RESCUESPEECH: A GERMAN CORPUS FOR SPEECH RECOGNITION IN SEARCH AND RESCUE DOMAIN

Sangeet Sagar[1,4], Mirco Ravanelli[2], Bernd Kiefer[1,4], Ivana Kruijff-Korbayová[4], Josef van Genabith[1,4]

[1]Saarland University, Germany
[2]Concordia University, Mila-Quebec AI Institute, Canada
[4]German Research Center for Artificial Intelligence (DFKI), Germany

sangeetsagar2020@gmail.com, ravanellim@mila.quebec,
{bernd.kiefer, josef.van_genabith}@dfki.de, ivana.kruijff@rettungsrobotik.de

ABSTRACT

Despite the recent advancements in speech recognition, there are still difficulties in accurately transcribing conversational and emotional speech in noisy and reverberant acoustic environments. This poses a particular challenge in the search and rescue (SAR) domain, where transcribing conversations among rescue team members is crucial to support real-time decision-making. The scarcity of speech data and associated background noise in SAR scenarios make it difficult to deploy robust speech recognition systems.

To address this issue, we have created and made publicly available a German speech dataset called *RescueSpeech*. This dataset includes real speech recordings from simulated rescue exercises. Additionally, we have released competitive training recipes and pre-trained models. Our study highlights that the performance attained by state-of-the-art methods in this challenging scenario is still far from reaching an acceptable level.

Index Terms— speech recognition, search and rescue, noise robustness.

1. INTRODUCTION

Automatic speech recognition (ASR) can be crucial in situations like search and rescue (SAR) missions. These scenarios often involve making critical decisions in extremely hostile conditions, such as underground rescue operations, nuclear accidents, fire evacuation, or collapsed building after an earthquake. In such cases, rescue workers must act quickly and accurately to prevent the loss of lives and damage. Transcribing and automatically analyzing the conversations within the rescue team can provide useful support to help the team make the right decisions in a limited amount of time. The context of search and rescue missions poses significant challenges for current speech recognition technologies. Speech recognizers must be able to handle conversational speech that is fast, emotional, and spoken under stressful conditions. Additionally, the acoustic environment in which rescuers operate is often extremely noisy, and recordings may be corrupted by various non-stationary noises, such as engine noise, vehicle sirens, radio chatter, helicopter noise, and other unpredictable disturbances. In recent years, there has been a significant amount of research focused on addressing these challenges [1–3]. Advanced deep learning techniques, such as self-supervised learning coupled with large datasets [4], have been instrumental in achieving impressive performance improvements. One of the most intriguing aspects of the SAR domain is that all of the aforementioned challenges occur simultaneously, creating an incredibly difficult and complex task. This not only makes it an area of significant scientific interest but also underscores the urgent need for continued research and development in this field.

Developing a speech recognition system in this context is made even more challenging due to the limited availability of data in this critical domain. Collecting speech data related specifically to the SAR domain can be difficult, and privacy restrictions can often limit access to such data by the scientific community. To encourage research in this field, we have released RescueSpeech[1], a German dataset for the Search and Rescue Domain Speech. This dataset contains authentic speech recordings between members of a rescue team during several rescue exercises. To the best of our knowledge, we are the first to publicly release an audio dataset in the SAR domain. RescueSpeech contains approximately 2 hours of annotated speech material. Although this amount may seem limited, it is actually quite valuable and can be effectively used to fine-tune large pretrained models such as wav2vec2.0 [5], WavLM [6], and Whisper [7]. In fact, we demonstrate that this material is also suitable for training models from scratch when combined with proper data augmentation techniques and multi-condition training.

This paper presents a comprehensive collection of experimental evidence for the task at hand– noise-robust German speech recognition. It employs state-of-the-art methods for

---

[1]Available at: https://zenodo.org/record/8077622

both speech recognition and speech enhancement, as well as a combination of the two. Despite excelling in simpler scenarios, our results show that even modern ASR systems like Whisper [7], struggle to perform well in the demanding rescue and search domain. We have made our training recipes and pretrained models available to the community within the SpeechBrain toolkit [2] . With the release of the RescueSpeech dataset we hope to foster research in this field and establish a common benchmark. We believe that our effort can help raise awareness about the importance of the use of speech technology in SAR missions, and the need for continued research in this domain.

2. THE RESCUESPEECH DATASET

RescueSpeech contains a blend of microphone and radio-recorded speech that includes excerpts from communication among robot-assisted emergency response team members during several simulated SAR exercises, which involves real firefighters speaking in high-stress situations like fire rescue, explosion etc. that can elicit heightened emotions. The speakers involved in the exercises are native speakers of the German language where conversations were carried out between team members, radio operators and the team leader. These dialogues loosely adopt a typical radio style communication wherein the start/end of a conversation is indicated by the use of certain words, connection quality is relayed, and acceptance or rejection of requests are conveyed. The practical use case of our dataset is limited not only for robot control but also for speech recognition, with its main application being the support of decision-makers and process monitors in disaster situations. The ASR output is analyzed by a natural language understanding (NLU) component and fused with sensor data, including GPS coordinates from robots or drones. This way we extract mission-related information from conversations and use it to offer assistance later in the deployment of the full system

Initially captured at 44.1 kHz sampling rate, these recordings are down-sampled to 16 kHz, and further segmented to obtain a set of mono-speaker single-channel audio recordings. All utterances are also manually transcribed. The total length of the dataset is 1.6h with a total of 2412 sentences with 1591/245/576 sentences in train/valid/test set. We call it the RescueSpeech clean dataset. Figure 1 shows a histogram plot of the average length of the segmented utterances with an average length of 2.39 sec. We also created a noisy version of RescueSpeech by contaminating our dataset with noisy clips from the AudioSet dataset [8] that includes five noise types– *emergency vehicle siren*, *breathing*, *engine*, *chopper*, and *static radio noise*. We utilized both real and synthetic room-impulse responses (RIR) (SLR26, SLR27 [9]) to add reverberation as well. We then added noisy sequences
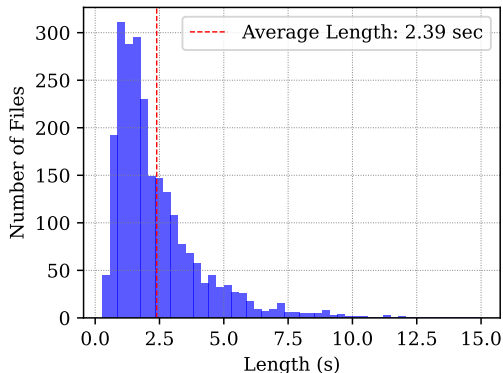
Fig. 1: Histogram plot illustrating the average length of utterances in RescueSpeech in secs.

to generate noisy utterances with different signal-to-noise ratios (SNR) (from -5 dB to 15 dB with a step of 1 dB). Each clean utterance is randomly corrupted with one of the noise types to generate 4500/1350/1350 train/valid/test utterances. We also ensure that a noise utterance used in the train set is only in this set. This randomness and exclusivity ensure that each split has an equal proportion for each noise type and that noises in each of the splits are different. This dataset provides a diverse set of noise and reverberation conditions that enable fine-tuning of our speech-enhancement model for improved accuracy on noisy RescueSpeech. We call this the RescueSpeech noisy dataset. Table 1 briefly shows the distribution of utterances and duration for the clean and noisy version of the dataset.

2.1. Related Corpora

To improve the accuracy of speech recognition systems in noisy and reverberant environments, several corpora have been developed, such as CHIME [10–13], DIRHA [14–17], AMI [18], VOiCES [19], and COSINE [20]. Among these, CHIME5 [12] and CHIME6 [13] are especially challenging because it contains conversational speech recorded during a dinner party in a domestic setting, where noise and reverberations are common. RescueSpeech also contains conversational speech recorded in challenging acoustic environments, but the scenario addressed in this corpus is unique and different from a dinner party. The acoustic conditions, emotions, and lexicon used in RescueSpeech are distinct, and thus provide an additional set of challenges for speech recognition systems.

The noisy version of RescueSpeech can be utilized to train speech enhancement systems that are robust in the acoustic conditions present in the Search and Rescue (SAR) domain. There are numerous datasets that have been released for speech enhancement purposes, including the deep-noise suppression (DNS) dataset [21], VoiceBank-DEMAND cor-

Table 1: Distribution of utterances and hours in the Rescue-Speech clean and noisy dataset.

	Clean		Noisy	
	Mins	#Utts.	HRS	#Utts.
Train	61.86	1591	7.20	4500
Valid	9.61	245	2.16	1350
Test	24.68	576	2.16	1350

pus [22], and WHAM! and WHAMR! corpora [23], all of which are helpful for training speech enhancement models. However, the key difference with RescueSpeech is that it has been specifically designed for the SAR domain, where characteristic sounds such as sirens, radio signals, helicopters, trucks, and others affect the recordings. This unique characteristic of RescueSpeech makes it an especially valuable resource for training speech enhancement systems that can perform well in SAR environments.

3. EXERIMENTAL SETUP

We explored multiple training strategies to perform noise robust speech recognition. Speech recognizers and enhancement models are trained on large corpora and then fine-tuned and evaluated on RescueSpeech data.

3.1. ASR training

We follow two approaches for ASR training: one based on sequence-to-sequence modeling (seq2seq) and another one based on the connectionist temporal classification (CTC) method. For the seq2seq model, we employ a CRDNN (convolutional, recurrent, and dense-neural network) architecture [24, 25]. The CRDNN encoder is trained on the full 1200h of the German CommonVoice corpus [26]. Decoding uses an attentional-GRU decoder and a beam search coupled with an RNN-based language model (LM). The LM is trained on Tuda-De[2] [27] (8M sents), Leipzig news corpus [28] (9M sents), and train transcripts of the CommonVoice corpus. For the CTC based models, we use wav2vec2.0, and WavLM architecture as encoders for the ASR pipeline. These encoders use self-supervised approach for learning high-level contextualized speech representation. It needs no language model and decoding is performed using greedy search. For wav2vec2.0 and WavLM we use pretrained encoders `facebook/wav2vec2-large-xlsr-53-german`[3] and `microsoft/wavlm-large`[4] respectively. Additionally, we also employ the pre-trained Whisper [7] model `openai/whisper`

`-large-v2`[5] to benchmark our systems against competitive state-of-the-art model.

CRDNN combines two blocks of CNN (each block with 2 CNN layers with a channel size (128, 256)), an RNN block (4 bidirectional LSTM layers with 1024 neurons in each layer), and a dense-neural network layer. The inputs are 40-dimensional mel-fiterbank features and the network is trained with an AdaDelta [29] optimizer with a learning rate (LR) of 1 (during fine-tuning we use LR 0.1). The model is trained for 25 epochs with a batch size of 8. During testing, beam search is used with a beam size of 80. Each epoch takes approximately 8h on a single RTXA6000 GPU with 48GB of memory. For wav2vec2.0 and WavLM CTC, training is performed for 45 and 20 epochs respectively with LR 1e-4 on a batch size 8 using an Adam [30] optimizer. Each epoch takes approximately 5.5h on a single RTXA6000 GPU with 48GB of memory. LR is annealed and the sampling frequency is set to 16 kHz for both approaches. More details on training and model parameters can be found in the repository.

3.2. Speech enhancement training

In this work, we perform speech enhancement using SepFormer [31]– a multi-head attention transformer-based source separation architecture. It uses a fully learnable masking-based architecture composed of an encoder, a masking network, and a decoder. The encoder and decoder blocks are essentially convolutional layers and we learn a deep-masking network based on self-attention which estimates element-wise masks. These masks are used by the decoder to reconstruct the enhanced signal in the time-domain. We use the DNS4 [7] dataset to synthesize the training and evaluation set. Using provided clean utterances, noisy clips (150 noises types), and RIRs, we generate 1300h of train and 6.7h of valid set at varying SNR (from -5 dB to 15 dB with a step of 1 dB), and a DNS-2022 baseline dev set is used as test set. Sampling rate is set to 16 kHz and only 30% of clean speech is convolved with RIR.

SepFormer employs an encoder and decoder with 256 convolution filters with kernel size 16, each with stride 8. The masking network has 2 layers of dual-composition block and a chunk length of 250. With each clean-noisy pairs fixed at 4s in length, the model is trained in a supervised fashion using scale-invariant SNR (SI-SNR) loss and Adam optimizer with LR of 1.5e-4. We utilize multi-GPU distributed data parallel (DDP) training scheme to train the network for 50 epochs with a batch size of 4. Each epoch takes approximately 9h on $8\times$ RTXA6000 GPU.

3.3. Training strategies

We use various training methods to create a robust speech recognition system that operates in the SAR (Search and Rescue) domain. These methods are described below:

1. *Clean training*: After pretraining the ASR and Language Model (LM) models, we fine-tune them on the RescueSpeech clean dataset. This process helps to adapt the models to our target domain. We keep the model and training parameters the same as described in Section 3.1.

2. *Multi-condition training*: Using the same pretrained model as above, we perform multi-condition training, which involves training the ASR model on an equal mix of clean and noisy audio from the RescueSpeech noisy dataset. By doing this, the model can learn to adapt to different noises present in the utterances, which helps it to perform speech recognition. This method forms the baseline for all our results. We set the learning rate (LR) to 0.1 and keep other parameters the same as above.

3. *Model-combination I: Independent training*: We pre-train a speech enhancement model and then fine-tune it on the RescueSpeech noisy dataset. This model is then integrated with the ASR model trained in the *clean training* stage to perform noise-robust speech recognition. In this stage, we freeze the enhancement model.

4. *Model-combination II: Joint training*: This is a continuation of the previous stage, where we follow a joint-training approach. We unfreeze the enhancement model and allow gradients from the ASR to propagate back to the speech enhancement model. Updating the weights of the model in this way enables it to generate output that is as clean as possible, as required by the ASR model.

4. RESULTS

4.1. ASR Performance

As a first attempt, we created a simple pipeline consisting solely of an ASR model, with no speech enhancement utilized in the front-end. Table 2 provides a comparison of different ASR models used on both clean and noisy audio recordings from the RescueSpeech dataset. The models included in the comparison are CRDNN, wav2vec2.0, WavLM, and Whisper. During the pre-training stage, all models (except Whisper) utilized only the CommonVoice dataset. However, during the clean training and multi-condition fine-tuning stage, the RescueSpeech dataset was used.

Unsurprisingly, the clean training approach is the most effective when tested on clean audio recordings. The top-performing model in this scenario is Whisper, which achieved

Table 2: Comparison of test WERs for CRDNN, wav2vec2.0-large, WavLM-large, and whisper-large-v2 models using different training strategies on clean and noisy speech inputs from the RescueSpeech dataset.

	ASR Model	clean	noisy
Pre-training	CRDNN	52.03	81.14
	Wav2vec2	47.92	76.98
	WavLM	46.28	73.84
	Whisper	27.01	50.85
Clean training	CRDNN	31.18	60.10
	Wav2vec2	27.69	62.60
	WavLM	23.93	58.28
	Whisper	**23.14**	46.70
Multi-cond. training	CRDNN	33.22	58.95
	Wav2vec2	29.89	57.98
	WavLM	25.22	52.75
	Whisper	24.11	**45.84**

Table 3: Speech enhancement performance on the Rescue-Speech noisy test inputs when combining speech enhancement and speech recognition (Model Comb. I vs Model Comb. II).

	Model Comb. I	Model Comb. II			
		CRDNN	wav2vec2	WavLM	Whisper
SI-SNRi	6.516	6.618	7.205	7.140	7.482
SDRi	7.439	7.490	7.765	7.694	8.011
PESQ	2.008	2.010	2.060	2.064	2.083
STOI	0.842	0.844	0.854	0.854	0.859

Table 4: Word-Error-Rate (WER%) achieved with independent training (Model Comb. I) and joint training (Model Comb. II) of the speech enhancement and ASR modules.

ASR Model	Model Comb. I	Model Comb. II
CRDNN	54.98	54.55
Wav2vec2	50.68	49.24
WavLM	48.24	46.04
Whisper	48.04	**45.29**

a WER of 23.14%. On the other hand, multi-condition training proved to be a superior strategy when dealing with noisy recordings. In this scenario, the best model is again Whisper, which achieved a WER of 45.84%. The performance gap with clean signals, highlights one more time the significant decline in recognition performance when dealing with challenging acoustic conditions, even for models that were pre-trained using state-of-the-art self-supervised techniques like wav2vec, WavLM, and Whisper (the latter of which is even semi-supervised).
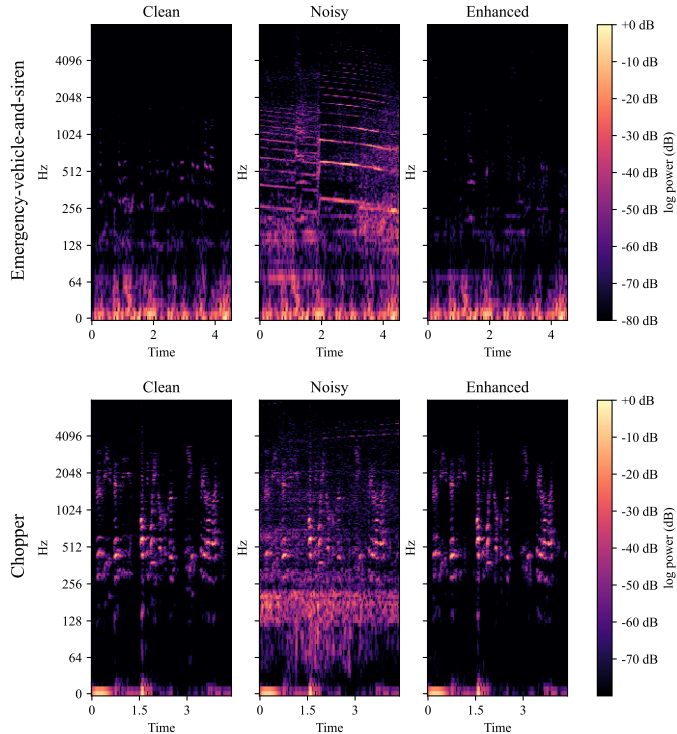
Fig. 2: Log-power spectrogram of clean, noisy, and SepFormer-enhanced utterances for *emergency vehicle siren* and *chopper* noise types at -5 dB SNR.

Fig. 3: PESQ, SDRi, SI-SNRi vs SNR of SepFormer enhanced utterances for two noise types– *emergency vehicle siren* and *chopper* noise.

4.2. Combining ASR and Speech Enhancement

In order to improve the ASR performance, we developed a speech enhancement system to clean up the recordings. To accomplish this, we utilized the SepFormer model, which has demonstrated competitive performance in speech separation and enhancement tasks [32]. Specifically, we trained the model on the DNS4 dataset, achieving SIG, BAK, and OVRL scores of 2.999, 3.076, and 2.437, respectively. Figure 2 shows the log-power spectrogram for two types of noisy audio recordings, *emergency vehicle siren* and *chopper noise*, both with an SNR of -5 dB, using the SepFormer model fine-tuned on the RescueSpeech noisy dataset. From a qualitative standpoint, it appears that SepFormer performs well on noises that impact the SAR domain. Figure 3 presents PESQ vs SNR and SI-SNRi, SDRi vs SNR for the same noise types. We observed that improvements in SI-SNR and SDR were greater for utterances with an SNR of -5 dB, indicating a more significant enhancement in speech intelligibility and reduction of distortion than for higher SNR utterances. This pattern is consistent across all noise types.

Table 3 displays the speech enhancement results obtained by incorporating a speech recognizer into the pipeline. In section 3.3, we explored two approaches: independent training (Model Comb. I) and joint training (Model Comb. II).
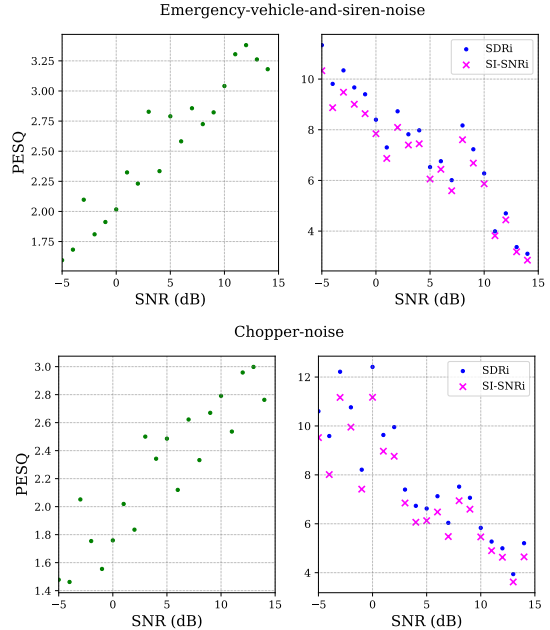
The joint training approach resulted in improvements across all considered speech enhancement metrics (SI-SNRi, SDRi, PESQ, STOI) and all ASR modules (CRDNN, Wav2vec2, WavLM, Whisper). Table 4 presents the final speech recognition output at the end of the pipeline.

As anticipated, the joint training approach outperformed a simple combination of independently trained speech enhancement and speech recognition modules. It is important to note that both speech enhancement and speech recognition models undergo fine-tuning using enhanced signals from the unfrozen Sepformer. We postulate that backpropagating the ASR gradient to the speech enhancement model enables the SepFormer to denoise utterances according to the specific requirements of the ASR model, facilitating better convergence. Training both models jointly allows the enhancement model to adapt its cleaning capabilities to align better with the needs of the ASR system. Overall, the best-performing model is the combination of SepFormer with Whisper ASR, which achieved a WER of 45.29%.

5. CONCLUSIONS

Our work addresses some major challenges that arise in the SAR domain: the lack of speech data, the need for robustness to SAR noises, and conversational speech. To overcome these challenges, we have introduced RescueSpeech, a new dataset of speech data in German that we use to perform robust speech recognition in a hostile noise-filled en-

vironment. To achieve this, we proposed multiple training strategies that involve fine-tuning pretrained models on our in-domain data. We tested different self-supervised models (e.g, Wav2Vec2, WavLM, and Whisper) for speech recognition. Despite leveraging these cutting-edge systems, our best model only achieves a WER of 45.29% on our test set. This result highlights the significant difficulty and the urgent need for further research in this crucial domain.

Overall, our work represents a step forward in addressing the challenges of speech recognition in the SAR domain. By introducing a new dataset, we hope to establish a useful benchmark and foster more studies in this field.

6. ACKNOWLEDGEMENTS

7. REFERENCES

[1] Christian Willms, Constantin Houy, Jana-Rebecca Rehse, Peter Fettke, and Ivana Kruijff-Korbayová, "Team Communication Processing and Process Analytics for Supporting Robot-Assisted Emergency Response," in *2019 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, 2019.

[2] Aylin Gözalan, Ole John, Thomas Lübcke, Andreas Maier, Maximilian Reimann, Jan-Gerrit Richter, and Ivan Zverev, "Assisting Maritime Search and Rescue (SAR) Personnel with AI-Based Speech Recognition and Smart Direction Finding," *Journal of Marine Science and Engineering*, vol. 8, no. 10, 2020.

[3] Saeid Mokaram and Roger K. Moore, "The Sheffield Search and Rescue corpus," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 5840–5844.

[4] Abdelrahman Mohamed, Hung yi Lee, Lasse Borgholt, Jakob D. Havtorn, Joakim Edin, Christian Igel, Katrin Kirchhoff, Shang-Wen Li, Karen Livescu, Lars Maaloe, Tara N. Sainath, and Shinji Watanabe, "Self-Supervised Speech Representation Learning: A Review," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1179–1210, oct 2022.

[5] Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli, "wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations," 2020.

[6] Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, Jian Wu, Long Zhou, Shuo Ren, Yanmin Qian, Yao Qian, Jian Wu, Michael Zeng, Xiangzhan Yu, and Furu Wei, "WavLM: Large-Scale Self-Supervised Pre-Training for Full Stack Speech Processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, oct 2022.

[7] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever, "Robust Speech Recognition via Large-Scale Weak Supervision," 2022.

[8] Jort F. Gemmeke, Daniel P. W. Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R. Channing Moore, Manoj Plakal, and Marvin Ritter, "Audio Set: An ontology and human-labeled dataset for audio events," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 776–780.

[9] 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)*, 2017, pp. 5220–5224.

[10] Jon Barker, Ricard Marxer, Emmanuel Vincent, and Shinji Watanabe, "The Third CHiME Speech Separation and Recognition Challenge," *Comput. Speech Lang.*, vol. 46, no. C, pp. 605–626, nov 2017.

[11] E. Vincent, S. Watanabe, A. Nugraha, J. Barker, and R. Marxer, "An analysis of environment, microphone and data simulation mismatches in robust speech recognition," *Computer Speech and Language*, vol. 46, pp. 535–557, 2017.

[12] Jon Barker, Shinji Watanabe, Emmanuel Vincent, and Jan Trmal, "The fifth 'CHiME' Speech Separation and Recognition Challenge: Dataset, task and baselines," in *Proc. of Interspeech*, 2018.

[13] Shinji Watanabe et al., "CHiME-6 Challenge: Tackling Multispeaker Speech Recognition for Unsegmented Recordings," in *Proc. 6th International Workshop on Speech Processing in Everyday Environments (CHiME 2020)*, 2020.

[14] Mirco Ravanelli, Luca Cristoforetti, Roberto Gretter, Marco Pellin, Alessandro Sosi, and Maurizio Omologo, "The DIRHA-ENGLISH corpus and related tasks for distant-speech recognition in domestic environments," in *Proc. of ASRU*, 2015.

[15] Marco Matassoni, Ramón Fernandez Astudillo, Athanasios Katsamanis, and Mirco Ravanelli, "The DIRHA-GRID corpus: baseline and tools for multi-room distant speech recognition using distributed microphones," in *Proc. of Interspeech*, 2014.

[16] Mirco Ravanelli and Maurizio Omologo, "On the selection of the impulse responses for distant-speech recognition based on contaminated speech training," in *Proc. of Interspeech*, Haizhou Li, Helen M. Meng, Bin Ma, Engsiong Chng, and Lei Xie, Eds., 2014.

[17] Mirco Ravanelli and Maurizio Omologo, "Contaminated speech training methods for robust DNN-HMM distant speech recognition," in *Proc. of Interspeech*, 2015.

[18] Steve Renals, Thomas Hain, and Herve Bourlard, "Recognition and interpretation of meetings: The AMI and AMIDA projects," in *Proc. of ASRU*, 2007.

[19] Colleen Richey, Maria A. Barrios, Zeb Armstrong, Chris Bartels, Horacio Franco, Martin Graciarena, Aaron Lawson, Mahesh Kumar Nandwana, Allen Stauffer, Julien van Hout, Paul Gamble, Jeff Hetherly, Cory Stephenson, and Karl Ni, "Voices Obscured in Complex Environmental Settings (VOICES) corpus," 2018.

[20] Alex Stupakov, Evan Hanusa, Deepak Vijaywargi, Dieter Fox, and Jeff A. Bilmes, "The design and collection of COSINE, a multi-microphone in situ speech corpus recorded in noisy environments," *Comput. Speech Lang.*, vol. 26, no. 1, pp. 52–66, 2012.

[21] Harishchandra Dubey, Vishak Gopal, Ross Cutler, Ashkan Aazami, Sergiy Matusevych, Sebastian Braun, Sefik Emre Eskimez, Manthan Thakker, Takuya Yoshioka, Hannes Gamper, and Robert Aichner, "ICASSP 2022 Deep Noise Suppression Challenge," 2022.

[22] Christophe Veaux, Junichi Yamagishi, and Simon King, "The voice bank corpus: Design, collection and data analysis of a large regional accent speech database," in *2013 International Conference Oriental COCOSDA held jointly with 2013 Conference on Asian Spoken Language Research and Evaluation (O-COCOSDA/CASLRE)*, 2013, pp. 1–4.

[23] Gordon Wichern, Joe Antognini, Michael Flynn, Licheng Richard Zhu, Emmett McQuinn, Dwight Crow, Ethan Manilow, and Jonathan Le Roux, "WHAM!: Extending Speech Separation to Noisy Environments," in *Proc. Interspeech*, Sept. 2019.

[24] Tara N. Sainath, Oriol Vinyals, Andrew Senior, and Haşim Sak, "Convolutional, Long Short-Term Memory, fully connected Deep Neural Networks," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 4580–4584.

[25] Yusheng Xiang, Tian Tang, Tianqing Su, Christine Brach, Libo Liu, Samuel S. Mao, and Marcus Geimer, "Fast CRDNN: Towards on Site Training of Mobile Construction Machines," *IEEE Access*, vol. 9, pp. 124253–124267, 2021.

[26] Rosana Ardila, Megan Branson, Kelly Davis, Michael Henretty, Michael Kohler, Josh Meyer, Reuben Morais, Lindsay Saunders, Francis M. Tyers, and Gregor Weber, "Common Voice: A Massively-Multilingual Speech Corpus," 2019.

[27] Benjamin Milde and Arne Koehn, "Open Source Automatic Speech Recognition for German," in *Speech Communication; 13th ITG-Symposium*, 2018, pp. 1–5.

[28] Dirk Goldhahn, Thomas Eckart, and Uwe Quasthoff, "Building Large Monolingual Dictionaries at the Leipzig Corpora Collection: From 100 to 200 Languages," in *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, Istanbul, Turkey, May 2012, pp. 759–765, European Language Resources Association (ELRA).

[29] Matthew D. Zeiler, "ADADELTA: An Adaptive Learning Rate Method," 2012.

[30] Diederik P. Kingma and Jimmy Ba, "Adam: A Method for Stochastic Optimization," 2014.

[31] Cem Subakan, Mirco Ravanelli, Samuele Cornell, Mirko Bronzi, and Jianyuan Zhong, "Attention is All You Need in Speech Separation," 2020.

[32] Cem Subakan, Mirco Ravanelli, Samuele Cornell, Francois Grondin, and Mirko Bronzi, "On Using Transformers for Speech-Separation," 2022.