

# BAST: Binaural Audio Spectrogram Transformer for Binaural Sound Localization

Sheng Kuang<sup>b</sup>, Kiki van der Heijden<sup>c</sup>, Siamak Mehrkanoon<sup>a,b,\*</sup>

<sup>a</sup>Information and Computing Sciences, Utrecht University, Utrecht, The Netherlands

<sup>b</sup>Department of Data Science and Knowledge Engineering, Maastricht University, The Netherlands.

<sup>c</sup>Donders Institute for Brain Cognition and Behavior, Radboud University, Nijmegen, The Netherlands

---

## Abstract

Accurate sound localization in a reverberation environment is essential for human auditory perception. Recently, Convolutional Neural Networks (CNNs) have been utilized to model the binaural human auditory pathway. However, CNN shows barriers in capturing the global acoustic features. To address this issue, we propose a novel end-to-end Binaural Audio Spectrogram Transformer (BAST) model to predict the sound azimuth in both anechoic and reverberation environments. Two modes of implementation, i.e. BAST-SP and BAST-NSP corresponding to BAST model with shared and non-shared parameters respectively, are explored. Our model with subtraction interaural integration and hybrid loss achieves an angular distance of 1.29 degrees and a Mean Square Error of  $1e-3$  at all azimuths, significantly surpassing CNN based model. The exploratory analysis of the BAST's performance on the left-right hemifields and anechoic and reverberation environments shows its generalization ability as well as the feasibility of binaural Transformers in sound localization. Furthermore, the analysis of the attention maps is provided to give additional insights on the interpretation of the localization process in a natural reverberant environment.

**Keywords:** Transformer, Sound localization, Binaural integration

---

## 1. Introduction

Sound source localization is a fundamental ability in everyday life. Accurately and precisely localizing incoming auditory streams is required for auditory perception and social communication. In the past decades, the biological basis and neural mechanism of sound localization have been extensively explored [1, 2, 3, 4]. Normal hearing listeners extract spatial acoustic cues by mainly relying on the interaural level differences (ILD) and interaural time differences (ITD) of the auditory input. These cues are encoded through the human auditory subcortical pathway, in which the olivary complex and associated physiological structures integrate and convey the binaural signals from cochleas to the auditory cortex [2, 4]. However, sound localization is frequently affected by the noise and reverberations in the complex real-word environment, such as at a cocktail party. These noisy signals may undercut or distort the spatial cues [5]. Although many studies have focused largely on speech extraction from background noise, it is still not clear how the spatial position of acoustic signals in complex listening environments is extracted by the human brain.

Recently, Deep Learning (DL) [6] has been proposed to model auditory processing and has achieved great success.

These approaches enable optimization of auditory models for real-life auditory environment [7, 8, 9, 10]. In the early attempts, DL methods were combined with conventional feature engineering to deal with the noise and reverberation problems [7, 11]. For instance, in [7], binaural spectral and spatial features were separately extracted, providing complementary information for a two-layer Deep Neural Network (DNN). Similarly, in [8, 12, 13], deep neural networks were used to de-noise and de-reverberate complex sound stimuli. In a CNN-based azimuth estimation approach, researchers utilized a Cascade of Asymmetric Resonators with Fast-Acting Compression to analyze sound signals and used onsite-generated correlograms to eliminate the echo interference [14]. However, most of these approaches highly depend on feature selection. To reduce this constraint, end-to-end Deep Residual Network (DRN) was recommended [15, 16]. Instead of selecting features from acoustic signal, raw spectrograms of sound was utilized in Deep Residual Network for azimuth prediction [16]. DRN was shown to be robust even in the presence of unknown noise interference at the low signal-to-noise ratio. Subsequently, [9] proposed a pure attention-based Audio Spectrogram Transformer (AST) and achieved the state-of-the-art results for audio classification on multiple datasets. Although these DL-based methods have yielded promising results, however due to a lack of similar architectures to the human binaural auditory pathway, they may not resemble the neural processing of sound localization.

To encode the neural mechanisms underlying sound localization, the performance of deep learning methods is commonly

---

\*Corresponding author

Email addresses:

s.kuangsheng@student.maastrichtuniversity.nl (Sheng Kuang),

kiki.vanderheijden@donders.ru.nl (Kiki van der Heijden),

siamak.mehrkanoon@maastrichtuniversity.nl; s.mehrkanoon@uu.nl

(Siamak Mehrkanoon)

compared to the human behavioral characteristics [17, 18, 19]. For instance, [19] systematically explored the performance of binaural sound clips localization of a CNN in a real-life listening environment, however, its architecture does not resemble the structure of human auditory pathway. This issue has been addressed by utilizing a hierarchical neurobiological-inspired CNN (NI-CNN) to model the binaural characteristics of human spatial hearing [17]. This unique hierarchical design, models the binaural signal integration process and is shown to have brain-like latent feature representations. However, NI-CNN [17] is not an end-to-end model as it leverages a cochlear method to generate auditory nerve representations as model input. Furthermore, considering the wide range of frequencies of sound input, the convolution operations in NI-CNN mainly extract local-scale features and therefore may exhibit limitations for extracting global features in the acoustic time-frequency spectrogram.

In this study, we build on the success and barriers of previously proposed deep neural networks at localizing sound sources to further develop an end-to-end transformer based model, which captures the global acoustic features from auditory spectrograms. We aimed at (i) investigating the performance of a pure Transformer-based hierarchical binaural neural network for addressing real-life sound localization; (ii) exploring the effect of various loss functions and binaural integration methods on the localization acuity at different azimuths; (iii) visualizing the attention flow of the proposed model to demonstrate the localization process.

## 2. Related Works

Binaural auditory models utilize head-related binaural sound waves to accomplish specific acoustic tasks. Gaussian Mixture Model (GMM), commonly used in machine learning-based studies, calculates the probability distribution of the source location in reverberant environments [20]. Subsequently, model-based methods have been encouraged to extract ILD and ITD cues for DNN training [7]. However, the performance of these hybrid techniques remains unstable since the feature extraction routine varies across different datasets. NI-CNN, as a data-driven DL method, is able to learn latent features for azimuth prediction from human auditory nerve representations [17]. However, it still requires a computational neuroscience-based cochlear model to generate these representations. In this paper, we design an end-to-end model to automatically learn features from spectrogram without using model-based methods.

Transformer was originally proposed in natural language processing to handle long-range dependencies [21, 22]. Recently, Transformer was successfully applied in the computer vision field by casting images into patch embedding sequences [23, 24]. Many hybrid models combined the Transformer with a CNN or Recurrent Neural Network (RNN) in audio processing and some studies even directly embedded attention blocks into CNN or RNN to capture global features in a parameter-efficient way [25, 26, 27].

Subsequently, the authors in [9] introduced AST which uses Transformer model and variable-length monaural spectrograms

to perform sound classification tasks. As a convolution-free, pure attention-based model, AST uses the overlapped-patch embedding generation policy to convert the intra-patch local features to inter-patch attention weights. AST has achieved the state-of-the-art results on multiple datasets [28, 29] for audio classification tasks.

## 3. Method

### 3.1. Model architecture

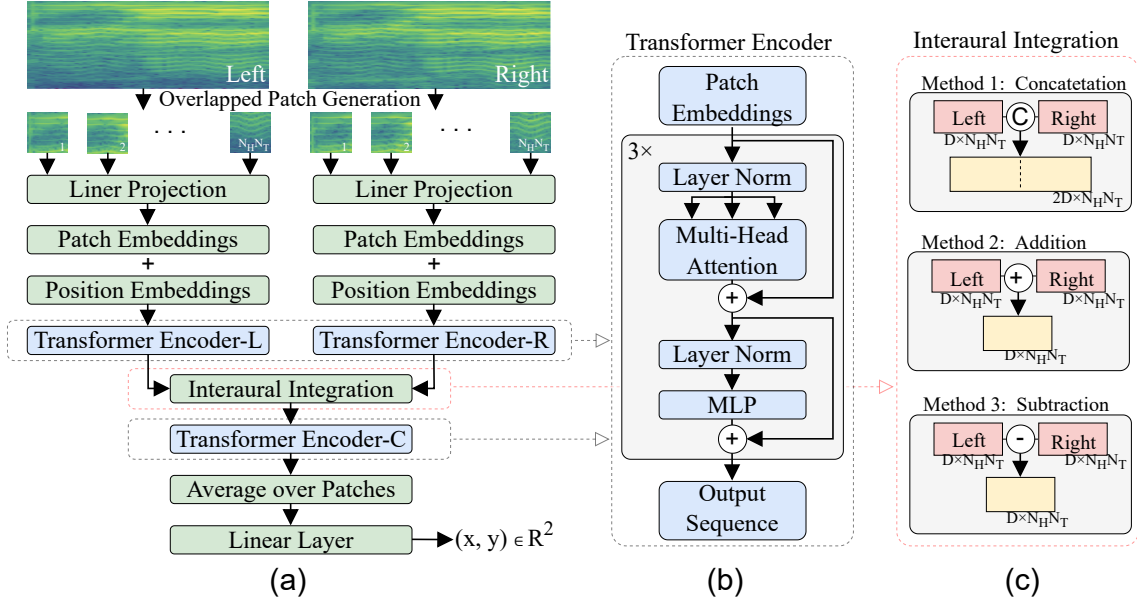
The proposed Binaural Audio Spectrogram Transformer (BAST), is illustrated in Fig. 1. Similar to NI-CNN [17], a dual-input hierarchical architecture is utilized to simulate the human subcortical auditory pathway. As opposed to NI-CNN which uses convolution layers, here three Transformer Encoders (i.e., left, right and center), hereafter called TE-L, TE-R and TE-C are utilized to construct a pure attention-based model. In particular, the pre-processed left and right sound waves are converted to left and right spectrograms denoted by  $x^L \in \mathbb{R}^{H \times T}$  and  $x^R \in \mathbb{R}^{H \times T}$ . Here,  $H$  indicates the frequency band and  $T$  indicates the number of Tukey windows (with shape parameter: 0.25).

In what follows, the TE-L path to process the input data is explained. The other path, i.e TE-R, follows the same process. At the beginning of patch embedding layer, the left spectrogram  $x^L \in \mathbb{R}^{H \times T}$  is first decomposed into an overlapped-patch sequence  $x_{patch}^L \in \mathbb{R}^{P^2 \times (N_H N_T)}$ , where  $P$  is the patch size,  $N_H$  and  $N_T$  are the number of patches in height and width respectively obtained as follows,

$$N_H = \left\lceil \frac{H - P + S}{S} \right\rceil, N_T = \left\lceil \frac{T - P + S}{S} \right\rceil. \quad (1)$$

In case  $H - P$  and  $T - P$  are not divisible by the stride  $S$  between patches, the spectrogram is zero-padded on the top and right respectively. A trainable linear projection is added to flatten each patch to a  $D$  dimensional latent representation, hereafter called patch embeddings [23]. Since our model outputs the sound location coordinates, the classification token in the Transformer encoder is removed. A fixed absolute position embedding [23] is added to the patch embeddings to capture the position information of the spectrogram in the Transformer. Here, the learnable position embedding is not used as it did not significantly change model performance compared to absolute position embedding[9]. The output of the left position embedding layer  $z_{in}^L \in \mathbb{R}^{D \times (N_H N_T)}$  is then fed to the Transformer encoder TE-L.

We use the identical Transformer encoder design in [23, 9], consisting of  $K$  stacked Multi-head Self-Attention (MSA) and Multi-Layer Perceptron (MLP) blocks. The BAST model performance is compared when using shared and non-shared parameters across the left and right Transformer encoders. Hereafter, BAST-SP refers to BAST model whose left and right Transformer encoders share parameters whereas in BAST-NSP the parameters of the left and right Transformer encoders are not shared. The output of left and right Transformer encoder, i.e.,  $z_{out}^L$  and  $z_{out}^R$ , represents neural signals underlying the initial



**Figure 1:** Architecture of the proposed Binaural Audio Spectrogram Transformer (BAST). (a) The architecture of the proposed model. (Here there are  $N_H N_T$  number of patches). (b) The architecture of a single Transformer encoder. (c) Three examined interaural integration methods: concatenation, addition and subtraction.

auditory processing stage along the left and right auditory pathway respectively. Subsequently, these binaural feature maps are integrated to simulate the function of the human olivary nucleus. Similar to NI-CNN, here three integration methods, i.e. addition, subtraction and concatenation, are investigated. Specifically, addition is the summation of feature maps of both sides; subtraction represents left feature map subtracted from the right feature map; concatenation is implemented by concatenating  $z_{out}^L$  and  $z_{out}^R$  along the the first dimension to produce  $z_{in}^C \in \mathbb{R}^{2D \times (N_H N_T)}$ . TE-C receives the integrated feature map  $z_{in}^C$  and output sequence  $z_{out}^C$ . Next, an average operation of the patch dimension and a linear transformer layer is applied to finally produce the sound location coordinates  $(x, y)$ . The last linear layer does not have any activation function, therefore the estimated coordinates can be any point on the 2D plane.

### 3.2. Loss Function

Three loss functions, i.e., Angular Distance (AD) loss [30], Mean Square Error (MSE) loss, as well as hybrid loss with a convex combination of AD and MSE, are explored in training the proposed model. Let  $C_i = (x_i, y_i)$  and  $\hat{C}_i = (\hat{x}_i, \hat{y}_i)$  denote the ground truth and predication coordinates for the  $i$ -th sample. MSE loss measures the squared difference of Euclidean distance between the prediction and the ground truth as follows:

$$\text{MSE} = \frac{1}{N} \sum_i^N \|C_i - \hat{C}_i\|_2^2, \quad (2)$$

where  $(\hat{x}_i, \hat{y}_i)$  is the predicted coordinates,  $(x_i, y_i)$  is the true sound location and  $N$  is the batch size. Note that MSE loss is able to penalize the large Euclidean distance error but is insensitive to the angular distance, which means that the azimuth may differ with the same MSE. In contrast to MSE, AD loss

merely measures the angular distance while ignoring the Euclidean distance:

$$\text{AD} = \frac{1}{\pi N} \sum_i^N \arccos \left( \frac{C_i \hat{C}_i^T}{\|C_i\|_2 \|\hat{C}_i\|_2} \right), \quad (3)$$

where  $C_i, \hat{C}_i \neq 0$ .

## 4. Experiments

### 4.1. Dataset

We use the binaural audio data in [17], which consists of a training dataset and an independent testing dataset. In the training dataset, 4600 real-life sound waves (duration:500 ms, sampling rate: 16000) are placed in 36 azimuth positions, respectively, with  $10^\circ$  azimuth resolution,  $0^\circ$  elevation, and 1-meter distance from the center point. In addition, sound waves are spatialized with two acoustic environments, i.e. an anechoic environment (AE) without reverberation and a  $10\text{m} \times 14\text{m}$  lecture hall with reverberation (RV). In particular, here the training and test sets contain data from both AE and RV environments. In total, the training dataset has 331200 binaural learning samples. Similarly, the independent testing dataset contains 400 new sound waves processed with the same method as described above, producing 28800 testing samples.

### 4.2. Model Evaluation

Models were evaluated by means of MSE and AD errors defined in Eq. (2) and (3). The lower AD and MSE errors the better localization performance is. Note that the MSE metric has no meaning when BAST is trained by AD loss because this loss does not optimize the Euclidean distance between the ground truth and the prediction, and that BAST has no constraints on the numerical range of the predicted coordinates.

Table 1: The performance of BAST-NSP and BAST-SP compared to NI-CNN and NI-CNN\* when different loss and binaural integration methods are used. The best performance model in AD and MSE are in bold. The ↓ indicates the lower the value of the metric, the better the model performance.

Model	Loss	Angular Distance(AD) ↓			Mean Squared Error (MSE) ↓		
		Concatenation	Addition	Subtraction	Concatenation	Addition	Subtraction
NI-CNN[17]	MSE	4.80°	4.80°	5.30°	0.011	0.013	0.014
	AD	3.70°	3.90°	5.20°	—	—	—
NI-CNN[17]*	MSE	8.92°	3.51°	3.67°	0.077	0.032	0.038
	AD	7.85°	1.97°	1.85°	—	—	—
	Hybrid	8.35°	3.53°	3.19°	0.074	0.033	0.031
BAST-NSP	MSE	2.78°	2.48°	2.42°	0.003	0.002	0.002
	AD	2.39°	1.30°	1.63°	—	—	—
	Hybrid	2.76°	1.83°	<b>1.29°</b>	0.004	0.002	<b>0.001</b>
BAST-SP	MSE	2.02°	4.97°	1.94°	0.002	0.018	0.002
	AD	2.66°	13.87°	1.43°	—	—	—
	Hybrid	1.98°	5.72°	2.03°	0.003	0.026	0.002

NI-CNN\* uses spectrogram as model input.

Table 2: The number of layers in each Transform encoder as well as the total number of trainable parameters of the proposed models. Tuple (·, ·, ·) indicates the number of layers in the left, right and center Transformer encoder respectively.

Model	Interaural Integration	Transformer Layers	Trainable Parameters
BAST-NSP	Concatenation	(3, 3, 3)	~76M
	Addition/ Subtraction	(3, 3, 3)	~57M
BAST-SP	Concatenation	(3, 3, 3)	~57M
	Addition/ Subtraction	(3, 3, 3)	~38M

### 4.3. Training Settings

As mentioned in 3.1, each sound wave are transformed to binaural spectrogram (size:  $2 \times 129 \times 61$ , frequency range: 0-8000Hz, window length: 128ms, overlap: 64ms) before training. In order to have balanced training samples, we randomly select 75% binaural spectrograms in each azimuth position and listening environments of training dataset. The remaining data is used as validation set. As stated before in the Dataset section, a separate test dataset is available for this study. This setting results in 248400 training samples and 82800 validation samples. Here, Adam optimizer [31] is used to train the model for 50 epochs with a batch size of 48 and a fixed learning rate of  $1e-4$ . In patch embedding layers, the stride of patches is set to 6, yielding 180 patches per spectrogram. Each Transformer encoder has three layers, with 1024 hidden dimensions (2048 dimensions when using concatenation as integration method in the last Transformer encoder TE-C), 16 attention heads in MSA blocks, 1024 dimensions in MLP blocks, 0.2 dropout rate in patch embeddings and MLP blocks. Our model implementation<sup>1</sup> is based on Python 3.8 and Pytorch 1.9.0, and are trained

<sup>1</sup>code available at <https://github.com/ShengKuangCN/BAST>



Figure 2: The angular distance (AD) error of the proposed BAST in each azimuth with different loss functions and interaural integration methods.

from scratch on  $2 \times$  NVIDIA GeForce GTX 1080Ti GPUs with 11GB of RAM. The empirically found and used number of layers in each Transformer encoder as well as the total number of trainable parameters are presented in Table 2.

## 5. Results

### 5.1. Overall Performance

The performance of the proposed BAST-NSP and BAST-SP models are compared with those of NI-CNN model. In particular, for NI-CNN model two modes of implementations corresponding to cochleogram and spectrogram as model inputs have been considered and are denoted by NI-CNN and NI-CNN\* respectively. The obtained results of the compared mod-

els with different combinations of binaural integration methods and loss functions are tabulated in Table 1. BAST-NSP has achieved the best AD error of 1.29° and the best MSE of 0.001 when using the subtraction binaural integration and the hybrid loss function. Compared to NI-CNN, BAST-NSP reduces AD error 65.4% from 3.70° to 1.29° and MSE 90.9% from 0.011 to 0.001. In addition, BAST-NSP outperforms NI-CNN\* (NI-CNN\*: AD=1.85°, MSE=0.031), although both models received the same input. BAST-SP has achieved AD=1.43°, MSE=0.002, therefore surpassing the performance of NI-CNN and NI-CNN\* models while it is still inferior to BAST-NSP. Overall, in this study, BAST-NSP outperforms the other tested models in performing binaural sound localization.

We further analyze the influence of different binaural integration methods on the BAST-NSP and BAST-SP performance. Here, the performances are compared in terms of AD error. Specifically, in both cases when the BAST-NSP is trained by AD loss and hybrid loss, binaural integration through addition and subtraction improved the model performance compared to concatenation (AD loss: Add.=1.30°, Sub.=1.63°, Concat.=2.39°; Hybrid loss: Add.=1.83°, Sub.=1.29°, Concat.=2.76°, see Table 1). In case of MSE loss, the performance across the three integration methods of BAST-NSP is similar. In BAST-SP, addition integration causes a huge AD error increment over BAST-NSP (BAST-SP: 4.97°, BAST-NSP: 1.30°), indicating that the left-right identical feature addition brings a great challenge to the model to predict the azimuth.

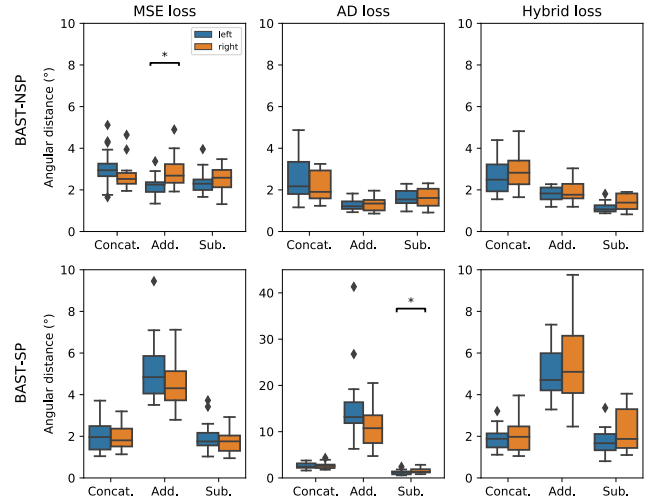
The effect of three types of loss functions on BAST-NSP and BAST-SP performance are not the same. In BAST-NSP, AD loss achieves lower AD when using concatenation or addition, but the hybrid loss yields the lowest AD in terms of subtraction. In BAST-SP, one can observe an interaction of loss function and binaural integration methods, i.e. the best loss function depends on the applied binaural integration methods.

### 5.2. Performance at different azimuths

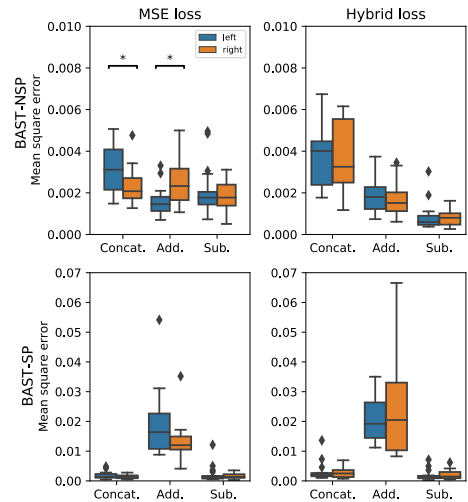
To better understand the localization performance of the models in different azimuth, the test AD error of each azimuth is shown in Fig. 2. The test AD error in BAST-NSP is much smaller when the sound source is located closer to the interaural midline. However, this pattern is not observed in BAST-SP.

### 5.3. Performance in left and right hemifield


To explore the symmetry of the model predictions, we further compare the evaluation metrics between the left-right hemifields. One can observe comparable model performance in left and right hemifield in Fig. 3 and 4. This result is confirmed by paired t-test (False Discovery Rate (FDR) corrected for multiple comparisons). More specifically, Fig. 3, shows an insignificant difference of AD error between the left and right hemifield in most conditions (corrected  $p > 0.05$ ). However, a minor but significant difference (corrected  $p < 0.05$ ) is observed in BAST-NSP trained with MSE loss and addition integration, and in BAST-SP trained with AD loss and subtraction. More precisely, the difference in AD error between the left and right hemifield was not significant in most conditions (corrected  $p > 0.05$ , Fig.



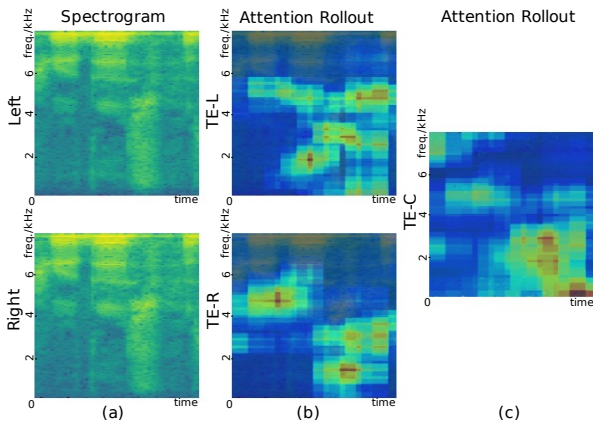
**Figure 3:** The AD error of the proposed BAST-NSP and BAST-SP in the left and right hemifield. The boxplot indicates quartiles of the metric distribution with respect to azimuths. The asterisk between two boxes indicates the statistical significance ( $p < 0.05$ , paired t-test with FDR correction) between the left and right hemifield.



**Figure 4:** The MSE of the proposed BAST-NSP and BAST-SP in the left and right hemifield. The boxplot indicates quartiles of the metric distribution with respect to azimuths. The asterisk between two boxes indicates the statistical significance ( $p < 0.05$ , paired t-test with FDR correction) between the left and right hemifield.



**Figure 5:** An example of the attention matrices in the proposed model (i.e., BAST-NSP, hybrid loss and subtraction). The corresponding sound clip was randomly selected in the category of human speech with reverberation. For each layer, we present the patch-to-patch attention matrix (size: 180×180) calculated by the rollout method in [32]. Note that we initialize the attention matrix at the first layer of TE-C by summing the attention matrices at the last layer of TE-L and TE-R.



**Figure 6:** Attention rollout corresponding to the spectrogram shown in Fig. 5. (a): Left and right spectrogram. (b): The left and right attention rollout obtained from the 3rd layer of TE-L and TE-R Transformers. (c): The left and right Attention rollout obtained from the 3rd layer of TE-C.

3), thus supporting the consistent symmetry of model predictions.

#### 5.4. Performance in different environments

We conduct two additional experiments to illustrate the generalization of the proposed model by training in one listening environment and testing in both environments separately, i.e., AE and RV. This analysis is conducted on the best performing model, i.e., BAST-NSP with hybrid loss and subtraction integration method. As shown in Table 3, the model that is trained using the data of both AE and RV environments, achieves the best test results compared to other models which are trained only using the data of one of the environment.

#### 5.5. Attention Analysis

To interpret the localization process, we utilize Attention Rollout [33] to visualize the attention maps of the proposed model (BAST-NSP with subtraction method and hybrid loss).

Table 3: The performance of the proposed BAST-NSP model in different listening environments. AE and RV indicate the anechoic and reverberation environments respectively.

Training Environment	Testing Environment	AD	MSE
AE	AE	1.14°	0.001
	RV	8.66°	0.027
RV	AE	16.70°	0.078
	RV	1.65°	0.002
AE+RV	AE	<b>1.10°</b>	<b>0.001</b>
	RV	<b>1.48°</b>	<b>0.001</b>

Rollout calculates the attention matrix by recursively multiplying the attention matrices along the forward propagation path. [32] enhanced this method by adding an additional identical matrix before multiplication to simulate the effect of residual connection of MSA. Due to the interaural integration layer, in BAST-NSP, we initialize the attention matrix of TE-C by summing the attention weights from both sides regardless of the integration method.

Fig. 5 shows the patch-to-patch attention matrices (size: 180×180) of a randomly selected spectrogram. In the first layer of each Transformer, most of the patches are self-focused and pay attention to some scattered patches. However, in the last layer, all patches yield nearly consistent attention weights to some specific patches. The attention rollout heat map with respect to the left and right spectrogram are depicted in Fig. 6. Although parameter-sharing setting is not used in BAST-NSP, one can still observe that the model focuses most of its attention on similar regions on both sides, see Fig. 6(b). The final attention map, Fig. 6(c), shows that the model further processes the attention after the integration layer and boosts the attention weights in bottom left regions.

## 6. Conclusion

In this paper, a novel Binaural Audio Spectrogram Transformer (BAST) for sound source localization is proposed. The obtained results show that this pure attention-based model leads to significant azimuth acuity improvement compared to NI-CNN model. In particular, subtraction interaural integration and hybrid loss is the best training combination for BAST. Additionally, we found that the performance and statistical significance in left-right hemifields vary with different combinations of training settings. In conclusion, this work contributes to a convolution-free model of real-life sound localization. The data and implementation of our BAST model are available at <https://github.com/ShengKuangCN/BAST>.

## References

- [1] D. W. Batteau, The role of the pinna in human localization, *Proceedings of the Royal Society of London. Series B. Biological Sciences* 168 (1011) (1967) 158–180.
- [2] J. O. Pickles, Auditory pathways: anatomy and physiology, *Handbook of clinical neurology* 129 (2015) 3–25.
- [3] K. van der Heijden, J. P. Rauschecker, B. de Gelder, E. Formisano, Cortical mechanisms of spatial hearing, *Nature Reviews Neuroscience* 20 (10) (2019) 609–623.
- [4] B. Grothe, M. Pecka, D. McAlpine, Mechanisms of sound localization in mammals, *Physiological reviews* 90 (3) (2010) 983–1012.
- [5] J. Blauert, S. Hearing, *The psychophysics of human sound localization*, in: *Spatial Hearing*, MIT Press, 1997.
- [6] Y. LeCun, Y. Bengio, G. Hinton, *Deep learning*, *nature* 521 (7553) (2015) 436–444.
- [7] X. Zhang, D. Wang, Deep learning based binaural speech separation in reverberant environments, *IEEE/ACM transactions on audio, speech, and language processing* 25 (5) (2017) 1075–1084.
- [8] S. Y. Lee, J. Chang, S. Lee, Deep learning-based method for multiple sound source localization with high resolution and accuracy, *Mechanical Systems and Signal Processing* 161 (2021) 107959.
- [9] Y. Gong, Y.-A. Chung, J. Glass, Ast: Audio spectrogram transformer, *arXiv preprint arXiv:2104.01778* (2021).
- [10] T.-D. Truong, C. N. Duong, H. A. Pham, B. Raj, N. Le, K. Luu, et al., The right to talk: An audio-visual transformer approach, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021*, pp. 1105–1114.
- [11] L. Perotin, R. Serizel, E. Vincent, A. Guérin, Crnn-based joint azimuth and elevation localization with the ambisonics intensity vector, in: *2018 16th International Workshop on Acoustic Signal Enhancement (IWAENC), IEEE, 2018*, pp. 241–245.
- [12] T. Yoshioka, S. Karita, T. Nakatani, Far-field speech recognition using cnn-dnn-hmm with convolution in time, in: *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP), IEEE, 2015*, pp. 4360–4364.
- [13] S. Park, Y. Jeong, H. S. Kim, Multiresolution cnn for reverberant speech recognition, in: *2017 20th Conference of the Oriental Chapter of the International Coordinating Committee on Speech Databases and Speech I/O Systems and Assessment (O-COCOSDA), IEEE, 2017*, pp. 1–4.
- [14] Y. Xu, S. Afshar, R. K. Singh, R. Wang, A. van Schaik, T. J. Hamilton, A binaural sound localization system using deep convolutional neural networks, in: *2019 IEEE International Symposium on Circuits and Systems (ISCAS), IEEE, 2019*, pp. 1–5.
- [15] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *Proceedings of the IEEE conference on computer vision and pattern recognition, 2016*, pp. 770–778.
- [16] N. Yalta, K. Nakadai, T. Ogata, Sound source localization using deep learning models, *Journal of Robotics and Mechatronics* 29 (1) (2017) 37–48.
- [17] K. van der Heijden, S. Mehrkanoon, Goal-driven, neurobiological-inspired convolutional neural network models of human spatial hearing, *Neurocomputing* 470 (2022) 432–442.
- [18] K. van der Heijden, S. Mehrkanoon, Modelling human sound localization with deep neural networks., in: *ESANN, 2020*, pp. 521–526.
- [19] A. Francl, J. H. McDermott, Deep neural network models of sound localization reveal how perception is adapted to real-world environments, *Nature Human Behaviour* 6 (1) (2022) 111–133.
- [20] N. Ma, J. A. Gonzalez, G. J. Brown, Robust binaural localization of a target sound source by combining spectral source models and deep neural networks, *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 26 (11) (2018) 2122–2131.
- [21] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, I. Polosukhin, Attention is all you need, *Advances in neural information processing systems* 30 (2017).
- [22] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, *arXiv preprint arXiv:1810.04805* (2018).
- [23] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al., An image is worth 16x16 words: Transformers for image recognition at scale, *arXiv preprint arXiv:2010.11929* (2020).
- [24] K. Han, Y. Wang, H. Chen, X. Chen, J. Guo, Z. Liu, Y. Tang, A. Xiao, C. Xu, Y. Xu, et al., A survey on vision transformer, *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2022).
- [25] Z. Zhang, S. Xu, S. Zhang, T. Qiao, S. Cao, Attention based convolutional recurrent neural network for environmental sound classification, *Neurocomputing* 453 (2021) 896–903.
- [26] Y.-B. Lin, Y.-C. F. Wang, Audiovisual transformer with instance attention for audio-visual event localization, in: *Proceedings of the Asian Conference on Computer Vision, 2020*.
- [27] Q. Kong, Y. Xu, W. Wang, M. D. Plumbley, Sound event detection of weakly labelled data with cnn-transformer and automatic threshold optimization, *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 28 (2020) 2450–2460.
- [28] K. J. Piczak, Esc: Dataset for environmental sound classification, in: *Proceedings of the 23rd ACM international conference on Multimedia, 2015*, pp. 1015–1018.
- [29] P. Warden, Speech commands: A dataset for limited-vocabulary speech recognition, *arXiv preprint arXiv:1804.03209* (2018).
- [30] X. Xiao, S. Zhao, X. Zhong, D. L. Jones, E. S. Chng, H. Li, A learning-based approach to direction of arrival estimation in noisy and reverberant environments, in: *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, 2015*, pp. 2814–2818.
- [31] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, *arXiv preprint arXiv:1412.6980* (2014).
- [32] H. Chefer, S. Gur, L. Wolf, Transformer interpretability beyond attention visualization, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021*, pp. 782–791.
- [33] S. Abnar, W. Zuidema, Quantifying attention flow in transformers, *arXiv preprint arXiv:2005.00928* (2020).