.0	93.0	83.1	71.5	59.5	45.8	83.8	72.7	59.5	44.3	33.4
SE-VCAE	85.6	76.4	61.9	54.7	39.7	78.0	68.3	56.8	46.3	39.3	81.3	71.1	61.3	53.6	42.7	61.6	53.9	44.9	35.8	31.0
Deep Xi	81.4	69.4	54.0	44.5	31.9	71.4	60.9	46.5	37.8	27.4	73.9	58.2	45.4	35.1	27.0	52.3	39.7	32.8	25.0	20.4
DEMUCS	70.3	58.0	41.8	32.3	25.4	58.6	45.5	33.7	25.6	21.5	60.5	45.4	34.2	28.1	22.8	38.9	31.6	27.4	20.3	16.9
Vanilla DAT	80.6	68.1	53.6	41.8	30.3	67.1	55.4	41.9	31.2	24.9	66.4	49.8	38.3	31.3	24.5	52.8	41.6	34.5	24.5	19.2
Soft-Freeze DAT	77.4	65.5	52.2	38.5	28.3	64.2	52.9	39.0	29.2	23.7	63.1	46.8	37.1	30.2	23.9	49.1	40.1	33.1	24.4	19.0
MTL	71.4	58.8	45.9	35.5	25.8	55.7	46.8	35.3	26.2	20.7	57.0	41.0	33.0	26.7	21.5	44.9	35.5	24.8	23.6	15.3
AvT	66.8	55.1	39.5	31.1	24.6	53.8	43.3	33.4	25.2	20.9	56.4	40.8	33.4	29.3	23.2	37.1	32.0	26.3	21.4	18.5

Method	WER under SNR (dB)															Clean WER
	Traffic					Metro					Car					
	0	5	10	15	20	0	5	10	15	20	0	5	10	15	20	
Baseline	72.4	62.5	50.2	41.0	33.6	68.4	54.4	46.4	34.9	27.6	35.0	28.1	24.3	21.7	16.7	10.3
SE-VCAE	60.0	51.9	44.2	39.7	32.5	54.0	43.6	38.6	33.0	29.6	35.4	32.7	28.3	27.3	26.0	15.9
Deep Xi	48.0	40.6	29.8	26.0	22.7	44.8	30.5	28.1	20.2	20.5	23.0	19.4	15.4	16.0	14.1	10.9
DEMUCS	38.2	30.3	25.3	20.6	17.9	35.6	24.9	22.6	17.1	15.9	20.5	18.1	14.6	13.8	13.1	10.9
Vanilla DAT	48.5	37.8	31.2	23.3	21.8	41.8	33.1	27.1	21.9	19.1	24.0	19.4	16.1	16.4	14.0	10.8
Soft-Freeze DAT	47.2	35.1	29.8	23.4	20.2	40.8	30.7	27.0	21.3	18.6	23.4	18.9	16.8	15.2	14.7	10.9
MTL	39.9	32.4	29.4	21.3	18.4	38.7	29.2	24.4	20.6	17.3	22.9	19.0	16.3	14.7	14.5	11.0
AvT	40.7	32.5	26.3	21.4	18.5	36.1	26.5	22.6	18.4	17.8	21.8	18.9	16.8	16.0	15.3	13.1

Table 2. Comparison of the performance (WER % after greedy decoding) of all techniques for various noise types and SNRs in the test set. The lowest SE and E2E WERs are shown in bold, and the lower WER among the two is highlighted in green.

Among the ML-based techniques, AvT is largely the best-performing technique across all noise types, with some exceptions. While adversarial training drives the representations to be noise-invariant and helps the noisy speech WERs (evident in low SNR conditions), it has an adverse effect on clean speech and high SNR conditions. AvT incurs the highest WERs on clean speech WER (after SE-VCAE) and sometimes comes second to MTL in high SNR conditions. 3) MTL and AvT are both significantly better than the DAT techniques. Summarily, either of DEMUCS or AvT might be a good choice for noise adaptation but the underlying noise type should also factor into the choice. Table 3 provides an ablation analysis justifying our choices of LSTM₃ and LSTM₄ for

Soft-Freeze DAT, LSTM₂ for MTL and $\lambda_r = 0.05$ for AvT.

6. CONCLUSIONS

In this work, we present a detailed comparison of three speech enhancement techniques and three model-based adaptation techniques for robust E2E ASR across a set of diverse noise types. We observe different trends for different noise types; while adversarial learning yields the largest improvements in performance on non-stationary noise types, a new SE technique DEMUCS gives the best results on relatively stationary noise types. In future work, we will study the interplay of noise types and accents when building robust ASR systems.

Method	Clean WER	WER under SNR (dB)													
		Babble		Airport/Station		Cafe		AC/Vacuum		Car		Metro		Traffic	
		0	15	0	15	0	15	0	15	0	15	0	15	0	15
Soft-Freeze DAT (LSTM ₄)	11.13	79.2	40.5	64.7	30.6	51.6	24.2	63.4	30.8	24.4	16.8	42.2	22.8	48.1	24.5
Soft-Freeze DAT (LSTM ₄ + LSTM ₃)	10.90	77.4	38.5	64.2	29.2	49.1	24.4	63.1	30.2	23.4	15.2	40.8	21.3	47.2	23.4
Soft-Freeze DAT (LSTM ₄ + LSTM ₃ + LSTM ₂)	10.97	79.4	40.4	66.3	30.8	52.0	24.3	63.4	29.4	22.9	15.2	41.5	22.1	47.8	24.0
MTL (noise classifier after LSTM ₁)	11.04	72.3	36.5	56.2	26.5	45.9	23.3	57.9	28.2	22.6	15.2	38.5	20.2	44.0	21.6
MTL (noise classifier after LSTM ₂)	11.05	71.4	35.5	55.7	26.2	44.9	23.6	57.0	26.7	22.9	14.5	38.7	20.6	39.9	21.3
MTL (noise classifier after LSTM ₃)	11.24	75.8	38.7	60.4	27.5	49.0	23.5	60.0	28.2	23.0	15.7	40.4	21.6	46.0	22.4
		15	20	15	20	15	20	15	20	15	20	15	20	15	20
AvT (recognition scaling factor $\lambda_r = 0.2$)	13.79	31.9	26.2	27.8	22.1	21.9	18.9	29.1	23.5	17.1	16.2	20.7	19.4	21.9	19.2
AvT (recognition scaling factor $\lambda_r = 0.1$)	13.37	30.8	24.9	25.4	21.7	21.7	17.8	27.9	23.3	16.2	15.6	20.3	17.1	21.8	19.0
AvT (recognition scaling factor $\lambda_r = 0.05$)	13.07	31.1	24.6	25.2	20.9	21.4	18.5	29.3	23.2	16.0	15.3	18.4	17.8	22.1	18.7

Table 3. WER % after greedy decoding under different settings for DAT, MTL and AvT. The best numbers are shown in bold.

7. REFERENCES

- [1] Davis Liang, Zhiheng Huang, and Zachary C Lipton, “Learning noise-invariant representations for robust speech recognition,” in *2018 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2018, pp. 56–63.
- [2] Shucong Zhang, Cong-Thanh Do, Rama Doddipatla, and Steve Renals, “Learning noise invariant features through transfer learning for robust end-to-end speech recognition,” in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7024–7028.
- [3] Zixing Zhang, Jürgen Geiger, Jouni Pohjalainen, Amr El-Desoky Mousa, Wenyu Jin, and Björn Schuller, “Deep learning for environmentally robust speech recognition: An overview of recent developments,” *ACM Transactions on Intelligent Systems and Technology (TIST)*, vol. 9, no. 5, pp. 1–28, 2018.
- [4] Michael L Seltzer, Dong Yu, and Yongqiang Wang, “An investigation of deep neural networks for noise robust speech recognition,” in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2013, pp. 7398–7402.
- [5] Arun Narayanan and DeLiang Wang, “Joint noise adaptive training for robust automatic speech recognition,” in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2014, pp. 2504–2508.
- [6] Yanmin Qian and Philip C Woodland, “Very deep convolutional neural networks for robust speech recognition,” in *2016 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2016, pp. 481–488.
- [7] David Pearce and J Picone, “Aurora working group: DSR front end LVCSR evaluation AU/384/02,” *Inst. for Signal & Inform. Process., Mississippi State Univ., Tech. Rep*, 2002.
- [8] Tom Ko, Vijayaditya Peddinti, Daniel Povey, Michael L Seltzer, and Sanjeev Khudanpur, “A study on data augmentation of reverberant speech for robust speech recognition,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 5220–5224.
- [9] Wei-Ning Hsu, Yu Zhang, and James Glass, “Unsupervised domain adaptation for robust speech recognition via variational autoencoder-based data augmentation,” in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2017, pp. 16–23.
- [10] Bin Liu, Shuai Nie, Shan Liang, Wenju Liu, Meng Yu, Lianwu Chen, Shouye Peng, and Changliang Li, “Jointly adversarial enhancement training for robust end-to-end speech recognition,” *Proc. Interspeech*, pp. 491–495, 2019.
- [11] Santiago Pascual, Antonio Bonafonte, and Joan Serra, “Segan: Speech enhancement generative adversarial network,” *Proc. Interspeech*, pp. 3642–3646, 2017.
- [12] Daniel T Braithwaite and W Bastiaan Kleijn, “Speech enhancement with variance constrained autoencoders,” in *Proc. Interspeech*, 2019, pp. 1831–1835.
- [13] Aaron Nicolson and Kuldip K Paliwal, “Deep xi as a front-end for robust automatic speech recognition,” *arXiv preprint arXiv:1906.07319*, 2019.
- [14] Dario Amodei, Sundaram Ananthanarayanan, Rishita Anubhai, Jingliang Bai, Eric Battenberg, Carl Case, Jared Casper, Bryan Catanzaro, Qiang Cheng, Guoliang Chen, et al., “Deep speech 2: End-to-end speech recognition in english and mandarin,” in *International conference on machine learning*, 2016, pp. 173–182.
- [15] Alexandre Defossez, Gabriel Synnaeve, and Yossi Adi, “Real time speech enhancement in the waveform domain,” in *Proc. Interspeech*, 2020.
- [16] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, “Librispeech: an asr corpus based on public domain audio books,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [17] Archiki Prasad and Preethi Jyothi, “How accents confound: Probing for accent information in end-to-end speech recognition systems,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 3739–3753.
- [18] Joachim Thiemann, Nobutaka Ito, and Emmanuel Vincent, “DEMAND: a collection of multi-channel recordings of acoustic noise in diverse environments,” in *Proc. Meetings Acoust*, 2013.
- [19] Frederic Font, Gerard Roma, and Xavier Serra, “Freesound technical demo,” in *Proceedings of the 21st ACM international conference on Multimedia*, 2013, pp. 411–412.
- [20] Yaroslav Ganin and Victor Lempitsky, “Unsupervised domain adaptation by backpropagation,” in *International conference on machine learning*, 2015, pp. 1180–1189.
- [21] David Snyder, Guoguo Chen, and Daniel Povey, “Musan: A music, speech, and noise corpus,” *arXiv preprint arXiv:1510.08484*, 2015.