

A Curriculum Learning Method for Improved Noise Robustness in Automatic Speech Recognition

Stefan Braun, Daniel Neil, Shih-Chii Liu

Institute of Neuroinformatics

University of Zurich and ETH Zurich

Zurich, Switzerland

Email: brauns@student.ethz.ch, dneil@ini.uzh.ch, shih@ini.uzh.ch

Abstract—The performance of automatic speech recognition systems under noisy environments still leaves room for improvement. Speech enhancement or feature enhancement techniques for increasing noise robustness of these systems usually add components to the recognition system that need careful optimization. In this work, we propose the use of a relatively simple curriculum training strategy called **accordion annealing (ACCAN)**. It uses a multi-stage training schedule where samples at signal-to-noise ratio (SNR) values as low as 0dB are first added and samples at increasing higher SNR values are gradually added up to an SNR value of 50dB. We also use a method called **per-epoch noise mixing (PEM)** that generates noisy training samples online during training and thus enables dynamically changing the SNR of our training data. Both the ACCAN and the PEM methods are evaluated on an end-to-end speech recognition pipeline on the Wall Street Journal corpus. ACCAN decreases the average word error rate (WER) on the 20dB to -10dB SNR range by up to 31.4% when compared to a conventional multi-condition training method.

I. INTRODUCTION

The performance of automatic speech recognition (ASR) systems has increased significantly with the use of deep neural networks (DNNs) [1]. However, their performance in noisy environments still leaves room for improvement. Over the past decades, a multitude of methods to improve noise robustness of ASR systems has been proposed [2], with many methods being applicable to DNNs. These methods enhance noise robustness at various levels and are applied prior to feature extraction, at the feature level, and during training.

Some example enhancement methods applied prior to feature extraction include denoising methods [3] and source separation methods [4] [5]. Methods applied at the feature level include methods that produce auditory-level features [6] and feature space adaptation methods [7]. Other approaches use DNNs, e.g. feature denoising with deep autoencoders [8] [9] or feature extraction from the raw waveform via convolutional neural networks (CNNs) [10] [11]. Many of these strategies add components to the speech recognition system that need careful optimization.

The training method itself can have a major influence on the performance of a neural network under noisy conditions. Training on noisy data is an established method of increasing the noise robustness of a network. Noisy training sets with a range of SNR values e.g. 10 dB - 20 dB [12] or 0 dB - 30 dB [10] are used during training. Other training methods such as

dropout [13] - originally intended to improve regularisation - have been shown to also improve noise robustness [12]. The same is true for model adaptation/noise aware training techniques [12].

This paper presents general training methods for improving noise robustness in recurrent neural network (RNN)-based recognizers. RNNs are used here because they have demonstrated state-of-the-art performance in tasks such as the common sequence labelling task in speech recognition [14] [15].

In particular, we introduce a new training strategy called **accordion annealing (ACCAN)** which exploits the benefits of curriculum-based training methods. By first training the network on low SNR levels down to 0dB and gradually increasing the SNR range to encompass higher SNR levels, the trained network shows better noise robustness when tested under a wide range of SNR levels.

This work also investigates the usefulness of adding noise both at the acoustic waveform level and at the feature representation level during training. In particular, we exploit a method called **per-epoch noise mixing (PEM)**, which is a waveform-level data augmentation method. It enables us to generate a new training set for every epoch of training, i.e. each training sample is mixed with a newly sampled noise segment at a randomly chosen SNR in every epoch. This form of data augmentation prevents networks from relying on constant noise segments for classification and helps in creating the necessary training samples over a wide SNR range. These steps lead to improved generalization and noise robustness of the trained network. Our results are evaluated on the *Wall Street Journal* corpus in a large-vocabulary continuous speech recognition (LVCSR) task. The testing is carried out on a large SNR range going from clean conditions (> 50 dB) down to -20 dB.

The paper is organized as follows: Section II presents our training methods for improved noise robustness. The evaluation setup is detailed in Section III, with results given in Section IV followed by discussions in Section V and concluding remarks in Section VI.

II. TRAINING METHODS FOR IMPROVED NOISE ROBUSTNESS

A. Baseline

Our baseline method takes advantage of multi-condition training [16] in order to increase the noise robustness of the network. Pink noise is added to a clean dataset to create samples with the desired SNR. Each training sample is randomly chosen to be of an SNR level in the range 0 to 50 dB with 5 dB steps. This wide range is larger than the SNR ranges used in previous work (e.g. 0 to 30 dB as in [10]). Our exhaustive simulations show that using such a large range resulted in the best performance on the test datasets. The noise mixing is done once at the waveform-level before filterbank audio features are computed. This one set of training data is presented to the network over all training epochs. The resulting networks will be referred to as “noisy-baseline”. For completeness, we also include a “clean-baseline”, i.e. a network that is only trained on clean speech.

B. Gaussian noise injection

Gaussian noise injection is a well-known method for improving generalisation in neural networks [17]. It is used here to improve the noise robustness of the network.

During training, artificial Gaussian noise is added to the filterbank features created from the different SNR samples. The amount of additive noise is drawn from a zero-centered Gaussian distribution. Using a Gaussian with a standard deviation of $\sigma = 0.6$ yielded the best results. This method is referred to as the “Gauss-method” in the rest of the paper.

C. Per-epoch noise mixing (PEM)

PEM is a method for adding noise to the waveform level during training. In every training epoch, each training sample is mixed with a randomly sampled noise segment at a randomly sampled SNR. The training procedure consists of the following steps:

- 1) Mix every training sample with a randomly selected noise segment from a large pink noise pool to create a resulting sample at a randomly chosen SNR level between 0 to 50 dB.
- 2) Extract audio features (e.g. filterbank features) for the noise-corrupted audio to obtain the training data for the current epoch.
- 3) *Optional*: add Gaussian noise to the audio features.
- 4) Train on newly generated training data for one epoch.
- 5) Discard this training data after the epoch to free up storage.
- 6) Repeat from step 1 until training terminates.

This method has several key advantages over conventional pre-training preprocessing methods. Firstly, it enables unlimited data augmentation on large speech datasets. With conventional methods, augmenting training data at the waveform level with real-world noise at various SNR values is prohibitively expensive in terms of processing time and training data size. PEM allows use to overcome these restrictions by training

on the GPU and pre-processing the next-epoch training data in parallel on the CPU. After an epoch was trained on, the training data gets discarded to free storage for the next epoch.

Secondly, PEM shows the network more unique training data: every training sample is presented at a selection of SNRs and with as many noise samples as can be extracted from the noise file and as needed by the number of epochs to reach a steady-state accuracy level. Thirdly, other noise types, different SNR training ranges, and even different audio features could be quickly tested as the training data can be easily augmented online. Finally, PEM enables us to dynamically change the SNR level during training, which renders advanced training paradigms such as curriculum learning (section II-D) feasible.

In contrast to the Gauss-method, PEM permits more control over the training data. Real-world noise is added to the acoustic waveform at controlled SNRs, ensuring that the training data corresponds to realistic noise corruption with results that can be evaluated. Of course, PEM can be combined with Gaussian noise addition (optional step three in Section II-C). We refer to PEM without Gaussian noise injection as “Vanilla-PEM” and to PEM with Gaussian noise injection as “Gauss-PEM”.

D. Curriculum learning

Neural networks have been shown to optimize on the SNR they are trained on [16]. A network trained on clean conditions thus fares worse than a network trained on noisy conditions. Also, networks trained on a vast SNR range generally do worse on a single SNR than networks optimized for this specific SNR. In order to achieve high accuracy under both high and low SNR with a single network, we explored novel training paradigms based on curriculum learning. While curriculum learning has been used in image classification (scheduled denoising autoencoders, [18]) as well as speech recognition (SortaGrad [15], a method for faster accuracy convergence), this is the first work targeted at LVCSR under noisy conditions.

Our novel ACCAN training method applies a multi-stage training schedule: in the first stage, the neural network is trained on the lowest SNR samples. In the following stages, the SNR training range is expanded in 5 dB steps towards higher SNR values. A typical schedule is shown in Table I. In every stage, training repeats until the WER on the development set no longer improves. At the end of each stage, the weights of the best network are stored and used as the starting point for the next stage. Both training and validation sets share the same SNR range. The ACCAN approach seems counter-intuitive as the network is first trained on noisy data that is difficult to classify. However, the noise allows the network to explore the parameter space more extensively at the beginning [19]. We also evaluated a method called “ACCAN-reversed”, which expands from high SNR to low SNR, but the results were very close to the standard “Gauss-PEM” approach.

III. SETUP

Audio database: All experiments were carried out on the Wall Street Journal (WSJ) corpus (LDC93S6B and

TABLE I
THE ACCAN TRAINING STRATEGY: SNR RANGES [dB] OF THE TRAINING STAGES

Method	Stage 1	Stage 2	...	Stage 10
ACCAN	[0]	[0, 5]	...	[0, 5, ..., 45, 50]
ACCAN reversed	[50]	[50, 45]	...	[50, 45, ..., 5, 0]

LDC94S13B) in the following configuration:

- training set: *train-si84* (7138 samples, 15h of speech),
- development set: *test-dev93* (503 samples, 1h of speech),
- test set: *test-eval92* (333 samples, 0.7h of speech).

For noise corruption, we used two different noise types: *pink noise* generated by the Audacity [20] software and *babble noise* from the NOISEX database [21].

Data preparation and language model: The labels and transcriptions were extracted with EESSEN [14] routines. All experiments were character-based and used 58 labels (letters, digits, punctuation marks etc.) During test time, the network output was decoded with the Weighted Finite State Transducer (WFST) approach from the EESSEN framework, which allows us to apply a trigram language model. The language model used an expanded vocabulary in order to avoid out-of-vocabulary words occurring in the standard WSJ language model.

Audio features: We used 123-dimensional filterbank features that consisted of 40 filterbanks, 1 energy term and their respective first and second order derivatives. The features were generated by preprocessing routines from EESSEN [14]. Each feature dimension is zero-mean and unit-variance normalized.

Neural network configuration: Our recognition pipeline is a end-to-end solution that involves a RNN as the acoustic model. In order to automatically learn the alignments between speech frames and label sequences, the Connectionist Temporal Classification (CTC) [22] objective was adopted. The Lasagne library [23] enabled us to build and train our 5-layer neural network. The first 4 layers consisted of bidirectional long short-term memory (LSTM) [24] units with 250 units in each direction. The fifth and final layer was a non-flattening dense layer with 59 outputs, corresponding to the character labels + the blank label required by CTC. The network contained 8.5M tunable parameters. All layers were initialized with the Glorot uniform strategy [25]. Every experiment started with the exact same weight initialization. During training, the Adam [26] stochastic optimization method was used. To prevent overfitting and to increase noise robustness, dropout [13] was used (dropout probability=0.3). Every epoch of training, the WER on the development set was monitored with a simple best-path decoding approach.

With all training strategies except ACCAN, the network was trained for a generous 150 epochs. The networks weights from the epoch with the lowest WER were kept for evaluation. Generally, the improvements in WER saturated well before 150 epochs were reached. The ACCAN method used a patience of 5 as to switch between SNR stages, i.e. if the WER did not improve for 5 epochs on the current SNR stage, the training continued on the next SNR stage. By respecting the stage-switching policy, ACCAN reached the final SNR stage with

the full SNR range at epoch 190. Saturation kicked in at epoch 240. While ACCAN trained for more epochs than the others, it only trained for 50 epochs on the full SNR range.

IV. RESULTS

The reported results are given for the 'test-eval92' evaluation set for the *Wall Street Journal* corpus. The evaluation set was tested in clean condition and with added pink noise or babble noise at 15 SNR levels from 50dB to -20 dB in dB steps. We report in Table II the average WER over following SNR ranges:

- Full SNR range: [clean signal, 50dB to -10dB]
- High SNR range: [50dB to 0dB]
- Low SNR range: [0dB to -10dB]
- range of interest (ROI): [20dB to -10dB]

We choose to include the ROI, as our hearing tests showed that this range seems to well reflect common scenarios in public environments, where a clean speech signal is most often not found. Detailed results for each SNR individually are given in Table III. Results for -15dB and -20dB are reported too, but should be considered as extreme cases. WER improvements are given as relative improvements in the text.

TABLE II
AVERAGE ABSOLUTE WER [%] FOR GIVEN SNR RANGES AFTER DECODING. PRINTED BOLD: LOWEST WER.

Testing against pink noise				
Method	Full	High	Low	ROI
Clean-baseline	54.7	29.0	109.6	67.9
Noisy-baseline	46.0	23.3	88.6	51.7
Gauss	37.4	19.8	71.1	42.1
Vanilla-PEM	35.6	17.8	70.6	40.8
Gauss-PEM	34.1	16.6	64.7	37.2
ACCAN	34.4	18.1	59.5	36.0
ACCAN reversed	35.2	17.8	66.3	38.8
Testing against babble noise				
Method	Full	High	Low	ROI
Clean-baseline	53.0	32.0	113.7	72.1
Noisy-baseline	53.3	29.9	114.0	68.4
Gauss	45.4	25.4	96.3	56.9
Vanilla-PEM	41.0	22.8	88.3	52.3
Gauss-PEM	39.5	21.6	83.7	49.0
ACCAN	39.6	21.5	80.2	47.0
ACCAN reversed	39.5	21.5	82.9	48.2

A. Noise addition methods

This section summarizes results from the *baseline*, *Gauss*, *Vanilla-PEM* and *Gauss-PEM* methods, all trained on the SNR range from 0dB to 50dB. Our network trained on clean speech only (clean-baseline) achieves 13.8% WER with a trigram language model and our 8.5M parameter network, while in literature [27], a 13.5% WER was achieved using a trigram language model and 3x larger (26.5M) parameter network. This confirms that our end-to-end speech recognition pipeline is fully functional.

Baseline: The noisy-baseline network starts with a 25% higher WER on the clean test set than our clean-baseline network. For SNRs lower than 25dB, the noisy-baseline is significantly more noise robust. The WER seems to drastically

TABLE III
ABSOLUTE WER [%] ON SINGLE SNRS AFTER DECODING. PRINTED BOLD: LOWEST WER.

Testing against pink noise																
Method / SNR[dB]	clean	50.0	45.0	40.0	35.0	30.0	25.0	20.0	15.0	10.0	5.0	0.0	-5.0	-10.0	-15.0	-20.0
Clean-baseline	13.8	14.4	14.0	13.8	13.7	13.7	16.1	18.9	25.9	40.1	61.8	86.4	109.0	133.4	147.2	152.8
Noisy-baseline	17.3	17.4	17.3	16.9	16.5	16.4	16.2	16.8	19.0	23.4	36.5	59.8	90.0	116.2	126.7	129.5
Gauss	15.7	15.8	15.7	15.6	14.8	14.4	14.5	15.3	16.9	20.2	28.9	45.5	72.8	94.9	99.0	98.9
Vanilla-PEM	13.3	13.2	13.2	12.7	12.6	12.6	12.9	13.6	15.1	18.9	26.2	45.0	73.5	93.4	96.9	97.1
Gauss-PEM	13.6	13.5	13.6	13.4	13.2	12.6	12.4	12.8	14.2	17.0	22.3	37.6	66.3	90.2	95.9	96.8
ACCAN	15.9	15.8	15.4	15.3	15.0	15.0	15.2	15.9	16.1	18.5	22.9	33.7	58.8	85.9	95.6	96.2
ACCAN reversed	14.6	14.4	14.3	14.1	13.3	13.4	13.9	14.4	15.4	18.5	24.4	40.2	67.8	90.9	97.0	97.2

Testing against babble noise																
Method / SNR[dB]	clean	50.0	45.0	40.0	35.0	30.0	25.0	20.0	15.0	10.0	5.0	0.0	-5.0	-10.0	-15.0	-20.0
Clean-baseline	13.8	14.2	14.2	13.9	14.2	14.5	15.7	18.8	26.6	43.9	74.2	102.2	116.6	122.4	122.3	121.4
Noisy-baseline	17.3	17.1	16.9	16.7	16.1	15.7	15.8	17.8	23.1	35.5	60.6	94.1	119.4	128.4	129.3	129.2
Gauss	15.7	15.6	15.7	15.7	15.3	15.0	15.4	16.5	19.5	27.5	45.9	77.4	102.6	109.0	109.8	110.6
Vanilla-PEM	13.3	13.2	12.9	12.7	12.3	12.7	12.8	14.0	17.4	25.6	44.2	72.7	93.2	99.1	99.5	99.8
Gauss-PEM	13.6	13.8	13.7	13.4	13.4	13.3	13.7	14.6	16.9	22.9	37.4	64.1	89.8	97.3	97.3	97.4
ACCAN	15.9	15.7	15.3	14.9	15.1	15.1	15.0	15.5	17.5	21.8	33.4	57.2	86.1	97.2	98.8	99.1
ACCAN reversed	14.6	14.4	14.2	14.0	14.1	14.0	14.0	14.6	16.5	21.9	35.5	63.5	88.7	96.6	97.6	97.6

increase at 25dB for the clean-baseline, while the noisy-baseline sees the increase at a lower 10dB SNR. However, all other methods outperform the noisy-baseline by a significant margin at high and low SNRs.

Vanilla-PEM vs Gauss: Compared to the noisy-baseline, Vanilla-PEM achieves a 23% decrease in WER on high SNR, while Gauss only reduces WER by 15% (both pink noise and babble noise). This results in vanilla-PEM being able to outperform the clean-baseline on clean speech, while Gauss is not able to do the same. On low SNR, both methods reduce WER by around 20% on the pink noise test set. On babble noise, PEM results in a higher 22.5% WER decrease compared to the 15.5% decrease provided by Gauss.

Gauss-PEM: The Gauss-PEM method achieves the overall lowest WER on the high and low SNR range. It beats the noisy-baseline method by between 26.5% and 28.7% on high SNR, on low SNR and on the ROI for both pink noise and babble noise. The results on the high SNR range are notable: Gauss-PEM is able to outperform the clean-baseline network at every single SNR step in the high SNR range, even on clean speech. The network is much more noise robust while at the same time it even improves clean speech scores. Around 35dB to 25dB, Gauss-PEM (other methods, too) reaches its minimum WER. This is expected, as the mean SNR of the training SNR range is 25dB and the network seems to optimize for SNR levels close to this value [16].

B. Curriculum learning

To further increase the noise robustness, we developed a curriculum learning strategy for the Gauss-PEM method, resulting in our novel ACCAN method. We compare our results to Gauss-PEM, as this was the most noise robust non-curriculum method. Our test results show increased noise robustness for ACCAN on pink noise and babble noise: the WER decreases between 3.3% (ROI, pink noise) and 4.1% (ROI babble noise). For pink noise, the biggest decrease is seen at 0dB (10.5% WER decrease) and -5dB (11.3% decrease). For

babble noise, the biggest WER decreases are found at 10dB (4.9%), 5dB (10.9%) and 0dB (10.7%).

The average WER of ACCAN on the high SNR range is worse on pink noise (relative 8.8% increase in WER), but better on babble noise (relative 0.4% decrease in WER). Ultimately, the absolute WER in clean speech of ACCAN (15.9%) is better than the noisy-baseline (17.3%) but worse than Gauss-PEM (13.6%).

V. DISCUSSION

All proposed training methods lead to networks with increased noise robustness in the low SNR range in comparison with the standard noisy-baseline method. The noise robustness is increased on a network level and it does not rely on complex preprocessing frameworks.

We see increased noise robustness against both tested noise types. This is remarkable, as the networks only saw the pink noise type during training. The results show that waveform-level noise mixing (PEM) is especially strong in transferring noise robustness to noise types not seen in training. The feature level noise addition (Gauss) is less effective on unseen noise types. Also, PEM enabled us to train noise robust networks that - at the same time - achieve lower WER on clean speech than a network trained only on clean speech. The uncompromising data augmentation by PEM should be a decisive factor to achieve these results. While the noisy-baseline trained on 1.7GB of unique data (waveform level), the PEM-enabled networks trained on up to 408GB (240 epochs for ACCAN * 1.7GB) of unique training data (waveform level). By permanently sampling different noise segments, we force the network to not rely only on constant noise features for classification, but to develop a better internal representation of the speech data. This representation could be refined further by using other noise types besides pink noise for training, such as babble noise, street noise, restaurant noise. Also, SNR steps smaller than 5dB could be used to allow more than 11 different SNR levels during training.

PEM enabled us to dynamically change the SNR during training. This facilitated the implementation of our novel ACCAN training strategy, that achieved the best noise robustness performance. The multi-stage training starts at low SNRs, where annealed networks are able to explore the parameters space with moderate influence of the speech signal. During gradual exposure to higher SNRs in the training process, accordion annealed networks refine their internal model of speech step by step, while they seemingly acquire higher noise robustness at the lower SNR levels. The inverse way of going through the SNR range does not yield increased noise robustness. The immediate presence of clean speech signals seems to force the network to converge faster to a complex acoustic model instead of exploring the parameter space.

VI. CONCLUSION

This work proposes new training methods for improving the noise robustness of RNNs for a LVCSR task. The networks are trained for a wide SNR range with the use of the Vanilla-PEM training method which adds noise at the waveform level and the Gauss method which injects Gaussian noise at the feature level. By combining the Gauss and Vanilla-PEM methods into the Gauss-PEM method, we achieve on average a 28% WER reduction on the 20dB to -10 dB SNR range when compared to a conventional multi-condition training method. At the same time, we achieve lower WER on clean speech than a network that is trained solely on clean speech. The ACCAN training strategy enhances the Gauss-PEM method with a curriculum learning strategy. The ACCAN training strategy results in performance up to 11.3% lower WER at low SNRs compared to Gauss-PEM method.

VII. ACKNOWLEDGEMENTS

The authors acknowledge Enea Ceolini and Joachim Ott for discussions on RNNs. They also acknowledge Ying Zhang and Yoshua Bengio (both from University of Montreal) for help in the setup of the language model. This work was partially supported by Samsung Advanced Institute of Technology and EU H2020 COCOHA #644732.

REFERENCES

- [1] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath *et al.*, "Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups," *IEEE Signal Processing Magazine*, vol. 29, no. 6, pp. 82–97, 2012.
- [2] J. Li, L. Deng, Y. Gong, and R. Haeb-Umbach, "An overview of noise-robust automatic speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 22, no. 4, pp. 745–777, 2014.
- [3] S. F. Boll, "Suppression of acoustic noise in speech using spectral subtraction," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 27, no. 2, pp. 113–120, 1979.
- [4] Y. Fujita, R. Takashima, T. Homma, R. Ikeshita, Y. Kawaguchi, T. Sumiyoshi, T. Endo, and M. Togami, "Unified asr system using lgm-based source separation, noise-robust feature extraction, and word hypothesis selection," in *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, 2015, pp. 416–422.
- [5] T. Hori, Z. Chen, H. Erdogan, J. R. Hershey, J. Roux, V. Mitra, and S. Watanabe, "The merl/sri system for the 3rd chime challenge using beamforming, robust feature extraction, and advanced speech recognition," in *Proc. IEEE ASRU*, 2015.
- [6] H. Hermansky, N. Morgan, A. Bayya, and P. Kohn, "Compensation for the effect of the communication channel in auditory-like analysis of speech (rasta-plp)," in *Second European Conference on Speech Communication and Technology*, 1991.
- [7] T. Hori, Z. Chen, H. Erdogan, J. R. Hershey, J. Roux, V. Mitra, and S. Watanabe, "The merl/sri system for the 3rd chime challenge using beamforming, robust feature extraction, and advanced speech recognition," in *Proc. IEEE ASRU*, 2015.
- [8] A. L. Maas, Q. V. Le, T. M. O’Neil, O. Vinyals, P. Nguyen, and A. Y. Ng, "Recurrent neural networks for noise reduction in robust asr," in *INTERSPEECH*, 2012, pp. 22–25.
- [9] X. Feng, Y. Zhang, and J. Glass, "Speech feature denoising and dereverberation via deep autoencoders for noisy reverberant speech recognition," in *2014 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. IEEE, 2014, pp. 1759–1763.
- [10] D. Palaz, R. Collobert *et al.*, "Analysis of cnn-based speech recognition system using raw speech as input," in *Proceedings of Interspeech*, 2015.
- [11] G. Trigeorgis, F. Ringeval, R. Brueckner, E. Marchi, M. A. Nicolaou, S. Zafeiriou *et al.*, "Adieu features? end-to-end speech emotion recognition using a deep convolutional recurrent network," in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 5200–5204.
- [12] M. L. Seltzer, D. Yu, and Y. Wang, "An investigation of deep neural networks for noise robust speech recognition," in *2013 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*. IEEE, 2013, pp. 7398–7402.
- [13] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [14] Y. Miao, M. Gowayed, and F. Metze, "Eesen: End-to-end speech recognition using deep rnn models and wfst-based decoding," in *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, 2015, pp. 167–174.
- [15] D. Amodei *et al.*, "Deep speech 2: End-to-end speech recognition in english and mandarin," *CoRR*, vol. abs/1512.02595, 2015. [Online]. Available: <http://arxiv.org/abs/1512.02595>
- [16] S. Yin, C. Liu, Z. Zhang, Y. Lin, D. Wang, J. Tejedor, T. F. Zheng, and Y. Li, "Noisy training for deep neural networks in speech recognition," *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2015, no. 1, pp. 1–14, 2015.
- [17] P. Koistinen and L. Holmström, "Kernel regression and backpropagation training with noise," in *1991 IEEE International Joint Conference on Neural Networks*, 1991, pp. 367–372.
- [18] K. J. Geras and C. A. Sutton, "Scheduled denoising autoencoders," *CoRR*, vol. abs/1406.3269, 2014. [Online]. Available: <http://arxiv.org/abs/1406.3269>
- [19] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proceedings of the 26th Annual International Conference on Machine Learning*, ser. ICML ’09. New York, NY, USA: ACM, 2009, pp. 41–48. [Online]. Available: <http://doi.acm.org/10.1145/1553374.1553380>
- [20] A. team, "Audacity," 2016. [Online]. Available: <http://www.audacityteam.org/>
- [21] A. Varga and H. J. Steeneken, "Assessment for automatic speech recognition: li. noise92: A database and an experiment to study the effect of additive noise on speech recognition systems," *Speech Communication*, vol. 12, no. 3, pp. 247–251, 1993.
- [22] A. Graves, S. Fernández, F. Gomez, and J. Schmidhuber, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proceedings of the 23rd International Conference on Machine Learning*. ACM, 2006, pp. 369–376.
- [23] L. team, "Lasagne: First release." Aug. 2015. [Online]. Available: <http://dx.doi.org/10.5281/zenodo.27878>
- [24] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [25] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *International Conference on Artificial Intelligence and Statistics*, 2010, pp. 249–256.
- [26] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2014. [Online]. Available: <http://arxiv.org/abs/1412.6980>
- [27] A. Graves and N. Jaitly, "Towards end-to-end speech recognition with recurrent neural networks," in *ICML*, vol. 14, 2014, pp. 1764–1772.