



Vision Transformer Segmentation for Visual Bird Sound Denoising

Sahil Kumar¹, Jialu Li², Youshan Zhang¹

¹Department of Artificial Intelligence and Computer Science, Yeshiva University, New York, NY, USA

²School of Public Policy, Cornell University, Ithaca, NY, USA.

skumar4@mail.yu.edu, jl4284@cornell.edu, youshan.zhang@yu.edu

Abstract

Audio denoising, especially in the context of bird sounds, remains a challenging task due to persistent residual noise. Traditional and deep learning methods often struggle with artificial or low-frequency noise. In this work, we propose ViTVS, a novel approach that leverages the power of the vision transformer (ViT) architecture. ViTVS adeptly combines segmentation techniques to disentangle clean audio from complex signal mixtures. Our key contributions encompass the development of ViTVS, introducing comprehensive, long-range, and multi-scale representations. These contributions directly tackle the limitations inherent in conventional approaches. Extensive experiments demonstrate that ViTVS outperforms state-of-the-art methods, positioning it as a benchmark solution for real-world bird sound denoising applications. Source code is available at: <https://github.com/aiai-4/ViTVS>.

Index Terms: Audio Denoising, Transformer, Segmentation

1. Introduction

Audio denoising remains a persistent challenge in modern applications such as teleconferences, social media speech-to-text, and hearing aids [1]. In the context of bird sound processing, the intricate and varied nature of avian vocalizations presents unique challenges. When addressing these challenges, traditional and deep learning methods are always constrained by issues like artificial noise and low-frequency noise. This has fueled the exploration of innovative solutions tailored for effective bird sound denoising [2].

Audio denoising methods encompass a spectrum of representations including, time-domain, frequency-domain, and time-frequency domain approaches. Traditional methods, such as the Gaussian mixture model [3], rely on statistical estimation to extract clean audio from noisy signals. Time-domain algorithms utilize techniques such as Short-Time Fourier Transform (STFT) [4] and Inverse Short-Time Fourier Transform (ISTFT) [5]. Additionally, denoising performance is enhanced by applying Wiener filters and LSA estimators [6, 7]. Deep learning models, such as full convolutional neural networks (FCN), IBM-based DNNs, and Noise2Noise methods, have demonstrated superiority in audio denoising [8]. These models exhibit effective noise pattern capture and improved learning capabilities, especially with limited training samples [9, 10, 11, 12, 13]. The vision transformer (ViT) [14], such as CleanUNet, swin transformer, and CrossVit [15, 16], has found success in various domains, including audio adaptation tasks. The audio spectrogram transformer showed the potential of attention-based models in audio classification [17]. Gul and Khan [8] explored various audio enhancement algorithms, utilizing image U-Net for quality improvement across multiple domains.

In recent years, deep learning applications in bird sound denoising have been extensively explored. Li et al., [18] introduced the Deeplabv3+ vision transformer, emphasizing perceptual quality enhancement for noisy audio signals. Zhang and Li [19] proposed a complex image generation SwinTransformer network, addressing challenges in high-performance audio denoising while emphasizing the quality of generated time-frequency representations. Additionally, Zhang et al., [20] explored the influence of recording devices and environmental noise on acoustic index scores, providing implications for bird sound-based assessments. These studies collectively contribute to advancing bird sound denoising through innovative deep-learning approaches [21, 22, 23, 24, 25]. We assessed our model using ground truth provided by [20]. The dataset was sourced from natural environments, introducing a potential bias if noise bandwidths perfectly match bird sounds. Most noises have distinct shapes from bird sounds and are removed by our model.

We present a vision transformer for visual (ViTVS) bird sound denoising model, as shown in Fig. 1. By transforming audio denoising into an image segmentation problem, ViTVS harnesses the capabilities of the vision transformer (ViT) architecture. The model includes essential modules: Short-Time Fourier Transform (STFT) for audio-to-image conversion, a vision transformer-based ViTVS model for segmentation, and Inverse Short-Time Fourier Transform (ISTFT) for audio reconstruction. ViTVS generates comprehensive, long-range, and multi-scale representations, distinguishing itself from conventional methods. The transformer-based encoder-decoder architecture, exemplified by PtDeepLab [18], effectively captures and fuses multi-scale features, showcasing its prowess in audio image segmentation. Our contributions are three-fold:

1. We introduce a vision transformer (ViTVS) model for bird sound denoising. Unlike existing ViT models, ViTVS redefines audio denoising as an image segmentation challenge, showcasing its distinctive ability to handle intricate bird sound patterns. The encoder-decoder architecture of ViTVS effectively captures and fuses multi-scale features, surpassing the capabilities of traditional ViT methods.
2. ViTVS encompasses important components by adding a PatchEmbedding layer to transform input audio into image-like patches and developing a SelfAttentionEncoderBlock for multi-head attention.
3. ViTVS achieves state-of-the-art bird sound denoising performance and intricate pattern learning. Leveraging negative log-likelihood loss enhances denoising capabilities, establishing ViTVS as a benchmark.

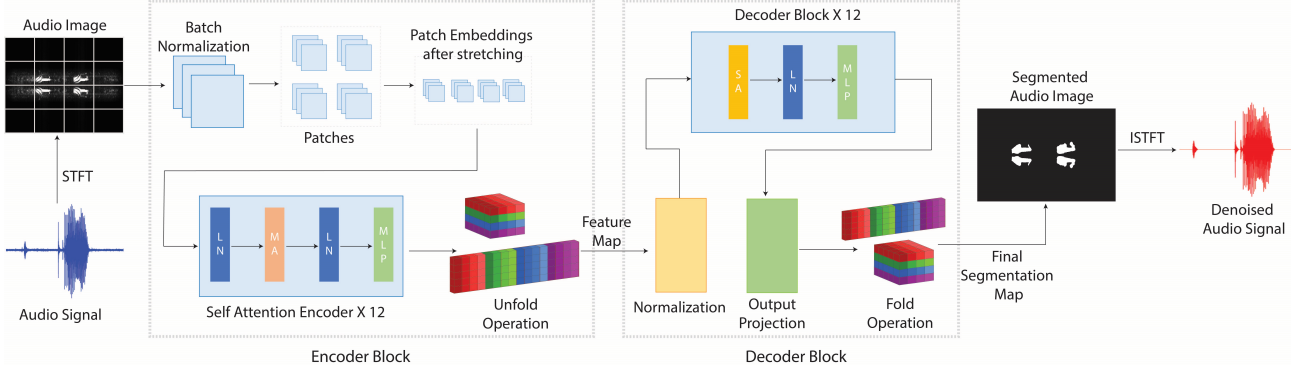


Figure 1: The overview of our ViTVS architecture. The encoder comprises a sequence of self-attention encoder blocks, each executing normalization, patch creation, and embedding layers. The decoder mirrors the encoder with additional operations, including unfolding and output projection, culminating in the final segmentation map. Both encoder and decoder consist of 12 blocks.

2. Methods

2.1. Problem

A noisy audio signal $y(t)$ is modeled as the sum of a clean audio signal $x(t)$ and additive noise $\varepsilon(t)$ [26]:

$$y(t) = x(t) + \varepsilon(t) \quad (1)$$

The objective of audio denoising is to learn a mapping \mathcal{F} that extracts the clean audio component X from the mixture audio signal $Y = \{y_i\}_{i=1}^N$, minimizing the approximation error between the estimated denoised audio $\mathcal{F}(Y)$ and the clean audio X . Our approach also transforms audio denoising into an image segmentation task, aiming to minimize the error between predictions of audio images $\mathcal{I} = \{I_i\}_{i=1}^N$ from our model and their corresponding ground truth labeled masks $U = \{u_i\}_{i=1}^N$.

2.2. Motivation

Our motivation stems from the limitations of previous bird sound denoising methods [4], which struggle to effectively separate bird calls from background noise. Traditional filtering techniques and general-purpose deep learning models are not designed specifically for bird sound tasks. To address these challenges, we propose the ViTVS model, tailored to capture the unique characteristics of avian vocalizations. By transforming audio denoising into an image segmentation task using a novel vision transformer model, we aim to improve denoising performance and accuracy. Our research seeks to advance techniques for analyzing and processing bird vocalizations.

2.3. Preliminary

Our model utilizes Short-Time Fourier Transform (STFT) and Inverse Short-Time Fourier Transform (ISTFT) for analyzing and reconstructing audio signals [6]. We use the absolute value of STFT to segment and identify noise masks, while preserving the original phase in ISTFT for accurate time-domain signal reconstruction. This technique ensures effective noise removal by processing the audio into frames of short-term spectra.

Audio Image Construction: STFT is suitable for non-stationary audio signals since it is capable of dividing the signal into frames and analyzing their short-term spectra:

$$I = |\text{STFT}_y(t, f)|. \quad (2)$$

where (I) represents the construction of an audio image at time t and frequency f . We generate audio images using STFT that can convert bird sound audio to image representations.

Denoised Audio Reconstruction: After noise removal in the frequency domain, denoised audio is reconstructed using ISTFT. Filtering noise areas in the frequency domain is guided by predicted masks ($\tilde{u} = F(I)$) [4]:

$$\mathcal{O}' = \mathcal{O}, \quad \text{and} \quad \mathcal{O}'[\tilde{u} < 1] = 0, \quad (3)$$

where \mathcal{O} is the original audio image, \mathcal{O}' will get the predicted denoised masks, Eq. (4) reconstructs denoised audio ($\hat{y}(t)$) using ISTFT after frequency domain noise removal.

$$\hat{y}(t) = \text{ISTFT}(\mathcal{O}') \quad (4)$$

3. Model Architecture

3.1. Input Transformation

Firstly, the model undergoes a meticulous input transformation involving batch normalization and image-to-patch conversion. The input image I experiences normalization through Batch Normalization (BN), as depicted in Eq. (5):

$$\text{BN}(X) = \frac{X - \mu}{\sqrt{\sigma^2 + \epsilon}} \cdot \gamma + \beta, \quad (5)$$

where X is the input, μ and σ are mean and standard deviation, ϵ is a stability constant, γ is the scale parameter, and β is the shift parameter. Subsequently, I undergoes linear embedding, resulting in a curated sequence $X \in \mathbb{R}^{N \times D}$, as expressed in Eq. (6):

$$X = \text{Linear}(\text{ITP}(\text{BN}(I))), \quad (6)$$

where ITP denotes ImageToPatches operation.

Image-to-Patch Conversion: The essential image-to-patch function transforms I into non-overlapping patches of size $p \times p$, as defined in Eq. (7):

$$\text{ITP}(I)_{mn} = \text{flatten}(I_{(pm):(pm+p-1), (pn):(pn+p-1), :}), \quad (7)$$

where $I_{(pm):(pm+p-1), (pn):(pn+p-1), :}$ selects the patch at position (m, n) in the image, and **flatten** reshapes this patch into a flattened vector. We ensure proper handling of patch edges during the reassembly of these patches into a complete image.

3.2. Encoder

In ViTVS model, the encoder and decoder play crucial roles in their architectures. The encoder serves as the backbone, consisting of self-attention-based blocks meticulously organized into

12 layers. **Self-Attention Encoder Block:** Each encoder block (i) includes self-attention mechanisms and feedforward neural networks [27], complemented by layer normalization. The progression within an encoder block is expressed in Eqs.(8), (9), and (10):

$$\text{Attention_Output}_i = \text{MultiheadAttention}(\text{LN}(X_i)), \quad (8)$$

$$\text{MLP_Output}_i = \text{MLP}(\text{LN}(\text{Attention_Output}_i)), \quad (9)$$

$$X_{i+1} = X_i + \text{MLP_Output}_i, \quad (10)$$

Multihead Self-Attention: The intricate operation of queries Q , keys K , and values V unfolds through linear projections of X . The nuances of this process are encapsulated in Eqs. (11) and (12):

$$Q = X \cdot W_Q, \quad K = X \cdot W_K, \quad V = X \cdot W_V, \quad (11)$$

$$\text{Attention_Scores} = \text{softmax} \left(\frac{QK^T}{\sqrt{D_k}} \right) \cdot V, \quad (12)$$

where X is the input sequence at a certain layer, W_Q , W_K , and W_V are learnable weight matrices and D_k is the dimension of queries and keys [27].

Feedforward Neural Network (MLP): The Multi-Layer Perceptron (MLP) works with two linear layers and GELU activation, as portrayed in Eq. (13):

$$\text{MLP}(X) = \text{GELU}(X \cdot W_1 + b_1) \cdot W_2 + b_2, \quad (13)$$

where W_1 , W_2 , b_1 , and b_2 represent the weights matrix and biases vectors that the model learns through training to make predictions based on the input X . The subscripts 1 and 2 denote the layers of the MLP.

3.3. Decoder

The decoder mirrors the encoder structure but incorporates an output projection layer, also featuring 12 layers.

Decoder Block: Each decoder block entails a self-attention encoder block, layer normalization, and an MLP. The baton passes from one block to the next, as expressed in Eqs. (14), (15), and (16):

$$\text{Attention_Output}_i = \text{SelfAttentionEncoderBlock}(\text{LN}(X_i)), \quad (14)$$

$$\text{MLP_Output}_i = \text{MLP}(\text{LN}(\text{Attention_Output}_i)), \quad (15)$$

$$X_{i+1} = X_i + \text{MLP_Output}_i, \quad (16)$$

Output Projection: The final step involves projecting the output to the original image dimensions through a linear layer, followed by a folding process, as defined in Eq. (17):

$$\text{Output} = \text{Fold}(\text{Linear}(X)). \quad (17)$$

In short, both the encoder and decoder architectures comprise 12 layers, each housing self-attention blocks, layer normalization, and multi-layer perceptrons (MLPs).

3.4. Loss Function and Algorithm

Loss Function: We train our ViTVS model using a negative log-likelihood loss (NLL) [28], expressed in Eq. (18):

$$\text{Loss} = -\frac{1}{N} \sum_{i=1}^N \log \left(\frac{\exp(\text{Output}_i[M_i])}{\sum_{j=1}^C \exp(\text{Output}_i[j])} \right), \quad (18)$$

where N denotes the number of pixels, C signifies the number of classes, j is a summation index representing each class (1 to C). and M_i represents the ground truth label for pixel i .

This loss function quantifies dissimilarity between predicted and ground truth labels, fostering the model to generate accurate and contextually relevant segmentation maps. The negative log-likelihood [28] penalizes incorrect predictions, contributing to a robust learning process. The log-softmax ensures numerical stability during computation.

Training Algorithm. We train our ViTVS model using Alg. 1.

Algorithm 1 Vision Transformer Segmentation for Visual Bird Sound Denoising

Require: Noisy audio signals $Y = \{y_i\}_{i=1}^N$ and labeled mask images $M = \{m_i\}_{i=1}^N$, where N is the total number of audio samples.

Ensure: Denoised audio signals $\mathcal{F}(Y)$

- 1: Generate audio images \mathcal{I} using Eq. (2)
 - 2: Initialize vision transformer model with encoder (12 blocks), decoder (12 blocks), and likelihood loss function.
 - 3: **for** iter = 1 **to** II **do**
 - 4: **for** $k = 1$ **to** n_B **do**
 - 5: Extract batch-wise data: I^k and M^k sampled from \mathcal{I} and M using Eqs. (5), (6) and (7).
 - 6: Encode audio images using the Vision Transformer encoder with 12 blocks using Eqs. (8), (9), and (10).
 - 7: Obtain segmentation using the vision transformer decoder with 12 blocks using Eq. (17).
 - 8: Calculate likelihood loss using the specified likelihood loss function using Eq. (18).
 - 9: Backward pass and update model parameters.
 - 10: **end for**
 - 11: **end for**
 - 12: Obtain the clean frequency domain using Eq. (3)
 - 13: Output the denoised audio signals using Eq. (4)
-

4. Experiments

4.1. Implementation details

To ensure the efficacy and robustness of our approach, we utilized PyTorch version 2.0.1 in conjunction with torchvision version 0.15.2. This framework was complemented by CUDA version 11.7, leveraging parallel processing capabilities on the NVIDIA GeForce RTX 4070 GPU. The learning rate is set to $5e-5$ with an AdamW optimizer, and the weight decay is adjusted to $5e-4$. Further, we resize input images to $256 \times 256 \times 3$ with a mini-batch size of 8 for 100 epochs.

4.2. Datasets

Our model is evaluated using the BirdSoundsDenoising dataset, which contains 14,120 audio signals lasting one to fifteen seconds each, split into training (10,000), validation (1,400), and testing (2,720) sets [4]. This dataset stands out by featuring natural environmental noises—like wind, waterfalls, and rain—instead of artificial noise.

4.3. Results

We evaluate our ViTVS model against existing methods in terms of learning efficacy, capability, and qualitative outcomes.




Figure 2: Segmentation results comparisons. The leftmost column is the original audio image. The ground truth is the labeled mask.

We use three key metrics—F1, IoU, and Dice—for image segmentation performance assessment [4], and the signal-to-distortion ratio (SDR) for audio denoising efficacy, as outlined in [18]. Higher scores in these metrics indicate better segmentation and denoising performance, reflecting closer alignment with ground truth, which calculated based on labeled masks rather than real-world clean bird sounds. As shown in Tab. 1, Fig. 1, and Fig. 2, our model significantly enhances denoising performance, paving the way for applications in bird species classification and audio signal enhancement.

4.4. Performance Comparison

Fig. 2 shows the comparisons of five best baseline segmentation models. The performance of various audio denoising methods is detailed in Tab. 1, including *F1* score, Intersection over Union (*IoU*), Dice score, and *SDR* on both validation and test sets. Among these methods, ViTVS stands out, demonstrating robust denoising capabilities. ViTVS achieves the highest *F1*, *IoU* and *Dice* scores on both validation and test sets, highlighting its superior segmentation accuracy. The *SDR* values reinforce ViTVS’s excellence in denoising, positioning it as a compelling choice for audio denoising. ViTVS consistently outperforms state-of-the-art methods across multiple metrics, emphasizing its efficacy and potential for real-world applications.

To investigate architectural variations and determine the optimal configuration, we conducted an ablation study on our Vision Transformer Segmentation (ViTVS) model. The study focused on comparing variants, specifically the number of blocks in both the encoder and decoder. We trained five variants with block configurations ranging from 4 to 20, evaluating their performance on validation and test sets using *IoU*, *Dice*, and *F1* scores. Notably, increasing blocks from 4 to 12 substantially improved denoising capabilities, emphasizing the pivotal role of block count. Intriguingly, further experimentation with 16 blocks did not yield additional benefits, suggesting 12 blocks strike an optimal balance between complexity and performance. The ViTVS 12-block variant showcased superior results compared to its 4, 5, 9, 16, and 20-block counterparts in terms of *IoU*, *Dice*, and *F1* scores, underscoring the significance of a finely tuned model for optimal denoising. These compelling findings, detailed in Tab. 2, highlight the critical importance of selecting the proper number of blocks for better results.

5. Discussion

Our ViTVS model demonstrates superior performance compared to the CNN-based PtDeepLab model [18] and eleven other state-of-the-art segmentation and denoising methods. It excels in segmentation results and enhances representation ca-

Table 1: Results comparisons of different methods (*F1*, *IoU*, and *Dice* scores are multiplied by 100).

Networks	Validation				Test			
	<i>F1</i>	<i>IoU</i>	<i>Dice</i>	<i>SDR</i>	<i>F1</i>	<i>IoU</i>	<i>Dice</i>	<i>SDR</i>
U ² -Net [29]	60.8	45.2	60.6	7.85	60.2	44.8	59.9	7.70
MTU-NeT [30]	69.1	56.5	69.0	8.17	68.3	55.7	68.3	7.96
Segmenter [31]	72.6	59.6	72.5	9.24	70.8	57.7	70.7	8.52
U-Net [32]	75.7	64.3	75.7	9.44	74.4	62.9	74.4	8.92
SegNet [33]	77.5	66.9	77.5	9.55	76.1	65.3	76.2	9.43
YOLOv8 [30]	74.5	62.1	74.5	8.33	73.6	60.9	73.3	8.12
FPN [34]	82.4	70.1	82.5	9.91	81.2	69.3	81.2	9.64
Unet++ [35]	83.0	71.0	83.1	9.98	82.1	69.8	82.1	9.87
MA-net [36]	83.5	71.7	83.5	10.12	82.4	70.9	82.4	9.92
DVAD [4]	82.6	73.5	82.6	10.33	81.6	72.3	81.6	9.96
PtDeepLab [18]	83.4	75.9	83.4	10.49	83.1	75.4	83.0	10.43
ViTVS	88.3	80.9	88.3	11.04	90.7	80.0	87.6	10.86

Table 2: Ablation study results comparing ViTVS variants.

Model Variant	Validation			Test		
	<i>IoU</i>	<i>Dice</i>	<i>F1</i>	<i>IoU</i>	<i>Dice</i>	<i>F1</i>
ViTVS 4-block	60.3	72.6	72.7	54.5	68.4	68.4
ViTVS 5-block	68.8	76.7	76.7	67.9	76.0	76.0
ViTVS 9-block	76.8	84.7	84.7	75.2	84.0	84.0
ViTVS 12-block	80.9	88.3	88.3	80.0	87.6	90.7
ViTVS 16-block	78.5	86.2	86.3	78.1	85.5	85.5
ViTVS 20-block	80.6	88.1	76.7	79.9	87.5	87.5

capacity by effectively capturing multi-scale information for detailed context patterning. The ablation study results, shown in Tab. 2, and the segmented mask quality in Fig. 2 validate our model’s robustness in preserving essential clean signals. While converting audio denoising to segmentation proposed in [4], our unique bird vision-transformed segmentation model marks the first of its kind to notably enhance denoising tasks.

6. Conclusion

In this paper, we introduce a vision transformer for a visual bird sound denoising model (ViTVS). Our ViTVS model is designed with a self-attention mechanism, patch-based processing, layer normalization, and multi-layer perceptrons to capture long-range dependencies and effectively handle diverse input sizes. ViTVS exhibits superior performance to state-of-the-art methods, as evidenced by extensive experimental evaluations, including comparisons with CNN-based self-attention and diffusion models. Therefore, our developed ViTVS model is suitable for bird sound denoising.

7. References

- [1] Y. Ephraim and D. Malah, "Speech enhancement using a minimum mean-square error log-spectral amplitude estimator," *IEEE transactions on acoustics, speech, and signal processing*, vol. 33, no. 2, pp. 443–445, 1985.
- [2] N. Krishnamurthy and J. H. Hansen, "Babble noise: modeling, analysis, and applications," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 17, no. 7, pp. 1394–1407, 2009.
- [3] L. Wang, W. Zheng, X. Ma, and S. Lin, "Denosing speech based on deep learning and wavelet decomposition," *Scientific Programming*, vol. 2021, 07 2021.
- [4] Y. Zhang and J. Li, "Birdsoundsdenoising: Deep visual audio denoising for bird sounds," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023, pp. 2248–2257.
- [5] K. Wang, B. He, and W.-P. Zhu, "Tstnn: Two-stage transformer based neural network for speech enhancement in the time domain," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 7098–7102.
- [6] L. Wang, W. Zheng, X. Ma, and S. Lin, "Denosing speech based on deep learning and wavelet decomposition," *Scientific Programming*, vol. 2021, 2021.
- [7] H. Lin, Y. Song, H. Wang, L. Xie, D. Li, G. Yang *et al.*, "Multimodal brain image fusion based on improved rolling guidance filter and wiener filter," *Computational and Mathematical Methods in Medicine*, vol. 2022, 2022.
- [8] S. Gul and M. S. Khan, "A survey of audio enhancement algorithms for music, speech, bioacoustics, biomedical, industrial and environmental sounds by image u-net," *IEEE Access*, 2023.
- [9] N. Alamdari, A. Azarang, and N. Kehtarnavaz, "Improving deep speech denoising by noisy2noisy signal mapping," *Applied Acoustics*, vol. 172, p. 107631, 2021.
- [10] F. G. Germain, Q. Chen, and V. Koltun, "Speech denoising with deep feature losses," *arXiv preprint arXiv:1806.10522*, 2018.
- [11] R. Xu, R. Wu, Y. Ishiwaka, C. Vondrick, and C. Zheng, "Listening to sounds of silence for speech denoising," 2020.
- [12] N. Saleem and M. I. Khattak, "Deep neural networks for speech enhancement in complex-noisy environments," *Int. J. Interact. Multim. Artif. Intell.*, vol. 6, pp. 84–90, 2020. [Online]. Available: <https://api.semanticscholar.org/CorpusID:198344045>
- [13] M. M. Kashyap, A. Tambwekar, K. Manohara, and S. Natarajan, "Speech denoising without clean training data: A noise2noise approach," in *Proc. Interspeech 2021*, 2021, pp. 2716–2720.
- [14] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," 2021.
- [15] Q. Kong, H. Liu, X. Du, L. Chen, R. Xia, and Y. Wang, "Speech enhancement with weakly labelled data from audioset," *arXiv preprint arXiv:2102.09971*, 2021.
- [16] C.-F. R. Chen, Q. Fan, and R. Panda, "Crossvit: Cross-attention multi-scale vision transformer for image classification," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 357–366.
- [17] Y. Gong, Y.-A. Chung, and J. Glass, "Ast: Audio spectrogram transformer," *arXiv preprint arXiv:2104.01778*, 2021.
- [18] J. Li, P. Wang, and Y. Zhang, "Deeplabv3+ vision transformer for visual bird sound denoising," *IEEE Access*, 2023.
- [19] Y. Zhang and J. Li, "Complex Image Generation SwinTransformer Network for Audio Denoising," in *Proc. INTERSPEECH 2023*, 2023, pp. 186–190.
- [20] C. Zhang, Y. Zhang, X. Zheng, X. Gao, and Z. Hao, "Influence of recording devices and environmental noise on acoustic index scores: Implications for bird sound-based assessments," *Ecological Indicators*, vol. 159, p. 111759, 2024.
- [21] Z. Li, B. Zhao, and Y. Yuan, "Bio-inspired audiovisual multi-representation integration via self-supervised learning," in *Proceedings of the 31st ACM International Conference on Multimedia*, 2023, pp. 3755–3764.
- [22] C. Zhang, K. He, X. Gao, and Y. Guo, "Automatic bioacoustics noise reduction method based on a deep feature loss network," *Ecological Informatics*, p. 102517, 2024.
- [23] J. Li, J. Li, P. Wang, and Y. Zhang, "Dcht: Deep complex hybrid transformer for speech enhancement," in *2023 Third International Conference on Digital Data Processing (DDP)*. IEEE, 2023, pp. 117–122.
- [24] A. C. P. Garcia, "Complex convolution neural network model (complex deeplabv3) on stft time-varying frequency components for audio denoising."
- [25] J. Li, P. Wang, J. Li, X. Wang, and Y. Zhang, "Dpatd: Dual-phase audio transformer for denoising," in *2023 Third International Conference on Digital Data Processing (DDP)*. IEEE, 2023, pp. 36–41.
- [26] A. Li, M. Yuan, C. Zheng, and X. Li, "Speech enhancement using progressive learning-based convolutional recurrent neural network," *Applied Acoustics*, vol. 166, p. 107347, 2020.
- [27] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," 2023.
- [28] D. Zhu, H. Yao, B. Jiang, and P. Yu, "Negative log likelihood ratio loss for deep neural network classification," 2018.
- [29] X. Qin, Z. Zhang, C. Huang, M. Dehghan, O. R. Zaiane, and M. Jagersand, "U2-net: Going deeper with nested u-structure for salient object detection," *Pattern Recognition*, vol. 106, p. 107404, 2020.
- [30] H. Wang, S. Xie, L. Lin, Y. Iwamoto, X.-H. Han, Y.-W. Chen, and R. Tong, "Mixed transformer u-net for medical image segmentation," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 2390–2394.
- [31] R. Strudel, R. Garcia, I. Laptev, and C. Schmid, "Segmenter: Transformer for semantic segmentation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 7262–7272.
- [32] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2015, pp. 234–241.
- [33] V. Badrinarayanan, A. Kendall, and R. Cipolla, "Segnet: A deep convolutional encoder-decoder architecture for image segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 12, pp. 2481–2495, 2017.
- [34] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," 2017.
- [35] Z. Zhou, M. M. R. Siddiquee, N. Tajbakhsh, and J. Liang, "Unet++: A nested u-net architecture for medical image segmentation," 2018.
- [36] R. Li, S. Zheng, C. Zhang, C. Duan, J. Su, L. Wang, and P. M. Atkinson, "Multiattention network for semantic segmentation of fine-resolution remote sensing images," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 60, p. 1–13, 2022. [Online]. Available: <http://dx.doi.org/10.1109/TGRS.2021.3093977>