

Exploiting Attention-based Sequence-to-Sequence Architectures for Sound Event Localization

Christopher Schymura*, Tsubasa Ochiai†, Marc Delcroix†, Keisuke Kinoshita†, Tomohiro Nakatani†, Shoko Araki†, Dorothea Kolossa*

* Institute of Communication Acoustics
Ruhr University Bochum, Bochum, Germany

Email: christopher.schymura@rub.de

† NTT Communication Science Laboratories
NTT Corporation, Kyoto, Japan

Abstract—Sound event localization frameworks based on deep neural networks have shown increased robustness with respect to reverberation and noise in comparison to classical parametric approaches. In particular, recurrent architectures that incorporate temporal context into the estimation process seem to be well-suited for this task. This paper proposes a novel approach to sound event localization by utilizing an attention-based sequence-to-sequence model. These types of models have been successfully applied to problems in natural language processing and automatic speech recognition. In this work, a multi-channel audio signal is encoded to a latent representation, which is subsequently decoded to a sequence of estimated directions-of-arrival. Herein, attentions allow for capturing temporal dependencies in the audio signal by focusing on specific frames that are relevant for estimating the activity and direction-of-arrival of sound events at the current time-step. The framework is evaluated on three publicly available datasets for sound event localization. It yields superior localization performance compared to state-of-the-art methods in both anechoic and reverberant conditions.

Index Terms—sound event localization, recurrent neural network, sequence-to-sequence model

I. INTRODUCTION

Sound event localization (SEL) considers estimating the spatial positions of sound events from audio signals. The definition of a sound event refers to a sound produced by a variety of sources, e.g. speech or a musical instrument, which oftentimes is a time-varying quantity with respect to volume and position. Many applications rely on accurate SEL, as it is an important processing step for speech enhancement in hearing aids [1], acoustic monitoring for industrial applications [2], robotics [3] and many others. The importance of localization and tracking algorithms for audio signal processing has also been stressed during the recent LOCATA challenge [4], where a wide range of contributions tackled this problem in a diverse set of challenging scenarios, cf. [5] for details.

However, many classical frameworks for SEL are based on parametric approaches, e.g. the steered response power phase transform (SRP-PHAT) [6] or multiple signal classification (MUSIC) [7] methods. Furthermore, approaches using learning-based techniques like logistic regression [8] or independent component analysis [9] have also been proposed.

In addition, approaches utilizing deep neural networks for SEL have been developed recently. For instance, the method

described in [10] proposed a convolutional neural network (CNN) operating on the phase of the audio input in the short-time Fourier transform (STFT) domain to predict the direction-of-arrival (DoA) of a sound event. The work of He et al. [11] introduced a similar framework for robotics applications, which used features derived from SRP-PHAT as input to a CNN-based SEL system. This approach was further developed in [12], where the complex STFT spectrum was used as input to a CNN, followed by a recurrent neural network (RNN) based on gated recurrent units (GRUs). A spatial pseudo spectrum, similar to classical parametric approaches like MUSIC, was obtained from the recurrent output of the network. This representation was used to obtain DoA estimates for each sound event. Further extensions towards sound event localization and detection (SELD) were proposed in [13], which considered a joint estimation of sound type and spatial location. This system was recently adapted to tackle dynamic scenarios with moving sound sources [14]. Models that exploit temporal context for SEL have shown superior performance compared to conventional feed-forward networks [13]–[15].

This paper proposes a recurrent architecture for SEL in scenarios with multiple, potentially overlapping sound sources, via an attention-based sequence-to-sequence approach to handle temporal dependencies. Such models were initially proposed in the context of neural machine translation [16], but have also been utilized for other tasks like automatic speech recognition [17]. Compared to a conventional recurrent network architecture, the attention mechanism facilitates an improved capturing of the temporal structure in the input sequence. This is achieved by exclusively focusing on parts of the audio signal that are relevant for predicting the spatial source locations at the current frame. This also enables the model to account for long-term dependencies when localizing sound events, which is generally not possible with classical methods that rely on the first-order Markov assumption [18].

II. SYSTEM OVERVIEW

The proposed SEL model is based on a recurrent encoder-decoder architecture with attentions [16]. Therefore, it is termed *Attention-based Deep REcurrent Neural-network for locALizing souNd Events* (ADRENALINE) in this work, to

stress the underlying design for the desired application. The model essentially consists of three building-blocks: a CNN-based feature extraction, the encoder RNN and the decoder RNN with an attention mechanism, which are depicted in Fig. 1. These components will be described in the following.

A. Feature extraction

To extract input features from a spectral representation of acquired audio signals, the CNN-based feature extractor from SELDNet [13] was utilized in this work. The ADRENALINE architecture operates on fixed-length chunks of audio data: Given an acoustic waveform with C channels, it is divided into non-overlapping chunks of 500 ms duration. An STFT, using a 2048-point Hamming window with a frame length of 40 ms and 20 ms shift, is performed within each chunk. The complex spectrum is separated into its magnitude and phase components, which are concatenated to an input tensor of dimension $K \times L \times 2C$, where $K = 25$ is the number of frames within one chunk and $L = 1024$ is the number of non-redundant bins in the complex spectrum.

The input tensor is fed into a CNN with three layers. Each layer is composed of 64 filters with kernel size 3×3 , followed by batch normalization, rectified linear units (ReLUs) and max-pooling with kernel size 1×8 in the first two layers and 1×2 in the third layer. The output of the third layer is reshaped to yield $K \times D_y$ feature sequences, where $D_y = 512$ is the feature dimension. The elements of the feature sequence will be denoted as $\mathbf{y}_{k'} \in \mathbb{R}^{D_y}$, where $k' = 1, \dots, K$ represent the encoder time indices.

B. Encoder

A conventional RNN-based encoder architecture is used in the ADRENALINE framework. The input feature sequence is fed into a bidirectional RNN, which yields a hidden state representation $\mathbf{h}_{E,k'} \in \mathbb{R}^{2D_h}$ at each time step, according to

$$\mathbf{h}_{E,k'} = f_E(\mathbf{h}_{E,k'-1}, \mathbf{y}_{k'}). \quad (1)$$

Due to the bidirectional implementation of the RNN, the hidden state is composed of the states obtained during the forward and backward passes [19], which is not explicitly shown in Eq. (1), cf. [16] for details. A GRU with a hidden dimension of $D_h = 64$ is used as the underlying RNN cell. It is initialized with zeros for $k' = 0$.

C. Decoder and attention mechanism

As shown in Fig. 1, at each decoder time index k , the decoder takes as input a vector \mathbf{z}_k , which is composed of a decoder output vector \mathbf{x}_{k-1} from the previous time-step and a context vector \mathbf{c}_k , generated by the attention mechanism. The context vector is computed according to

$$\mathbf{c}_k = \sum_{k'=1}^K \alpha_{k'} \mathbf{h}_{E,k'}, \quad (2)$$

where α_k denote the attention weights. Each attention weight is obtained via a scaled dot-product [20] between the encoder

hidden state and the corresponding previous decoder hidden state, followed by a softmax activation function

$$\alpha_{k'} = \text{SOFTMAX}\left(\frac{1}{\sqrt{2D_h}} \mathbf{h}_{E,k'}^T \mathbf{h}_{D,k-1}\right). \quad (3)$$

The decoder output \mathbf{x}_k represents the current information about detected sound events and their locations. It is composed of three elements: a vector $\gamma_k \in \mathbb{R}^S$ indicating source activity and vectors of corresponding azimuth and elevation angles, $\phi_k \in \mathbb{R}^S$ and $\vartheta_k \in \mathbb{R}^S$, respectively. Each DoA corresponds to the source activity indicator in the same row of the source activity vector. Their dimension S indicates the maximum number of sources that the framework can handle simultaneously. It is set to $S = 4$ in this work. This representation was inspired by the SELDNet architecture [13], which utilized a similar output structure, but additionally performed sound event detection, which is omitted in the framework proposed here. The joint decoder output $\mathbf{x}_k \in \mathbb{R}^{3S}$ is obtained by stacking the individual output elements into a single vector. Hence, by utilizing the result from Eq. (2), the decoder input vector can be represented as $\mathbf{z}_k = [\mathbf{c}_k \quad \mathbf{x}_{k-1}]^T \in \mathbb{R}^{2D_h+3S}$.

The decoding step for estimating source activities and locations can now be expressed via the decoder RNN

$$\mathbf{h}_{D,k} = f_D(\mathbf{h}_{D,k-1}, \mathbf{z}_k), \quad (4)$$

whose outputs are $\gamma_k = \sigma(\mathbf{W}^\gamma \mathbf{h}_{D,k})$, $\phi_k = \mathbf{W}^\phi \mathbf{h}_{D,k}$ and $\vartheta_k = \mathbf{W}^\vartheta \mathbf{h}_{D,k}$, with weights \mathbf{W}^γ , \mathbf{W}^ϕ and \mathbf{W}^ϑ , where $\sigma(\cdot)$ denotes the sigmoid nonlinearity.

D. Loss function

The decoder output representation introduced in the previous section requires a form of loss that is tailored to the SEL task. Therefore, a loss function with two components

$$\mathcal{L}_{\text{SEL}} = \mathcal{L}_{\text{ACT}}(\hat{\gamma}, \gamma) + \lambda \mathcal{L}_{\text{DOA}}(\hat{\phi}, \hat{\vartheta}, \phi, \vartheta, \gamma) \quad (5)$$

is proposed, where $\lambda \in \mathbb{R}_+$ is a scaling factor. The discrete decoder time index k is omitted here for notational convenience. Eq. (5) incorporates a binary cross-entropy loss term $\mathcal{L}_{\text{ACT}}(\hat{\gamma}, \gamma)$ [21, Chap. 5] evaluating the difference between estimated and ground-truth source activities $\hat{\gamma}$ and γ , respectively. The second term is based on the DoA error [12]

$$\xi = \arccos\left(\sin(\hat{\phi}) \sin(\phi) + \cos(\hat{\phi}) \cos(\phi) \cos(\vartheta - \hat{\vartheta})\right), \quad (6)$$

which measures the angle between the estimated azimuth $\hat{\phi}$ and elevation $\hat{\vartheta}$ and the ground-truth DoA, given by ϕ and ϑ . The trigonometric functions in Eq. (6) are applied element-wise, so that the resulting vector ξ contains a single DoA error for each potential source. The DoA error showed slightly better performance when dealing with angular values, compared to the more commonly used mean squared error (MSE) loss.

As only active sources shall be considered in the DoA loss, the DoA errors of non-active sources are omitted by incorporating the ground-truth source activity vector as

$$\mathcal{L}_{\text{DOA}}(\hat{\phi}, \hat{\vartheta}, \phi, \vartheta, \gamma) = \frac{1}{S} \xi^T \gamma, \quad (7)$$

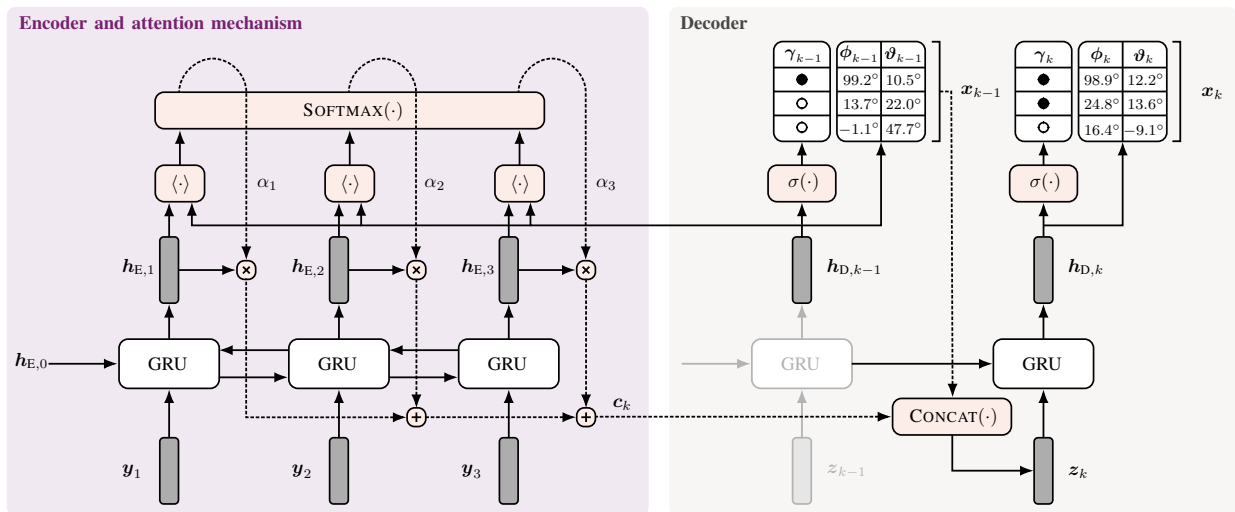


Fig. 1. General overview of the ADRENALINE model architecture. The box on the left shows an exemplary encoding process for three discrete time-steps from $k' = 1, \dots, 3$, where $\mathbf{y}_{k'}$ denotes the observed input features and $\mathbf{h}_{E,k'}$ corresponds to the hidden state of the encoder GRUs. The attention weights $\alpha_{k'}$ are computed via scaled dot products between the encoder hidden states and the corresponding decoder hidden state $\mathbf{h}_{D,k-1}$ from the previous decoding time-step. A context vector \mathbf{c}_k is derived as a weighted sum of the encoder hidden states, using the attention weights. The output of the decoder \mathbf{x}_k is composed of the source activity indicator γ_k and the corresponding source DoAs, comprising azimuth ϕ_k and elevation ϑ_k . A concatenation of the decoder output from the previous time-step and the current context vector serves as input to the decoder, as shown in the box on the right.

which effectively blocks the gradient flow for non-active sources during back-propagation.

It should be noted that the DoA loss presented in Eq. (7) does not take potentially occurring source permutations between consecutive time-steps into account. Therefore, the DoA error given in Eq. (6) is computed for all possible source permutations during the calculation of the loss. This is computationally feasible, as the maximum number of sources set to $S = 4$ is small and all computations can be run in parallel. The lowest resulting DoA error is then selected for evaluating Eq. (7). Similar approaches have also been used in other application domains, e.g. for speech separation in [22].

III. EVALUATION

The proposed ADRENALINE framework is evaluated on three publicly available datasets and compared to related SEL baseline methods. The program code is accessible online¹ and the detailed training procedure and experimental setup is described in the following.

A. Datasets

The ANSYN, RESYN and REAL datasets introduced in [13] were selected as suitable evaluation corpora in this study. These datasets provide a number of first order Ambisonic format recordings that is large enough to train deep learning models for SEL.

The first two were synthesized using simulated anechoic and reverberant impulse responses, whereas the latter utilized impulse responses recorded in real environments. All datasets are assembled with a similar structure: they comprise three subsets each, corresponding to either no, at most two, or

at most three temporally overlapping sources. Each subset provides three cross-validation splits with 300 audio files each. These files are further divided into 240 files for training and 60 for validation. All audio files are sampled with 44.1 kHz, have a duration of 30 s and are encoded in the Ambisonics format. The audio signals in the synthetic datasets were created using sound events from 11 different classes, covering the full azimuth range and an elevation range from -60° to 60° . For the recorded impulse response dataset, sound events from 8 classes were used for synthesizing the audio signals. Herein, again the full azimuth range was covered and the elevation range was restricted between -40° and 40° .

B. Baseline methods

Two baseline methods were employed in this work to compare the performance of the proposed model to similar SEL frameworks. The first baseline model builds upon the CNN feature extraction stage described in Sec. II-A. It simply adds a fully connected layer, which outputs source activity and DoA vectors. This is a feed-forward model with no recurrent elements and will be referred to as CNN in the experiments.

Additionally, a variant of the SELDNet architecture [13] is utilized as a second baseline. As the name suggests, it was initially proposed for SELD tasks, predicting location and class of a sound event, where it achieved state-of-the-art performance. It shares the same CNN-based feature extraction stage with the proposed ADRENALINE framework and is also based on bidirectional GRUs. The main difference between the methods is that SELDNet does not incorporate any attention mechanisms. Instead, it directly utilizes the output of the recurrent layers for estimation.

To enable a fair comparison to the proposed ADRENALINE framework, the SELDNet model is modified here by removing

¹<https://github.com/rub-ksv/adrenaline>

TABLE I
FRAME RECALL IN PERCENT GROUPED BY DATASET. THE METRIC IS
AVERAGED OVER ALL SUBSETS AND CROSS-VALIDATION FOLDS.

	ANSYN	RESYN	REAL
CNN	87.48	71.91	72.07
SELDNet(m)	85.78	72.46	69.63
ADRENALINE	84.83	71.18	72.08

information about the sound class from the output layer. This essentially yields a system output identical to the one proposed in Sec. II-C. The model is referred to as SELDNet(m) in the following, denoting the modified output.

C. Performance metrics

Frame recall and DoA error are used as metrics for evaluating SEL performance. The frame recall metric [13] describes the percentage of frames, where the estimated number of active sources matches the ground-truth. During evaluation, a source is considered active if the corresponding element in the source activity vector exceeds a threshold of 0.5. The DoA error is computed as described in Eq. (6), where only active sources are considered. The Hungarian algorithm [23] is used to solve the assignment problem between multiple estimated DoAs and the ground-truth, which yields the smallest possible DoA error.

D. Experimental setup

All models were implemented in PyTorch [24]. The baseline model parameterization was chosen as reported in [13]. The parameters of the proposed ADRENALINE model were specified as described in Sec. II. The SEL loss function given in Eq. (5) was used with a scaling factor of $\lambda = 1$ for optimizing all evaluated models. All model parameters were initialized using the Kaiming initialization method from [25]. Training was conducted using the AdamW optimizer [26] with batch-size 256 and a learning rate of 0.0002. The learning rate was varied by adopting the scheduling scheme proposed in [20], which resulted in an increased training stability. Separate models were trained for each cross-validation fold, utilizing early stopping with a maximum number of 200 epochs.

IV. RESULTS AND DISCUSSION

The experimental results are summarized in Tab. I and Fig. 2. The Mann-Whitney- U test [27] was used to show statistically significant differences in DoA error between the proposed ADRENALINE framework and the baselines. This nonparametric test was chosen, because the application of a Shapiro-Wilk test [28] on the resulting DoA errors indicated that the individual samples are not normally distributed.

The resulting DoA errors indicate that the proposed ADRENALINE framework outperforms both baseline methods on all three datasets, where the performance differences are statistically significant in most cases. This shows that employing a sequence-to-sequence architecture for SEL yields improved localization performance compared to standard recurrent architectures and feed-forward networks in anechoic

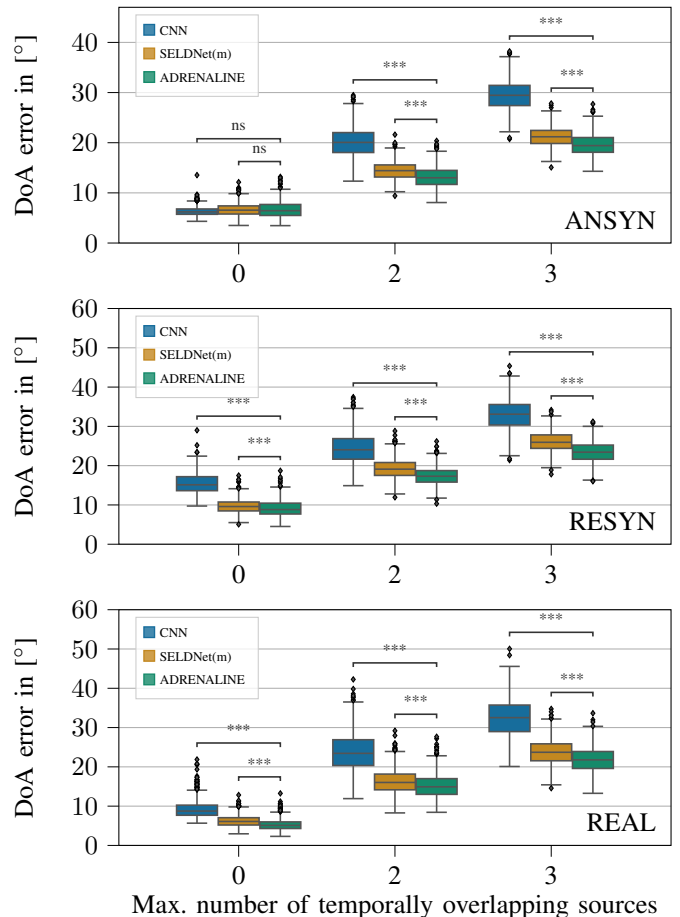


Fig. 2. DoA errors obtained in all cross-validation folds for all three datasets, shown as grouped box-plots. Asterisks indicate a statistically significant difference, where *** denotes $p < 0.001$.

as well as reverberant environments. The performance benefit of the ADRENALINE framework is especially prominent in scenarios with multiple overlapping sources.

Incorporating temporal information is crucial for effectively performing SEL. Even though conventional GRU or long-short term memory (LSTM) architectures can capture long-term temporal dependencies, the attention-based framework utilized here seems to provide a better utilization of temporal information for predicting source activity and DoAs. A possible reason for this is the ability of attention-based sequence-to-sequence models to incorporate multiple relevant time steps into the estimation, without taking into account their specific temporal order. This is not possible for (bidirectional) recurrent architectures, as they can only operate on the currently accumulated information in the forward and backward passes. Fig. 3 depicts exemplary attention maps for all three evaluation corpora, which show that the attention weights have a larger spread in reverberant environments. This seems reasonable, as DoA estimation under increased reverberation requires longer temporal contexts.

All methods yield comparable performance for estimating the correct number of sources. This is reflected by the frame




Fig. 3. Attention maps produced by the ADRENALINE model for exemplary chunks from all three evaluation datasets. The index k' denotes the encoder time index, whereas k is the decoder time index.

recall metric shown in Tab. I. The frame recall seems to degrade significantly in reverberant environments. This indicates that reverberation does not only provide a challenge for localization, but also hinders all investigated models in distinguishing between concurrent sound events.

V. CONCLUSION AND OUTLOOK

This paper presents an attention-based sequence-to-sequence architecture for sound event localization. Compared to previous models based on standard recurrent neural network architectures, the proposed model better utilizes temporal information for estimating the activity and direction-of-arrival of sound events. An experimental evaluation on publicly available datasets for sound event localization supports the initial hypothesis that exploiting attentions for temporal focusing aids estimating the location of sound events. In addition, the proposed framework still provides many interesting opportunities to develop this approach further. By acknowledging the latest progress in natural language processing, a transformer-based architecture might also be applicable as an alternative to sequence-to-sequence models for sound event localization. Additionally, the incorporation of probabilistic estimates through, e.g., differentiable Bayesian techniques might further broaden the application range of the proposed framework.

REFERENCES

- [1] R. Li, D. Pan, and S. Zhang, "Speech enhancement algorithm based on sound source localization and scene matching for binaural digital hearing aids," *Journal of Medical and Biological Engineering*, vol. 39, pp. 403–417, Apr 2018.
- [2] C. J. Grobler, C. P. Kruger, B. J. Silva, and G. P. Hancke, "Sound based localization and identification in industrial environments," in *43rd Annual Conference of the IEEE Industrial Electronics Society*, Oct 2017, pp. 6119–6124.
- [3] C. Evers and P. A. Naylor, "Acoustic SLAM," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 9, pp. 1484–1498, Sept 2018.
- [4] H. W. Löllmann, C. Evers, A. Schmidt, H. Mellmann, H. Barfuss, P. A. Naylor, and W. Kellermann, "The LOCATA challenge data corpus for acoustic source localization and tracking," in *2018 IEEE 10th Sensor Array and Multichannel Signal Processing Workshop (SAM)*, July 2018, pp. 410–414.
- [5] C. Evers, H. Löllmann, H. Mellmann, A. Schmidt, H. Barfuss, P. Naylor, and W. Kellermann, "The LOCATA challenge: Acoustic source localization and tracking," 2019.
- [6] J. H. DiBiase, H. F. Silverman, and M. S. Brandstein, *Robust Localization in Reverberant Rooms*, pp. 157–180, Springer, Berlin, Heidelberg, 2001.

- [7] R. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Transactions on Antennas and Propagation*, vol. 34, no. 3, pp. 276–280, March 1986.
- [8] I. Trowitzsch, C. Schymura, D. Kolossa, and K. Obermayer, "Joining sound event detection and localization through spatial segregation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 487–502, 2020.
- [9] H. Sawada, R. Mukai, S. Araki, and S. Makino, "Multiple source localization using independent component analysis," in *2005 IEEE Antennas and Propagation Society International Symposium*, July 2005, vol. 4B, pp. 81–84 vol. 4B.
- [10] S. Chakrabarty and E. A. P. Habets, "Broadband DOA estimation using convolutional neural networks trained with noise signals," in *2017 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, Oct 2017, pp. 136–140.
- [11] W. He, P. Motlicek, and J. Odobez, "Deep neural networks for multiple speaker detection and localization," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, May 2018, pp. 74–79.
- [12] S. Adavanne, A. Politis, and T. Virtanen, "Direction of arrival estimation for multiple sound sources using convolutional recurrent neural network," in *2018 26th European Signal Processing Conference (EUSIPCO)*, Sep. 2018, pp. 1462–1466.
- [13] S. Adavanne, A. Politis, J. Nikunen, and T. Virtanen, "Sound event localization and detection of overlapping sources using convolutional recurrent neural networks," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 1, pp. 34–48, March 2019.
- [14] S. Adavanne, A. Politis, and T. Virtanen, "Localization, detection and tracking of multiple moving sound sources with a convolutional recurrent neural network," 2019.
- [15] Z. Wang, J. Li, and Y. Yan, "Target speaker localization based on the complex Watson mixture model and time-frequency selection neural network," *Applied Sciences*, vol. 8, no. 11, pp. 2326, Nov 2018.
- [16] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," in *3rd International Conference on Learning Representations, ICLR 2015*, 2015.
- [17] Z. Tüske, K. Audhkhasi, and G. Saon, "Advancing sequence-to-sequence based speech recognition," in *INTERSPEECH 2019*, 2019.
- [18] S. Chakrabarty, K. Kowalczyk, M. Taseska, and E. A. P. Habets, "Extended Kalman filter with probabilistic data association for multiple non-concurrent speaker localization in reverberant environments," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2014, pp. 7445–7449.
- [19] M. Schuster and K. K. Paliwal, "Bidirectional recurrent neural networks," *IEEE Transactions on Signal Processing*, vol. 45, no. 11, pp. 2673–2681, Nov 1997.
- [20] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems 30*, pp. 5998–6008. Curran Associates, Inc., 2017.
- [21] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*, MIT Press, 2016, <http://www.deeplearningbook.org>.
- [22] D. Yu, M. Kolbæk, Z. Tan, and J. Jensen, "Permutation invariant training of deep models for speaker-independent multi-talker speech separation," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 241–245.
- [23] H. W. Kuhn and B. Yaw, "The Hungarian method for the assignment problem," *Naval Res. Logist. Quart.*, pp. 83–97, 1955.
- [24] A. Paszke et al., "PyTorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds., pp. 8024–8035. Curran Associates, Inc., 2019.
- [25] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," in *2015 IEEE International Conference on Computer Vision (ICCV)*, Dec 2015, pp. 1026–1034.
- [26] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *International Conference on Learning Representations*, 2019.
- [27] H. B. Mann and D. R. Whitney, "On a test of whether one of two random variables is stochastically larger than the other," *Annals of Mathematical Statistics*, vol. 18, no. 1, pp. 50–60, 1947.
- [28] S. S. Shapiro and M. B. Wilk, "An analysis of variance test for normality (complete samples)," *Biometrika*, vol. 52, no. 3/4, pp. 591–611, Dec. 1965.