


## Journal Pre-proof

Less is More: AMBER-AFNO - a New Benchmark for Lightweight 3D Medical Image Segmentation

Andrea Dosi, Semanto Mondal, Rajib Chandra Ghosh,  
Massimo Brescia, Giuseppe Longo

PII: S0957-4174(26)01431-4  
DOI: <https://doi.org/10.1016/j.eswa.2026.132518>  
Reference: ESWA 132518



To appear in: *Expert Systems With Applications*

Received date: 31 August 2025  
Revised date: 26 February 2026  
Accepted date: 16 April 2026

Please cite this article as: Andrea Dosi, Semanto Mondal, Rajib Chandra Ghosh, Massimo Brescia, Giuseppe Longo, Less is More: AMBER-AFNO - a New Benchmark for Lightweight 3D Medical Image Segmentation, *Expert Systems With Applications* (2025), doi: <https://doi.org/10.1016/j.eswa.2026.132518>

This is a PDF of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability. This version will undergo additional copyediting, typesetting and review before it is published in its final form. As such, this version is no longer the Accepted Manuscript, but it is not yet the definitive Version of Record; we are providing this early version to give early visibility of the article. Please note that Elsevier's sharing policy for the Published Journal Article applies to this version, see: <https://www.elsevier.com/about/policies-and-standards/sharing#4-published-journal-article>. Please also note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2026 Published by Elsevier Ltd.

## Less is More: AMBER-AFNO - a New Benchmark for Lightweight 3D Medical Image Segmentation

Andrea Dosi<sup>a,\*</sup> (andrea.dosi@unina.it), Semanto Mondal<sup>a</sup> (semanto.mondal@unina.it), Rajib Chandra Ghosh<sup>b</sup> (rajib.chandraghosh@unina.it), Massimo Brescia<sup>a,c</sup> (massimo.brescia@unina.it) and Giuseppe Longo (giulongo@unina.it)<sup>a</sup>

<sup>a</sup> Department of Physics E. Pancini, University of Naples Federico II, Via Cinthia 21, Naples, 80126, Italy

<sup>b</sup> Department of Electrical Engineering and Information Technology (DIETI), University of Naples Federico II, Via Claudio 21, Naples, 80125, Italy

<sup>c</sup> INAF, Astronomical Observatory of Capodimonte, Salita Moiariello 16, Naples, 80131, Italy

### \*Corresponding author

Andrea Dosi

Department of Physics E. Pancini, University of Naples Federico II, Via Cinthia 21, Naples, 80126, Italy

Tel: (0039) 3405794941

Email: andrea.dosi@unina.it

### Highlights

- First AFNO-based spectral mixing for 3D medical segmentation
- Eliminates quadratic token-to-token attention entirely
- Linear memory scaling for high-resolution 3D volumes
- 78% fewer parameters than UNETR++
- SOTA or near-SOTA on ACDC, Synapse, and BraTS

# Less is More: AMBER-AFNO - a New Benchmark for Lightweight 3D Medical Image Segmentation

Andrea Dosi<sup>a,\*</sup>, Semanto Mondal<sup>a</sup>, Rajib Chandra Ghosh<sup>b</sup>, Massimo Brescia<sup>a,c</sup> and Giuseppe Longo<sup>a</sup>

<sup>a</sup>Department of Physics E. Pancini, University of Naples Federico II, Via Cinthia 21, Naples, 80126, Italy

<sup>b</sup>Department of Electrical Engineering and Information Technology (DIETI), University of Naples Federico II, Via Claudio 21, Naples, 80125, Italy

<sup>c</sup>INAF, Astronomical Observatory of Capodimonte, Salita Moiaritello 16, Naples, 80131, Italy

## ARTICLE INFO

### Keywords:

3D Medical Image Segmentation  
Adaptive Fourier Neural Operators (AFNO)  
Lightweight transformer architecture  
Frequency-domain token mixing  
AMBER  
AMBER-AFNO  
ACDC Dataset  
Synapse Dataset  
BraTS Dataset  
Dice Similarity Coefficient (DSC)  
Hausdorff Distance (HD)

## ABSTRACT

We adapt the remote sensing-inspired AMBER model (Dosi et al., 2025) from multi-band image segmentation to 3D medical datacube segmentation. To address the computational bottleneck of the volumetric transformer, we propose the AMBER-AFNO architecture. This approach uses Adaptive Fourier Neural Operators (AFNO) instead of the multi-head self-attention mechanism. Unlike spatial pairwise interactions between tokens, global token mixing in the frequency domain avoids  $\mathcal{O}(N^2)$  attention-weight calculations. As a result, AMBER-AFNO achieves quasi-linear computational complexity and linear memory scaling.

This new way to model global context reduces reliance on dense transformers while preserving global contextual modeling capability. By using attention-free spectral operations, our design offers a compact parameterization and maintains a competitive computational complexity. We evaluate AMBER-AFNO on three public datasets: ACDC, Synapse, and BraTS. On these datasets, the model achieves state-of-the-art or near-state-of-the-art results for DSC and HD95. Compared with recent compact CNN and Transformer architectures, our approach yields higher Dice scores while maintaining a compact model size.

Overall, our results show that frequency-domain token mixing with AFNO provides a fast and efficient alternative to self-attention mechanisms for 3D medical image segmentation.

## 1. Introduction

According to the World Health Statistics 2024 report (World Health Organization, 2024), heart disease was the leading cause of death globally in 2021, accounting for approximately 9.1 million deaths. Kidney disease, lung cancer, lower respiratory infections, and stroke also ranked among the top ten global causes of mortality. In all cases, early diagnosis is crucial to reduce mortality rates and

optimize treatment strategies, including therapy or surgery. Such diagnosis relies on advanced imaging technologies such as MRI and CT scans, which generate three-dimensional (3D) volumetric data capturing anatomical structures at multiple depths (Tai et al., 2019; Florkow et al., 2022). Since a single 2D slice often fails to capture the full anatomical context, effective analysis of volumetric data cubes has become essential (Zhou et al., 2019).

In medical image segmentation tasks, the U-Net architecture (Ronneberger et al., 2015) has gained widespread popularity due to its encoder-decoder structure, which captures both local and global contextual representations through skip connections. Numerous variants have been proposed to improve multi-scale representation and robustness, including AFFU-Net (Zheng et al., 2022), OAU-Net (Song et al., 2023), ISTD-Net (Hou et al., 2022), MultiResU-Net (Lan et al., 2022), SAU-Net

\*Corresponding author

✉ andrea.dosi@unina.it (A. Dosi);  
semanto.mondal@unina.it (S. Mondal);  
rajib.chandraghosh@unina.it (R.C. Ghosh);  
massimo.brescia@unina.it (M. Brescia); giulongo@unina.it (G. Longo)  
ORCID(s): 0000-0002-5943-6867 (A. Dosi);  
0009-0000-2306-4478 (S. Mondal); 0009-0001-1137-3465 (R.C. Ghosh);  
0000-0001-9506-5680 (M. Brescia); 0000-0002-9182-8414 (G. Longo)

(Chen et al., 2024), KiU-Net (Valanarasu et al., 2022), and UCR-Net (Sun et al., 2022). However, convolutional neural networks (CNNs) inherently struggle to model long-range dependencies due to limited receptive fields determined by kernel size and stride, particularly in volumetric data where inter-slice relationships are critical.

To address these limitations, the Transformer Vaswani et al. (2017) was introduced into computer vision as Vision Transformers (ViT) Dosovitskiy et al. (2021). Transformer-based models rely on self-attention mechanisms to model long-range dependencies across the entire image. Consequently, various transformer-based architectures have been proposed for medical image segmentation, such as RT-UNet Li et al. (2022a), SWTRU Zhang et al. (2022), LMIS Zhu et al. (2024b), and SDV-TUNet Zhu et al. (2024a), which combine transformer blocks with U-Net-like structures to capture both global and local information.

However, in volumetric settings, transformer-based models introduce a significant computational bottleneck. In 3D medical image segmentation, feature maps grow cubically with spatial resolution, and the quadratic complexity of token-to-token self-attention results in high memory consumption, increased parameter counts, and longer inference time. Although efficient attention mechanisms have been proposed to approximate self-attention, most still rely on token-wise interaction operations whose complexity scales poorly with high-resolution volumes. Therefore, designing a 3D segmentation model that balances effective global context modeling with computational efficiency remains an open problem.

SegFormer (Xie et al., 2021a) addresses efficiency in semantic segmentation through a lightweight hierarchical transformer encoder and an MLP-based decoder. Building upon this foundation, AMBER (Dosi et al., 2025) extends SegFormer to multi-band image segmentation by incorporating three-dimensional convolutions and a Funnelizer layer to better model spatial-spectral relationships.

To explicitly address the efficiency-globality trade-off in 3D lightweight medical image segmentation, we propose AMBER-AFNO, where Adaptive Fourier Neural Operators (AFNO) (Guibas et al., 2021) replace multi-head self-attention within a

hierarchical transformer encoder. Instead of constructing quadratic attention matrices, AFNO performs global token mixing in the frequency domain through learnable spectral filters, enabling quasi-linear computational complexity and linear memory cost. The multi-head structural design is preserved through block-wise frequency mixing, allowing different frequency subspaces to capture distinct semantic patterns without pairwise token interactions. In contrast to attention approximation strategies or convolution-attention compression schemes, this approach reformulates global context modeling itself rather than merely reducing attention computation.

Building on the AMBER architecture, the proposed AMBER-AFNO adapts spectral-domain token mixing to full 3D volumetric segmentation. This design reduces the parameter count by nearly 78% compared with heavy transformer-based models such as UNETR++, while maintaining competitive FLOPs and segmentation performance. Unlike recent lightweight CNN-Transformer variants (e.g., LW-CTrans (Kuang et al., 2025)), whose efficiency is achieved by simplifying convolutional or attention modules, the lightweight property of AMBER-AFNO stems from spectral-domain global mixing.

We validate AMBER-AFNO on three public 3D medical segmentation benchmarks: ACDC (Bernard et al., 2018), Synapse (Landman et al., 2015), and BraTS (Menze et al., 2015). Using Dice Similarity Coefficient (DSC) and Hausdorff Distance (HD95) as evaluation metrics, the experimental results demonstrate that spectral-domain token mixing achieves a favorable accuracy-efficiency trade-off across diverse anatomical scenarios.

## 2. Related Work

Three-dimensional medical image segmentation has seen great progress since the adoption of encoder-decoder networks, such as U-Net and V-Net. Most recent studies can be roughly grouped into four non-mutually exclusive categories: (1) single-branch encoder-decoder networks, (2) hybrid and multi-branch encoders, (3) transformers and efficient attention models, and (4) lightweight 3D segmentation networks.

### Single-branch encoder-decoder networks.

Many classical CNN-based segmentation networks have been proposed, many of which evolved from the original U-Net (Ronneberger et al., 2015). Examples include UNet++ (Shaker et al., 2024), Attention U-Net (Ronneberger et al., 2015), CE-Net (Gu et al., 2019), DeepLab (Chen et al., 2018), and for volumetric data, V-Net (Milletari et al., 2016) that uses residual connections in 3D convolutional blocks and Dice loss, a loss function adapted to volumetric medical images. Later, nnU-Net (Isensee et al., 2021) demonstrated that adapting depth, patch size, and training settings to the dataset on the same architecture can lead to state-of-the-art performance on a multitude of tasks. Despite the high power of fully convolutional networks, they often use increasingly large receptive fields in deeper layers and thus cannot directly model long-range interactions across higher-resolution volumes.

#### **Hybrid and multi-branch encoders.**

CNN-Transformer frameworks were proposed to address locality limitations. TransUNet (Chen et al., 2021) incorporated ViT modules into the CNN encoder, and UNETR (Hatamizadeh et al., 2021) redefined U-Net with a transformer encoder and convolutional decoder. CoTr (Xie et al., 2021b) introduced deformable attention in 3D space to reduce computational burden. Subsequent works designed multi-branch encoders to disentangle semantic resolutions or image details. DS-TransUNet (Lin et al., 2022) employed dual-scale ViT encoders; DHR-Net (Bai et al., 2024) introduced semantic and detail branches; MILU-Net (He et al., 2024) and BSC-Net (Zhou et al., 2024) adopted dual-resolution designs. However, multi-branch models usually bring significantly more parameters and increase training complexity, limiting scalability.

#### **Transformer-based and efficient attention models.**

Global self-attention in Vision Transformers has proven effective for representation learning. In volumetric segmentation, UNETR (Hatamizadeh et al., 2021) and Swin-UNETR (Hatamizadeh et al., 2022) demonstrate the importance of hierarchical token interactions. nnFormer (Zhou et al., 2023) alternates convolution and attention blocks, and UNETR++ (Shaker et al., 2024) proposed Efficient Paired Attention (EPA) for separating spatial and channel attention. DS-UNETR++ (Jiang et al.,

2025) further introduced gated cross-attention mechanisms.

In its standard form, self-attention grows quadratically with the number of tokens, making it costly for larger 3D inputs. Linformer (Wang et al., 2020) and Performer (Choromanski et al., 2022) apply low-rank and kernel-based approximations, while FNet (Lee-Thorp et al., 2022) replaces attention with fixed Fourier transforms.

In contrast, Adaptive Fourier Neural Operators (AFNO) (Guibas et al., 2021) learn frequency filters and block diagonal channel mixing matrices for token mixing in the Fourier domain. By truncating high-frequency modes and applying adaptive frequency filters, AFNO achieves quasi-linear complexity and linear memory cost while modeling global context. AFNO was originally proposed for 3-channel images, and its potential for volumetric medical segmentation remains underexplored.

#### **Lightweight 3D medical image segmentation networks.**

For real-world clinical applications where computational resources are limited, efficient 3D segmentation aims to reduce parameters and floating-point operations while maintaining competitive performance. Efficient design principles from 2D image processing, such as depthwise separable convolution and inverted residual blocks (MobileNetV2 (Sandler et al., 2018)), and grouped pointwise convolution with channel shuffle (ShuffleNet (Zhang et al., 2018)), have inspired lightweight 3D models. LCOVNet (Zhao et al., 2021) utilizes separable spatiotemporal convolution with attention calibration, and ADHDC-Net (Liu et al., 2023) employs decoupled convolution to replace standard convolution in medical images. SlimUNETR (Pang et al., 2024) simplifies the attention mechanism, and LWCTrans (Kuang et al., 2025) combines lightweight convolution encoding with simplified transformer modules to significantly reduce parameter count while maintaining comparable Dice scores.

Most lightweight variants focus on convolutional filter redesign or attention simplification, with limited investigation of token mixing strategies themselves. In particular, spectral and domain-agnostic variants such as AFNO have not been systematically applied to volumetric medical image segmentation. Consequently, whether accurate 3D segmentation can be achieved with significantly reduced

parameter counts compared to current state-of-the-art methods remains unclear.

In summary, volumetric medical image segmentation architectures have evolved from CNN encoder–decoders to CNN–Transformer hybrids and fully attention-based models emphasizing global context modeling. While transformer-based approaches enhance long-range dependencies, their quadratic attention mechanisms are computationally intensive for high-resolution 3D images.

Unlike the above approaches, *AMBER-AFNO* bypasses pairwise attention mechanisms by performing global mixing in the frequency domain. By applying Fourier transforms and learning frequency filters over limited modes, *AMBER-AFNO* avoids constructing  $\mathcal{O}(N^2)$  attention matrices and enables quasi-linear complexity with linear memory usage, while significantly reducing parameter counts and preserving the ability to handle high-resolution 3D volumes.

### 3. Methodology

In this section, we introduce *AMBER-AFNO*, an extension of the *AMBER* model (Dosi et al., 2025), which replaces the traditional self-attention mechanism with Adaptive Fourier Neural Operators (AFNO) to reduce model complexity and improve computational efficiency. While staying close to the original *AMBER* design, *AMBER-AFNO* adapts the architecture to address the 3D volumetric data. As illustrated in Figure 1, the architecture consists of two main modules: a hierarchical Transformer encoder with 3D patch embedding and AFNO-based feature mixing; and a lightweight MLP decoder that fuses multi-scale features and predicts the final 3D segmentation mask. Unlike *AMBER*, which uses a dimensionality reduction layer (“*fun-nalizer*”) to collapse 3D features into 2D outputs, *AMBER-AFNO* operates entirely in 3D and directly outputs a volumetric segmentation mask of size  $D \times H \times W \times N_{cls}$ , where  $N_{cls}$  is the number of classes. Furthermore, the decoder integrates a deconvolutional layer to up-sample feature maps and recover the original spatial resolution.

#### 3.1. Hierarchical Transformer Encoder

We designed a series of Mix Transformer encoders (MiT) for semantic segmentation of 3D


images, replacing standard self-attention layers with Adaptive Fourier Neural Operators (AFNO) to reduce the number of parameters without compromising accuracy. Unlike conventional attention mechanisms—which are memory-intensive and scale quadratically with input size—AFNO performs global feature mixing in the frequency domain with quasi-linear complexity. This enables the encoder to capture both local and global context while significantly lowering computational and memory costs.

**Hierarchical Feature Representation.** The goal of this module is, given an input image, to generate CNN-like multi-level features. These features provide high-resolution coarse features and low-resolution fine-grained features that usually boost the performance of semantic segmentation (Xie et al., 2021a). More precisely, given an input 3D image with  $D \times H \times W$ , we perform patch merging to obtain a hierarchical feature  $F_1$  with a resolution of  $\frac{D_{qw}}{2^i} \times \frac{H}{2^i} \times \frac{W}{2^i} \times C_1$ , where  $i \in \{0, 1, 2, 3\}$  and  $C_{i+1}$  is larger than  $C_i$ .

**Overlapped Patch Merging.** We utilize merging overlapping patches to avoid the need for positional encoding. To this end, we define  $K$ ,  $S$ , and  $P$ , where  $K$  is the three-dimensional kernel size (or patch size),  $S$  is the stride between two adjacent patches, and  $P$  is the padding size. Unlike the original SegFormer, in our experiments, we set  $K = 3$ ,  $S = 1$ ,  $P = 1$ , and  $K = 3$ ,  $S = 2$ ,  $P = 2$  to perform overlapping patch merging. The patch size is intentionally kept small to preserve image details and avoid parameter explosion.  $S = 1$  preserves the original image spatial dimensions  $H$  and  $W$ , avoiding the reduction of spatial dimensions by 1/4.

**Adaptive Fourier Neural Operators (AFNO).** The self-attention mechanism, while effective in capturing global dependencies, is widely recognized as the primary computational bottleneck in transformer-based encoder architectures due to its quadratic complexity with respect to the input sequence length (Xie et al., 2021a).

In *AMBER* (Dosi et al., 2025), a multi-head self-attention scheme is used, where each attention layer incorporates a reduction factor  $R$  to mitigate this cost. This reduces the computational complexity from the standard  $\mathcal{O}(N^2)$  to  $\mathcal{O}\left(\frac{N^2}{R}\right)$ , offering a



**Figure 1: The Proposed AMBER-AFNO framework** consists of two main modules: A hierarchical Transformer encoder to extract coarse and fine features; and a lightweight MLP decoder to directly fuse these multi-level features and predict the semantic segmentation mask. **FFN** indicates a feed-forward network.

more tractable solution for high-resolution inputs (Xie et al., 2021a).

In the proposed AMBER-AFNO architecture, the entire attention mechanism is replaced with the AFNO block. AFNO uses token mixing in the frequency domain, leveraging the Fast Fourier Transform (FFT) to achieve global interactions with quasi-linear complexity. Here, the input tokens are transformed into the frequency domain using FFT, capturing global contextual information efficiently. In (Guibas et al., 2021) the authors introduced AFNO for the 2D image segmentation task. Using similar methodologies, we have extended the approach for the 3D image segmentation task.

The input to the AFNO block is a 5-D tensor which can be represented as  $x \in \mathbb{R}^{B \times D \times H \times W \times C}$  where  $B$  is the batch size,  $D$ ,  $H$ , and  $W$  are the spatial dimensions (depth, height, and width), and  $C$  is the embedding dimension.

We then apply a real-valued 3D Fast Fourier Transform (RFFT) over the spatial dimensions:

$$\hat{x} = \text{RFFT}_3(x) \in \mathbb{C}^{B \times D \times H \times (W/2+1) \times C} \quad (1)$$

The channel dimension  $C$  is partitioned into  $K$  frequency blocks:

$$\hat{x} \rightarrow \hat{x}_{\text{blk}} \in \mathbb{C}^{B \times D \times H \times (W/2+1) \times K \times \frac{C}{K}} \quad (2)$$

Each frequency block undergoes a learnable complex-valued two-layer MLP:

$$\hat{x}_{\text{blk}}^{(i)} \leftarrow W_2^{(i)} \cdot \phi \left( W_1^{(i)} \cdot \hat{x}_{\text{blk}}^{(i)} + b_1^{(i)} \right) + b_2^{(i)}, \quad \forall i \in \{1, \dots, K\} \quad (3)$$

After the operation of two layer MLP, the MLP output is reshaped back to  $\hat{x} \in \mathbb{C}^{B \times D \times H \times (W/2+1) \times C}$  and then soft shrinkage is applied to attenuate small-magnitude frequency responses:

$$\hat{x}' = \text{SoftShrink}(\hat{x}) \quad (4)$$

Finally, the inverse 3D FFT (IRFFT) is applied, and the result is combined with the original input via residual addition:

$$\tilde{x} = \text{IRFFT}_3(\hat{x}') + x \quad (5)$$

Thus, the final output of the AFNO-3D block is as follows:

$$\tilde{x} = \text{IRFFT}_3 \left( \text{SoftShrink} \left( \text{MLP}_C \left( \text{RFFT}_3(x) \right) \right) \right) + x \quad (6)$$

**Algorithm 1** shows the pseudo-implementation of the AFNO-3D block in the AMBER-AFNO architecture.

**Mix-FFN.** Likewise, in the AMBER (Dosi et al.,

**Algorithm 1:** AFNO-3D with Adaptive Weight Sharing and Adaptive Masking**Input:** Tensor  $x \in \mathbb{R}^{b \times d \times h \times w \times c}$ **Output:** Transformed tensor  $y$  $bias \leftarrow x$  $x \leftarrow \text{RFFT3}(x)$  $x \leftarrow \text{reshape}(x, [b, d, h, \frac{w}{2} + 1, k, \frac{c}{k}])$  $x \leftarrow \text{BlockMLP}(x)$  $x \leftarrow \text{reshape}(x, [b, d, h, \frac{w}{2} + 1, c])$  $x \leftarrow \text{SoftShrink}(x)$  $x \leftarrow \text{IRFFT3}(x)$ **return**  $x + bias$ **Function**  $\text{BlockMLP}(x)$ : $x \leftarrow xW_1 + b_1$  $x \leftarrow \text{ReLU}(x)$ **return**  $xW_2 + b_2$ 


2025) and SegFormer (Xie et al., 2021a), we also used the Mix-FFN, which considers the effect of zero padding using a  $3 \times 3 \times 3$  Conv in the feed-forward network (FFN). We used Mix-FFN instead of positional encoding because Mix-FFN combines both Depthwise Convolution and MLP layers to capture both local and global context in the seen (Xie et al., 2021a).

$$x_{\text{out}} = \text{MLP} \left( \text{GELU} \left( \text{Conv}_{3 \times 3 \times 3} \left( \text{MLP}(x_{\text{in}}) \right) \right) \right) + x_{\text{in}} \quad (7)$$

where  $x_{\text{in}}$  is the feature from the self-attention module. Mix-FFN mixes a  $3 \times 3 \times 3$  convolution and an MLP into each FFN.

**3.2. Lightweight All-MLP Decoder**

The four feature maps produced by the MiT encoder are first channel-projected to a common embedding size  $d$  by position-agnostic MLP layers. Then every projected tensor is trilinearly upsampled to the finest encoder resolution and concatenated along the channel axis. A  $1 \times 1 \times 1$  convolution with RELU activation and 3-D batch normalization fuses this aggregate into a compact representation. A single transposed 3-D convolution subsequently doubles the spatial dimensions, bringing the volume back to the native voxel grid while reducing the channels to  $N_{\text{cls}}$ . A final  $1 \times 1 \times 1$  convolution sharpens



**Figure 2:** Simplified Work Flow Diagram of AMBER-AFNO Architecture

the logits, yielding the segmentation mask  $M \in \mathbb{R}^{N_{\text{cls}} \times D \times H \times W}$ . Compared with the original five-stage AMBER decoder, this variant omits the dedicated spectral-reduction block and integrates the final MLP into the transposed convolution, lowering memory cost without sacrificing accuracy.

The figure 2 shows a more simplified and straightforward workflow diagram of the AMBER-AFNO Architecture.

**4. Experiments**

To assess the effectiveness of the proposed AMBER-AFNO architecture, we compare against the following state-of-the-art convolutional and transformer-based segmentation models: U-Net (Ronneberger et al., 2015), nnU-Net (Isensee et al., 2021), TransUNet (Chen et al., 2021), Swin-UNet (Cao et al., 2023), UNETR (Hatamizadeh et al., 2021), UNETR++ (Shaker et al., 2024), MISSFormer (Huang et al., 2023), Swin-UNETR (Hatamizadeh et al., 2022), nnFormer (Zhou et al., 2023), TransBTS (Wang et al., 2021), CoTr (Xie et al., 2021b), PCC-Trans (Feng et al., 2024), and LW-CTrans (Kuang et al., 2025).

The last one mentioned, LW-CTrans (Kuang et al., 2025), is a newly introduced lightweight 3D CNN-Transformer hybrid that shares the same goal as our approach (i.e., fewer model parameters) for volumetric segmentation. This baseline is presented to not only evaluate the performance of our method against the mainstream choice of a lightweight 3D CNN-Transformer hybrid but also test whether the proposed frequency-domain token mixing (AFNO) can deliver a better trade-off between performance and efficiency, given a similar compact model size constraint.

Although it is not comprehensive, this list offers

a competitive and diverse set of solid baselines. These cover purely convolutional, convolutional-recurrent, attention-based, and lightweight models. All models are tested on three publicly available datasets: the ACDC challenge dataset for cardiac MR image segmentation, the Synapse dataset for multi-organ abdominal CT segmentation, and the BraTS dataset. The next three sub-sections describe the datasets and data pre-processing steps, the training and inference protocols used in the study, and the two quantitative metrics (HD95 and DSC) used for comparative analysis.

#### 4.1. Dataset Overview

##### Automated cardiac diagnostic segmentation dataset (ACDC)

The ACDC dataset (Bernard et al., 2018) consists of 3D cardiac MRI images with multi-class annotations. A total of 200 labeled samples were used, split into 160 for training and 40 for testing. The annotated classes include the right ventricle (RV), myocardium (MYO), and left ventricle (LV). In this study, the Dice Similarity Coefficient (DSC) was used as the evaluation metric for model comparison.

##### Multi-organ CT Segmentation Dataset (Synapse)

The Synapse dataset (Landman et al., 2015) contains 30 abdominal CT scans and is commonly used for benchmarking multi-organ segmentation methods. Following the UNETR++ preprocessing protocol (Shaker et al., 2024), the dataset is split into 18 training volumes and 12 validation volumes. It provides annotations for eight abdominal organs: spleen, right kidney, left kidney, gallbladder, liver, stomach, aorta, and pancreas. For quantitative evaluation, we adopt two standard metrics in medical image segmentation: the Dice Similarity Coefficient (DSC) to measure volumetric overlap, and the 95th percentile Hausdorff Distance (HD95) to assess boundary accuracy.

##### Brain Tumor Segmentation (BraTS)

The BraTS dataset (Menze et al., 2015) contains multimodal 3D brain MRI scans of patients with gliomas, which is a type of brain tumor. It contains 387 MRI scan images for training and 73 scans for testing. Here, each image consists of four channels, such as FLAIR, T1w, T1gd, and T2w. The

images have four classes, including whole tumor (WT), enhancing tumor (ET), tumor core (TC) and background. In this experiment, the Dice Similarity Coefficient (DSC) is used to measure volumetric overlap, and the 95th percentile Hausdorff Distance (HD95) to assess boundary accuracy.

## 5. Experimental Settings

For a fair comparison with existing transformer-based baseline methods, we follow the same preprocessing and training scheme as UNETR++ (Shaker et al., 2024), UNETR (Hatamizadeh et al., 2021), DS-UNETR++ (Jiang et al., 2025) and SAM-based medical segmentation models (Cen et al., 2024). For each dataset, we use the same preprocessing procedures, including resampling to the same voxel spacing, intensity normalization, Z-score normalization, and the same cropping and padding protocols. No additional training data are used.

All models are trained using the same input resolutions, as reported in UNETR++, to guarantee strict comparability. For Synapse, the input size is  $128 \times 128 \times 64$ ; for BTCV,  $96 \times 96 \times 96$ ; for ACDC,  $160 \times 160 \times 16$  and for BraTS,  $128 \times 128 \times 128$ . All other data augmentation strategies and optimization details follow nnFormer (Zhou et al., 2023) to maintain consistency across baselines.

Training is conducted on a single NVIDIA Tesla V100-SXM2 32GB GPU (300W TDP) using CUDA 12.2. All models are trained for 1,000 epochs with an initial learning rate of 0.01 and weight decay of  $3 \times 10^{-5}$ . During inference, we employ a sliding-window strategy with 50% overlap and report the Dice score of a single model without ensemble techniques.

Deep supervision (Li et al., 2022b) is employed during training of *AMBER-AFNO* on the ACDC and Synapse datasets, thereby improving convergence stability and segmentation accuracy by enforcing auxiliary supervision at intermediate decoder stages. In contrast, for the BraTS dataset, empirical validation showed that removing deep supervision improved performance; thus, the reported BraTS results are from a model trained without deep supervision. Our loss function is a sum of Dice and Cross-Entropy losses, the same loss function used for UNETR++ and other transformer-based

**Table 1**  
High-level description of the datasets.

Dataset	Spatial Dimension	Dimension	Depth Dimension	Modality	Class	Training Samples	Test Samples
ACDC	320 × 320 (re-sampled)		90–130 slices	1 (Cine MRI)	4	160 cases	40 cases
Synapse	512 × 512		75–250 slices	1 (CT)	9	18 scans	12 scans
Brain Tumor (BraTS)	240 × 240		155 slices	4 (MRI)	4	387 scans	73 scans

baseline models. A detailed formulation of the loss function is provided in the following section.

## 6. Loss Function

For the semantic segmentation task, one possible problem that arises is class imbalance. We have considered three datasets with multilabel segmentation problems. All the datasets have class imbalance issues. The background class dominates over other classes, as the ROI for medical images is very low compared to the original shape of the image. One loss function alone, such as focal loss, dice loss or cross entropy loss can not handle this situation. To address this issue we have used a custom weighted loss function using the concept of Deep Supervision, which is a combination of cross-entropy loss and dice loss. Instead of relying on the final output of the model, the intermediate feature maps or predictions at different resolutions of the decoder block are also considered while calculating the loss and backpropagation.

$$L(G, P) = 1 - \frac{2}{J} \sum_{j=1}^J \frac{\sum_{i=1}^I G_{i,j} P_{i,j}}{\sum_{i=1}^I G_{i,j}^2 + \sum_{i=1}^I P_{i,j}^2} - \frac{1}{I} \sum_{i=1}^I \sum_{j=1}^J G_{i,j} \log P_{i,j} \quad (8)$$

Where  $\mathbf{G}$  refers to the set of ground truth labels, and  $\mathbf{P}$  refers to the set of predicted probabilities.  $P_{i,j}$  and  $G_{i,j}$  represent the predicted probability and the one-hot encoded true value of class  $j$  at voxel  $i$ , respectively.  $I$  denotes the total number of voxels, and  $J$  denotes the number of classes (Jiang et al., 2025).

## 7. Evaluation Metrics

We have adopted two primary evaluation metrics to assess segmentation performance: the Dice Similarity Coefficient (DSC), which quantifies the overlap between predicted and ground truth regions, and the 95th percentile Hausdorff Distance (HD95), which measures the spatial distance between boundary surfaces while mitigating the influence of outliers. Detailed definitions and computation procedures for these metrics are provided in the following subsection.

### Hausdorff Distance (HD95)

The HD 95 is a boundary-based metric that evaluates segmentation quality by computing the 95th-percentile distance between the predicted volume's boundary voxels and those of the ground-truth segmentation.

$$HD_{95}(Y, P) = \max(d_{95}(Y, P), d_{95}(P, Y)) \quad (9)$$

Here,  $d_{95}(Y, P)$  is the maximum 95th percentile distance between the ground truth and predicted voxels, and  $d_{95}(P, Y)$  is the maximum 95th percentile distance between the predicted and ground truth voxels.

### Dice Similarity Coefficient (DSC)

The Dice Similarity Coefficient (DSC) measures the similarity between two sets, returning values from 0 to 1, with 1 representing perfect similarity. It is computed using the following formula:

$$DSC(G, P) = \frac{2|G \cap P|}{|G| + |P|} = \frac{2 \sum_{i=1}^I G_i P_i}{\sum_{i=1}^I G_i + \sum_{i=1}^I P_i} \quad (10)$$

Where  $G$  is the set of real results,  $P$  refers to the set of predicted results,  $G_i$  and  $P_i$  represent the true and predicted values of the voxel  $i$ , respectively, and  $I$  is the number of voxels.

## 8. Results

In this section, we evaluate AMBER-AFNO on three benchmark datasets: ACDC, Synapse, and BraTS. Performance is compared against state-of-the-art CNN and transformer models, including UNet (Ronneberger et al., 2015), nnU-Net (Isensee et al., 2021), TransUNet (Chen et al., 2021), SwinUNet (Cao et al., 2023), UNETR (Hatamizadeh et al., 2021), UNETR++ (Shaker et al., 2024), MISSFormer (Huang et al., 2023), Swin-UNETR (Hatamizadeh et al., 2022), nnFormer (Zhou et al., 2023), TransBTS (Wang et al., 2021), CoTr (Xie et al., 2021b), PCCTrans (Feng et al., 2024) and LW-CTrans (Kuang et al., 2025).

In particular, LW-CTrans (Kuang et al., 2025) is a lightweight CNN-Transformer baseline designed to reduce parameter count and computational cost while maintaining competitive segmentation accuracy. Its baseline comparison makes it convenient to validate whether our proposed token-mixing approach in the frequency domain can achieve the same or higher segmentation accuracy while maintaining a lightweight design.

We will perform quantitative (Dice, HD95) and qualitative (visual) assessment of the accuracy of the segmentations and contours, as well as the anatomical variability between the datasets.

### 8.1. ACDC Dataset

As shown in Tab. 2 and Tab. 3, AMBER-AFNO ranks first in the overall Dice score (92.85) on the ACDC validation set, with an improvement over UNETR++ (92.83), and significantly fewer parameters (14.77M vs. 66.8M). With nearly 4 times fewer parameters, the proposed model still outperforms UNETR++ in segmentation accuracy, validating that token mixing in the spectral domain can effectively eliminate redundancy in the network.

In contrast to the latest CNN-Transformer model, LW-CTrans (92.62%), AMBER-AFNO achieves a higher Dice score with lower computational complexity (163.27G vs. 275.49G FLOPs). It suggests that frequency-domain token mixing can yield

more effective global context modeling than compressed attention-based hybrid models under similar lightweight settings. At the class level, our approach yields the highest Dice score for the myocardium (90.74) and the second-highest for the right (91.60) and left (96.21) ventricles, demonstrating its overall balanced performance across all cardiac classes.

These results show that, on the ACDC dataset, the AFNO-based encoder coupled with a lightweight SegFormer-style decoder delivers state-of-the-art Dice performance while maintaining a favorable accuracy–efficiency trade-off compared to both heavy transformer-based models and recent lightweight CNN–Transformer architectures.

### 8.2. Synapse Dataset

According to Tab. 4 and Tab. 5, the proposed AMBER-AFNO yields an average Dice of 83.76% on the Synapse validation set, which is ranked 3rd among the methods in comparison. Although UNETR++ and nnFormer gain higher Dice of 87.22% and 86.57%, respectively, both require significantly more model parameters of 42.96M and 150.5M, respectively, in comparison to 14.86M in AMBER-AFNO, indicating a good balance between accuracy and efficiency.

Notably, compared to the state-of-the-art lightweight CNN–Transformer baseline LW-CTrans (73.34%), AMBER-AFNO increases the Dice score by over 10 percentage points with comparable model compactness and computational complexity (161.24G vs. 275.92G FLOPs). This implies that frequency-domain global token mixing is more scalable than drastic convolution–attention compression when multiple anatomical classes and intricate organ boundaries are present.

On the Synapse dataset, the task is more difficult than ACDC and BraTS, as the model must learn the probability distributions of multiple abdominal organs with varying shapes and sizes. On this task, our AFNO-based encoder with fewer parameters still achieves consistent performance, demonstrating a favorable trade-off between representational capacity and computational cost for modeling global contextual information in the spectral domain. Therefore, this evidence further supports the use of AMBER-AFNO for 3D multi-organ segmentation under the low-latency constraint.

**Table 2**

Dice scores (%) for the right ventricle (RV), myocardium (Myo), and left ventricle (LV), together with the mean DSC (%) on the ACDC validation set. **Bold** values denote the best result in each column, while underlined values denote the second best.

Methods	RV	Myo	LV	DSC
TransUNet(Chen et al., 2021)	88.86	84.54	95.73	89.71
Swin-UNet(Cao et al., 2023)	88.55	85.62	95.83	90.00
UNETR(Hatamizadeh et al., 2021)	85.29	86.52	94.02	86.61
MISSFormer(Huang et al., 2023)	86.36	85.75	91.59	87.90
nnFormer(Zhou et al., 2023)	90.94	89.58	95.65	92.06
UNETR++(Shaker et al., 2024)	<b>91.89</b>	<u>90.61</u>	96.00	<u>92.83</u>
PCCTrans(Feng et al., 2024)	90.55	90.57	<u>96.22</u>	92.45
LW-CTrans (Kuang et al., 2025)	91.07	90.50	<b>96.24</b>	92.62
<b>AMBER-AFNO (ours)</b>	<u>91.60</u>	<b>90.74</b>	96.21	<b>92.85</b>

**Table 3**

Comparison on ACDC. AMBER-AFNO achieves best segmentation results(DSC), while being efficient (Params in millions)

Methods	Params	FLOPs	DSC
UNETR++(Shaker et al., 2024)	66.8	43.71	92.83
LW-CTrans (Kuang et al., 2025)	4.42	275.49	92.62
<b>AMBER-AFNO (ours)</b>	14.77	163.27	<b>92.85</b>

### 8.3. BraTS Dataset

From Tab. 6, we see that AMBER-AFNO yields the best mean Dice score on the BraTS validation set, with 82.82%, which is slightly better than UNETR++ with 82.75%. This is a strong result, as UNETR++ is one of the heavier transformer baselines that yield the best performance.

Specifically, AMBER-AFNO achieves the highest Dice score (80.33%) in the Enhancing Tumor (ET) region, which is generally the most challenging subregion due to its unclear boundary and limited spatial coverage. This implies that spectral-domain global mixing can model long-range contextual information while retaining the ability to detect

**Table 4**

D

Dice scores (%) for eight abdominal organs and HD95 (mm) on the Synapse validation set. **Bold** values denote the best result in each column, while underlined values denote the second best.

Methods	Spl	RKid	LKid	Gal	Liv	Sto	Aor	Pan	HD95	DSC
U-Net(Ronneberger et al., 2015)	86.67	68.60	77.77	69.72	93.43	75.58	89.07	53.98	–	76.85
TransUNet(Chen et al., 2021)	85.08	77.02	81.87	63.16	94.08	75.62	87.23	55.86	31.69	77.49
Swin-UNet(Cao et al., 2023)	90.66	79.61	83.28	66.53	94.29	76.60	85.47	56.58	21.55	79.13
UNETR(Hatamizadeh et al., 2021)	85.00	84.52	85.60	56.30	94.57	70.46	89.80	60.47	18.59	78.35
MISSFormer(Huang et al., 2023)	91.92	82.00	85.21	68.65	94.41	80.81	86.99	65.67	18.20	81.96
nnFormer(Zhou et al., 2023)	90.51	86.25	86.57	<u>70.17</u>	<b>96.84</b>	<b>86.83</b>	<u>92.04</u>	<b>83.35</b>	10.63	<u>86.57</u>
Swin-UNETR(Hatamizadeh et al., 2022)	<u>95.37</u>	<u>86.26</u>	86.99	66.54	95.72	77.01	91.12	68.80	<u>10.55</u>	83.48
UNETR++(Shaker et al., 2024)	<b>95.77</b>	<b>87.18</b>	<b>87.54</b>	<b>71.25</b>	<u>96.42</u>	<u>86.01</u>	<b>92.52</b>	<u>81.10</u>	<b>7.53</b>	<b>87.22</b>
PCCTrans(Feng et al., 2024)	88.84	82.64	85.49	68.79	93.45	71.88	86.59	66.31	17.10	80.50
LW-CTrans (Kuang et al., 2025)	89.48	82.84	85.11	37.52	93.60	63.10	81.75	53.30	31.47	73.34
<b>AMBER-AFNO (ours)</b>	87.82	<u>86.26</u>	<u>87.36</u>	61.33	96.02	80.50	91.42	79.36	16.96	83.76

**Table 5**

Comparison on Synapse. AMBER-AFNO achieves the third best segmentation results (DSC), while being more efficient (Params in millions)

Methods	Params	FLOPs	DSC
TransUNet(Chen et al., 2021)	96.07	88.91	77.49
UNETR(Hatamizadeh et al., 2021)	92.49	75.76	78.35
Swin-UNet(Cao et al., 2023)	62.83	384.2	83.48
nnFormer(Zhou et al., 2023)	150.5	213.4	86.57
UNETR++(Shaker et al., 2024)	42.96	47.98	<b>87.22</b>
LW-CTrans (Kuang et al., 2025)	4.42	275.92	73.34
<b>AMBER-AFNO (ours)</b>	14.86	161.24	83.76

**Table 6**

Dice scores (%) for Whole Tumor (WT), Enhancing Tumor (ET), Tumor Core (TC), and overall DSC (%), together with HD95 (mm) on the BraTS validation set. **Bold** values denote the best result in each column, while underlined values denote the second best.

Methods	WT	ET	TC	DSC	HD95
UNETR (Hatamizadeh et al., 2021)	90.35	76.30	77.02	81.22	6.61
TransBTS (Wang et al., 2021)	90.91	77.86	76.10	81.62	5.80
Swin-UNETR (Cao et al., 2023)	91.12	77.65	<u>78.41</u>	82.39	5.33
CoTr (Xie et al., 2021b)	91.01	77.52	<u>77.43</u>	81.99	5.78
nnU-Net (Isensee et al., 2021)	91.21	77.96	78.05	82.41	5.58
nnFormer (Zhou et al., 2023)	<u>91.23</u>	77.84	77.91	82.34	<u>5.18</u>
UNETR++ (Shaker et al., 2024)	91.27	<u>78.39</u>	<b>78.60</b>	<u>82.75</u>	<b>5.05</b>
LW-CTrans (Kuang et al., 2025)	89.58	73.83	75.38	79.60	6.91
<b>AMBER-AFNO (ours)</b>	90.86	<b>80.33</b>	77.26	<b>82.82</b>	6.42

detailed tumour structure.


Compared with the lightweight CNN-Transformer baseline LW-CTrans (79.60%), AMBER-AFNO achieved a more than 3% increase in segmentation accuracy with a small number of parameters. This implies that frequency-domain token mixing works well for multimodal and multi-region segmentation tasks, further validating the generalizability of the proposed framework across various clinical applications.

#### 8.4. Discussion: Performance vs Efficiency Trade-off

The experimental results in Tables 2, 3, 4, 5, and 6 clearly show that our AMBER-AFNO model consistently outperforms both heavy transformer models and recent lightweight CNN-Transformer models in key metrics.

On the ACDC dataset, AMBER-AFNO achieves the best Dice Similarity Coefficient (92.85%), which is superior to UNETR++ (92.83%) by a small margin with much fewer parameters (14.77M vs. 81.55M). Compared to the light-weight model LW-CTrans (92.62%), AMBER-AFNO improves the performance with much lower FLOPs (163.27G vs. 275.49G). This shows that the proposed spectral-domain token mixing could yield stronger representation power than current CNN-Transformer hybrids for lightweight models.

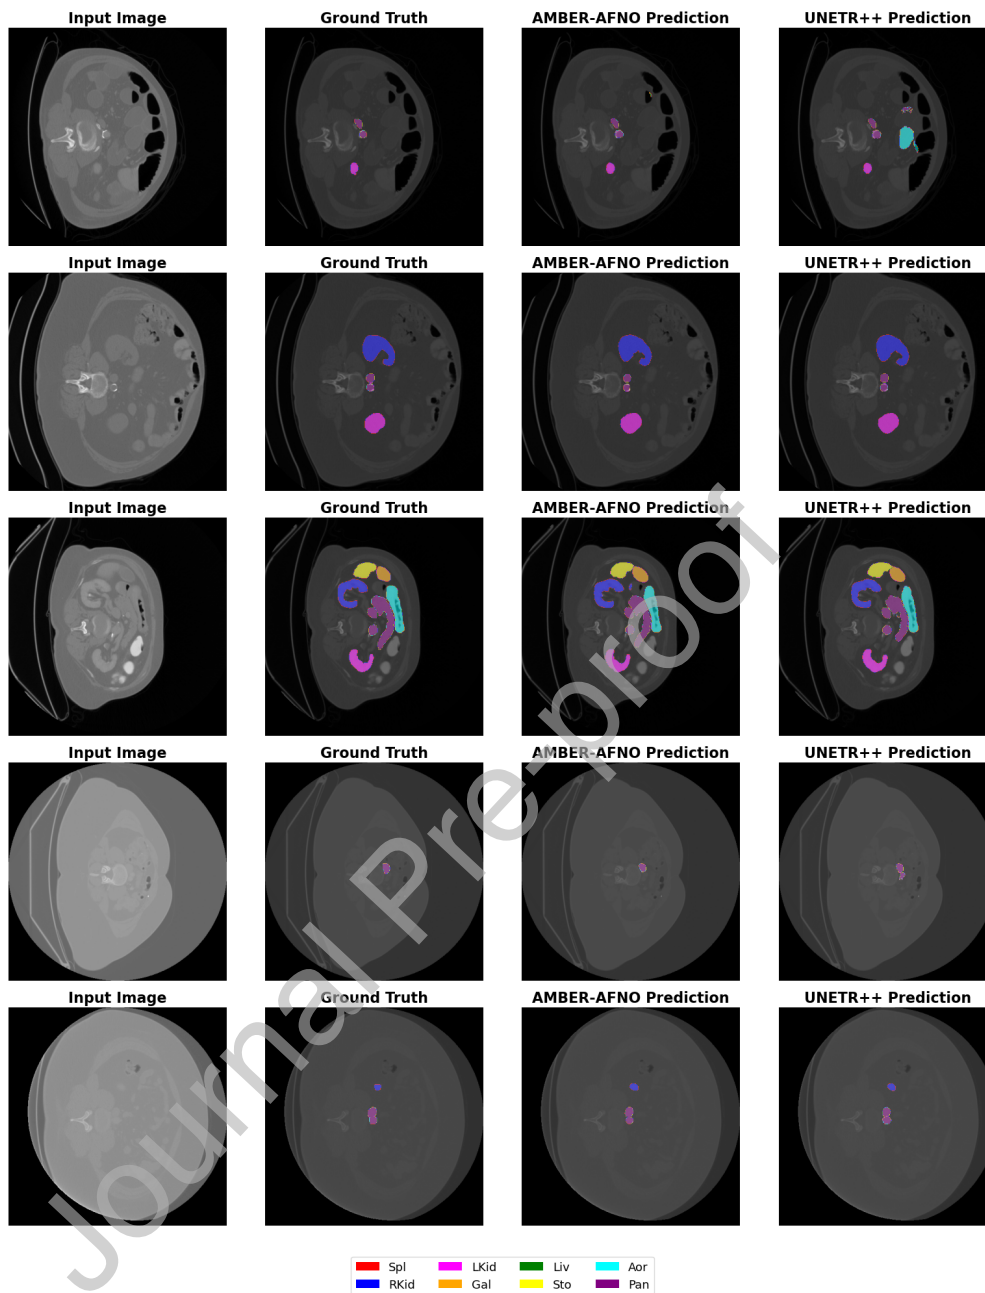
On the multi-organ Synapse dataset, which consists of eight abdominal organs and higher anatomical diversity, AMBER-AFNO obtains a mean Dice of 83.76, ranking in third place. Though larger models, such as UNETR++ and nnFormer, achieve higher absolute Dice scores, they require significantly more model parameters. Furthermore, relative to LW-CTrans (73.34), AMBER-AFNO significantly



**Figure 3:** Qualitative comparison of AMBER-AFNO and UNETR++ segmentation predictions on representative ACDC validation samples. In the legend, **RV** denotes the Right Ventricular Cavity, **Myo** the Myocardium, and **LV** the Left Ventricular Cavity.

improves performance by more than 10 percentage points with a lower parameter count. Meanwhile, it incurs lower computational cost (161.24G

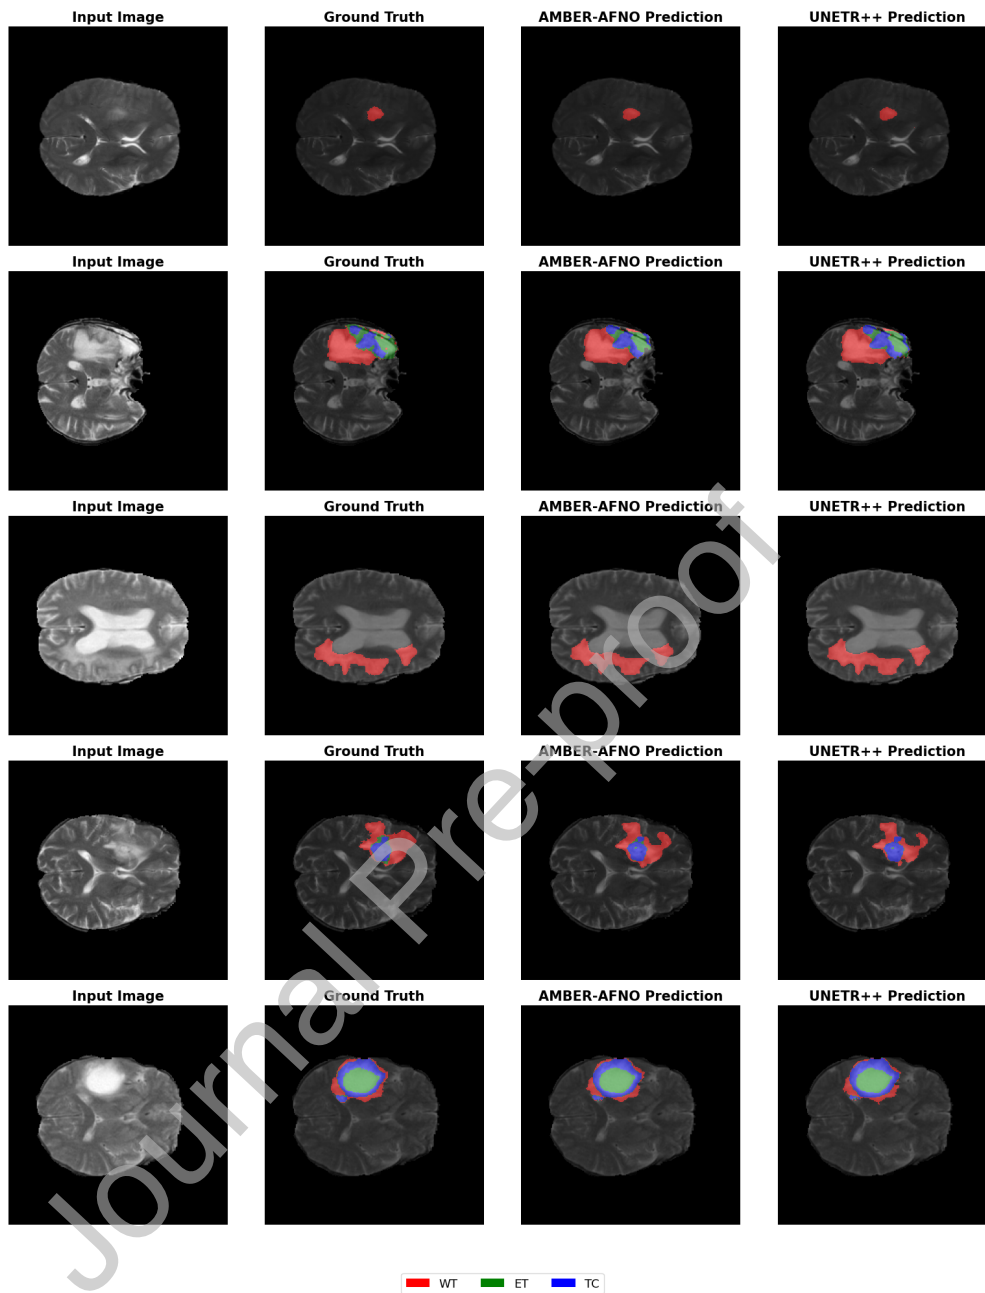
vs. 275.92G FLOPs), indicating that frequency-domain token mixing is more suitable than pure convolution-attention fusion for multi-class segmentation.



**Figure 4:** Visual Comparison of AMBER-AFNO and UNETR++ Model's Prediction on Synapse Dataset. In the legend, **Spl** denotes the Spleen, **RKid** the Right Kidney, **LKid** the Left Kidney, **Gal** the Gallbladder, **Liv** the Liver, **Sto** the Stomach, **Aor** the Aorta, and **Pan** the Pancreas.

On the challenging BraTS dataset with high intra-class variability and multi-region tumor segmentation task, the proposed AMBER-AFNO gains the

best average Dice (82.82%) among all methods. Specifically, it outperforms the recent UNETR++



**Figure 5:** Visual comparison of AMBER-AFNO and UNETR++ predictions on the BraTS dataset. In the legend, **WT** denotes the Whole Tumor, **ET** the Enhancing Tumor, and **TC** the Tumor Core.

(82.75%) while using a much smaller model. Meanwhile, compared with the lightweight counterpart LW-CTrans (79.60%), the proposed model improves Dice accuracy while remaining lightweight.

Moreover, our method achieves the highest Dice score (80.33%) on the Enhancing Tumor dataset, further verifying its superiority in identifying fine-grained tumor substructures.

From these three comparisons, the one against LW-CTrans is especially interesting because, although optimized for the fewest parameters and lowest computational cost, LW-CTrans always yields lower segmentation Dice scores than AMBER-AFNO. Instead, AMBER-AFNO maintains a moderate model size while always offering higher Dice scores on ACDC, Synapse, and BraTS.

This performance–efficiency trade-off demonstrates the effectiveness of our central hypothesis. Specifically, AMBER-AFNO removes token-to-token interactions and retains only global context modeling by using Adaptive Fourier Neural Operators instead of the quadratic self-attention mechanism, leading to a remarkable decrease in parameter growth and computational cost, without compromising segmentation accuracy.

All this is obtained without resorting to decoder branches, multi-scale fusion hierarchies, or deep stacks of attention modules. The overall architecture’s simplicity and modularity allow straightforward deployment across different datasets with negligible task-specific adaptations. All in all, these characteristics make AMBER-AFNO a workable and scalable solution compared to heavy transformer-based models and existing light CNN-Transformer models.

Tab. 7 reports the memory profile and the latency of AMBER-AFNO across four hardware configurations combining two modern GPUs (NVIDIA H100 and NVIDIA L40) and two different CPUs (AMD EPYC 7413-24c and AMD EPYC 9534-64c). AMBER-AFNO exhibits an extremely light memory footprint, requiring only 2.96 GB of GPU memory for full-resolution 3D inference, making the model deployable even on mid-range or shared accelerator environments. In terms of speed, the model achieves sub-100 ms latency on the NVIDIA L40 and remains below 160 ms on the H100, demonstrating high throughput for the BraTS 128<sup>3</sup> volumetric inputs. CPU-only execution is naturally slower, but still processes each volume within 3–3.5 s, which remains practical for non-real-time clinical workflows. Overall, the measurements highlight the efficiency of Amber-AFNO, combining low memory demand with fast inference suitable for scalable deployment.

As discussed in Section 9, along with the AMBER-AFNO, we have also performed an ablation study

on the lightest version of AMBER-AFNO with embedding dimension [32, 64, 128, 256]. Fig. 6 shows the efficiency–accuracy trade-off of different models on the ACDC dataset. The mean Dice score (DSC) is plotted against the total number of parameters. AMBER-AFNO achieves the highest DSC (92.85%) while maintaining a moderate number of parameters. AMBER-AFNO (light) and UNETR++ both achieve 92.83%, but AMBER-AFNO (light) uses significantly fewer parameters. Although LW-CTrans has fewer parameters, it does not surpass AMBER-AFNO or AMBER-AFNO (light) in performance.

## 9. Ablation Study

In order to better demonstrate the efficiency of the proposed AFNO, we make the encoder variants trained on ACDC dataset share all hyper-parameters including network depth, number of heads, MLP expansion ratio, deep supervision (see details in Section §3.1), as well as embedding dimension [32, 64, 128, 256]. The only difference between these encoders is the token-mixing block: for **AMBER (MHSA)**, the AFNO blocks are replaced with the conventional multi-head self-attention (MHSA); for **AMBER-AFNO (light)**, we use the proposed AFNO block with the above-mentioned setting.

Despite having less than half of the parameters and the FLOPs count of the MHSA module counterpart (**8.7M, 58.29 GFLOPs**) versus (**19.01M, 132.07 GFLOPs**), **AMBER-AFNO (light)** achieves a Dice score of **92.83%**, which is on par with the significantly larger UNETR++ (66.8M parameters).

However, the MHSA-based model achieves a relatively lower mean Dice score of **92.03%**, representing a performance decrease of about 0.8% while being more complex. Therefore, the superiority of spectral token mixing over the quadratic self-attention is straightforwardly illustrated here. Notably, a comparison with the latest, lightweight CNN-Transformer model, **LW-CTrans**, provides a clearer view of the trade-off between accuracy and computational complexity. Though LW-CTrans performs well in terms of the number of parameters, **4.42M**, it suffers from significantly higher computational complexity, **275.49 GFLOPs**, and a decreased Dice score of 92.62. Specifically, the proposed **AMBER-AFNO (light)** enjoys a higher Dice

**Table 7**

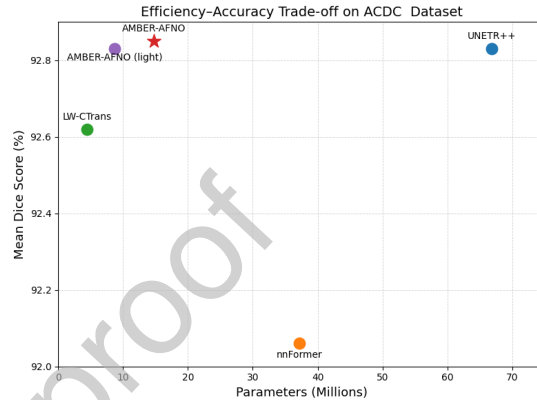
Computational performance of Amber-AFNO on the BraTS tumor input ( $1 \times 128 \times 128 \times 128$ ). GPU T. and CPU T. denote the average forward-pass latency over 20 runs. Mem indicates the peak GPU memory usage during inference.

GPU	CPU	Params (M)	FLOPs (G)	Mem (GB)	GPU T. (ms)	CPU T. (ms)
NVIDIA H100 PCIe	AMD EPYC 7413 (24c)	9.84	122.91	2.96	159.6	3305.0
NVIDIA L40	AMD EPYC 9534 (64c)	9.84	122.91	2.96	84.2	3453.2

score of 92.83 with almost  $5\times$  lower FLOPs compared to LW-CTrans, indicating a better performance-complexity ratio. AMBER-AFNO achieves the highest Dice score of 92.85% while maintaining a substantially smaller parameter count than UNETR++, although with higher FLOPs due to the full 3D spectral mixing.

The number of parameters, FLOPs, and DSC of all models are listed in Tab. 8. In general, these results demonstrate that the proposed AFNO-based frequency-domain token mixing offers a better trade-off between efficiency and accuracy than conventional MHSA and recent CNN/transformer models for 3-D medical image segmentation. Fig. 6 visually demonstrates how AMBER-AFNO performs better compared to other state-of-the-art models in terms of dice score and total number of parameters.

In order to test the robustness and reproducibility of the proposed model based on the AFNO, we conducted a sensitivity analysis where we tested different scenarios, including the impact of the number of epochs, learning rate, batch size, embedding dimension, dropout ratio, number of blocks in each stage, and deep supervision. The outcomes of the sensitivity analysis are summarized in the table. We can see from the table that the Dice score remains within the range of 0.91-0.93 for a wide range of hyperparameters, which suggests that the proposed model is not sensitive to the hyperparameters (within a reasonable range). Using more blocks or a relatively larger embedding dimension can yield marginal improvements in the Dice score, but it will also increase computational cost. In addition, deep supervision or a relatively small learning rate can speed up the training but has little impact on the Dice score. These results suggest that the proposed model is robust and can be easily generalized to other datasets (i.e., we do not need to re-tune many hyperparameters).



**Figure 6:** Efficiency-accuracy trade-off of different models on the ACDC Dataset

*Notes on parameters:* **Deep Supervision** indicates whether additional prediction heads are used, enabling loss computation at multiple resolutions. **Batch Size** denotes the number of 3D volumes processed per iteration; **Epochs** represents total training cycles; **LR** is the initial learning rate; **Embed Dim.** lists the embedding dimensions across encoder stages; **Dropout** specifies the dropout probability applied in the MLP layers; and **Blocks/Stage** refers to the number of AFNO blocks used in each hierarchical encoder stage.

## 10. Conclusions

In this study, we introduced AMBER-AFNO, a novel framework for 3D medical image segmentation. It uses Adaptive Fourier Neural Operators (AFNO) instead of quadratic self-attention. Unlike other approaches that approximate or sparsify attention, we directly reformulate the problem of modeling global context. We do this by removing token-to-token interactions and ensuring tokens are mixed in the spectral domain. Our design delivers

**Table 8**

Efficiency–accuracy trade-off on the ACDC validation set. AFNO delivers state-of-the-art Dice scores with markedly fewer parameters. Best results are highlighted in **bold**, and second-best results are underlined.

Method	Params	FLOPs	DSC (%)
UNETR++(Shaker et al., 2024)	66.8	43.71	<u>92.83</u>
nnFormer(Zhou et al., 2023)	37.20	399.25	92.06
LW-CTrans (Kuang et al., 2025)	4.42	275.49	92.62
AMBER-AFNO	14.77	163.27	<b>92.85</b>
AMBER-AFNO (light)	8.70	58.29	<u>92.83</u>
AMBER (MHSA)	19.01	132.07	92.03

**Table 9**

Sensitivity analysis of key hyperparameters on the ACDC validation set. The model demonstrates stable Dice scores across a wide range of configurations, confirming robustness to variations in training strategy and architecture depth.

Experiments	Deep Supervision	Batch Size	Epochs	LR	Embed Dim.	Dropout	Blocks/Stage	Dice
1	True	4	1000	0.0100	[32, 64, 128, 256]	0.1	[2, 2, 4, 2]	<b>0.93</b>
2	False	1	1500	0.0037	[32, 64, 128, 256]	0.1	[3, 4, 6, 3]	0.92
3	True	4	1000	0.0100	[64, 128, 320, 512]	0.0	[3, 4, 6, 3]	<b>0.93</b>
4	False	4	1000	0.0100	[64, 128, 320, 512]	0.0	[3, 4, 6, 3]	0.91
5	False	4	1000	0.0100	[64, 128, 320, 512]	0.0	[3, 4, 6, 3]	<b>0.93</b>
6	False	4	1000	0.0053	[64, 128, 320, 512]	0.0	[3, 6, 36, 3]	0.92

quasi-linear complexity and linear memory growth with respect to volume size. This effectively overcomes a major drawback of transformer-based 3D segmentation.

To the best of our knowledge, no prior work has integrated AFNO-based spectral token mixing into a transformer framework for 3D medical image segmentation. Other efficient CNN–Transformer variants mainly seek to reduce the complexity of attention or convolutional blocks. In contrast, AMBER-AFNO removes quadratic self-attention from the transformer encoder. It still preserves global receptive-field modeling through frequency truncation.

On three public benchmarks (ACDC, Synapse, BraTS), AMBER-AFNO matches or surpasses heavy transformer models (UNETR++ (Shaker et al., 2024), nnFormer (Zhou et al., 2023)) with substantially fewer parameters (nearly 78% less on the ACDC dataset). Compared to the light CNN–Transformer LW-CTrans (Kuang et al., 2025), AMBER-AFNO achieves higher Dice scores. It also has a smaller model size and lower FLOPs on ACDC and Synapse. This demonstrates a more favorable

balance between accuracy and efficiency versus quadratic or compressed attention-based methods.

Although performance on the Synapse dataset suggests that further enhancement of local structural modeling may be beneficial, the proposed framework establishes a new direction for efficient volumetric segmentation: global context modeling through spectral operators rather than attention matrices. Future work will explore hybrid spectral–spatial strategies, improved multi-scale integration, and transfer learning for domain adaptation.

In summary, AMBER-AFNO bridges the gap between computationally intensive transformer architectures and lightweight CNN–Transformer models by introducing an alternative global information-mixing strategy. Its compact parameterization and scalable complexity make it well-suited for deployment in resource-constrained clinical environments.

## Acknowledgments

This work has been funded by project code PIR01\_00011 “IBISCO”, PON 2014–2020, for all three entities (INFN, UNINA, and CNR).

We acknowledge the use of the ADHOC (Astrophysical Data HPC Operating Center) resources, within the project "Strengthening the Italian Leadership in ELT and SKA (STILES)", proposal nr. IR0000034, admitted and eligible for funding from the funds referred to in the D.D. prot. no. 245 of August 10, 2022 and D.D. 326 of August 30, 2022, funded under the program "Next Generation EU" of the European Union, "Piano Nazionale di Ripresa e Resilienza" (PNRR) of the Italian Ministry of University and Research (MUR), "Fund for the creation of an integrated system of research and innovation infrastructures", Action 3.1.1 "Creation of new IR or strengthening of existing IR involved in the Horizon Europe Scientific Excellence objectives and the establishment of networks".

### **Credit Author Statement New**

Andrea Dosi: Conceptualization, Methodology, Software, Writing- Original draft preparation Semanto Mondal.: Methodology , Software, Writing- Original draft preparation. Rajib Chandra Ghosh: Methodology , Software, Writing- Original draft preparation. Massimo Brescia: Supervision. Giuseppe Longo: Supervision.

### **Authors and their corresponding ORCID IDs**

• Andrea Dosi — 0000-0002-5943-6867 • Semanto Mondal — 0009-0000-2306-4478 • Rajib Chandra Ghosh — 0009-0001-1137-3465 • Massimo Brescia — 0000-0001-9506-5680 • Giuseppe Longo — 0000-0002-9182-8414

### **Declaration of interests**

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## References

- Bai, Q., Luo, X., Wang, Y., Wei, T., 2024. Dhynet: A dual-branch hybrid reinforcement network for semantic segmentation of remote sensing images. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 17, 4176–4193. doi:10.1109/JSTARS.2024.3357216.
- Bernard, O., Lalonde, A., Zotti, C., Cervenansky, F., Yang, X., Heng, P.A., Cetin, I., Lekadir, K., Camara, O., Gonzalez Ballester, M.A., Sanroma, G., Napel, S., Petersen, S., Tziritas, G., Grinias, E., Khened, M., Kollerathu, V.A., Krishnamurthi, G., Rohé, M.M., Pennec, X., Sermesant, M., Isensee, F., Jäger, P., Maier-Hein, K.H., Full, P.M., Wolf, I., Engelhardt, S., Baumgartner, C.F., Koch, L.M., Wolterink, J.M., Išgum, I., Jang, Y., Hong, Y., Patravali, J., Jain, S., Humbert, O., Jodoin, P.M., 2018. Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: Is the problem solved? *IEEE Transactions on Medical Imaging* 37, 2514–2525. doi:10.1109/TMI.2018.2837502.
- Cao, H., Wang, Y., Chen, J., Jiang, D., Zhang, X., Tian, Q., Wang, M., 2023. Swin-unet: Swin-like pure transformer for medical image segmentation, in: *Computer Vision – ECCV 2022 Workshops*, Springer Nature Switzerland, Cham. pp. 205–218.
- Chen, J., Fang, J., Zhou, Z., Yang, C., Xie, L., Zhang, X., Shen, W., Tian, Q., 2024. Segment anything in 3d with radiance fields. URL: <https://arxiv.org/abs/2304.12308>, arXiv:2304.12308.
- Chen, J., Lu, Y., Yu, Q., Luo, X., Adeli, E., Wang, Y., Lu, L., Yuille, A.L., Zhou, Y., 2021. Transunet: Transformers make strong encoders for medical image segmentation. URL: <https://arxiv.org/abs/2102.04306>, arXiv:2102.04306.
- Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L., 2018. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 40, 834–848. doi:10.1109/TPAMI.2017.2699184.
- Chen, M., Zhang, Y., Wang, L., Liu, J., Li, Q., 2024. Sau-net: A novel network for building extraction from high-resolution remote sensing images by reconstructing fine-grained semantic features. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 17, 6747–6761. doi:10.1109/JSTARS.2024.3371427.
- Choromanski, K., Likhoshesterov, V., Dohan, D., Song, X., Gane, A., Sarlos, T., Hawkins, P., Davis, J., Mohiuddin, A., Kaiser, L., Belanger, D., Colwell, L., Weller, A., 2022. Rethinking attention with performers. URL: <https://arxiv.org/abs/2009.14794>, arXiv:2009.14794.
- Dosi, A., Brescia, M., Cavuoti, S., D’Aniello, M., Delli Veneri, M., Donadio, C., Ettari, A., Longo, G., Rownok, A., Sannino, L., Zampella, M., 2025. AMBER: advanced segmenter for multi-band image segmentation—an application to hyperspectral imaging. *Neural Computing and Applications* 37, 17273–17291. URL: <https://doi.org/10.1007/s00521-025-11315-1>, doi:10.1007/s00521-025-11315-1.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houshy, N., 2021. An image is worth 16x16 words: transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 URL: <https://arxiv.org/abs/2010.11929>, arXiv:2010.11929.
- Feng, Y., Su, J., Zheng, J., Zheng, Y., Zhang, X., 2024. A parallelly contextual convolutional transformer for medical image segmentation. *Biomedical Signal Processing and Control* 98, 106674. URL: <https://www.sciencedirect.com/science/article/pii/S1746809424007328>, doi:10.1016/j.bspc.2024.106674.
- Florkow, M.C., Willemsen, K., Mascarenhas, V.V., Oei, E.H.G., van Stralen, M., Seevinck, P.R., 2022. Magnetic resonance imaging versus computed tomography for three-dimensional bone imaging of musculoskeletal pathologies: A review. *Journal of Magnetic Resonance Imaging* 56, 11–34. doi:10.1002/jmri.28067. epub 2022-01-19.
- Gu, Z., Cheng, J., Fu, H., Zhou, K., Hao, H., Zhao, Y., Zhang, T., Gao, S., Liu, J., 2019. Ce-net: Context encoder network for 2d medical image segmentation. *IEEE Transactions on Medical Imaging* 38, 2281–2292. doi:10.1109/TMI.2019.2903562.
- Guibas, J., Mardani, M., Li, Z., Tao, A., Anandkumar, A., Catanzaro, B., 2021. Adaptive fourier neural operators: Efficient token mixers for transformers. arXiv preprint arXiv:2111.13587 doi:10.48550/arXiv.2111.13587, arXiv:2111.13587.
- Hatamizadeh, A., Nath, V., Tang, Y., Yang, D., Roth, H.R., Xu, D., 2022. Swin unetr: Swin transformers for semantic segmentation of brain tumors in mri images, in: *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries*, Springer International Publishing, Cham. pp. 272–284.
- Hatamizadeh, A., Yang, D., Roth, H., Xu, D., 2021. Unetr: Transformers for 3d medical image segmentation. arXiv preprint arXiv:2103.10504 arXiv:2103.10504.
- He, Z., Li, X., Lv, N., Chen, Y., Cai, Y., 2024. Retinal vascular segmentation network based on multi-scale adaptive feature fusion and dual-path upsampling. *IEEE Access* 12, 48057–48067. doi:10.1109/ACCESS.2024.3383848.
- Hou, Q.Y., Zhao, Y.J., Zhou, H., Li, Y., Yu, L., 2022. Istdu-net: Infrared small-target detection u-net. *IEEE Geoscience and Remote Sensing Letters* doi:10.1109/LGRS.2022.3141584.
- Huang, X., Deng, Z., Li, D., Yuan, X., 2023. Missformer: An effective transformer for 2d medical image segmentation. *IEEE Transactions on Medical Imaging* 42, 1484–1494. doi:10.1109/TMI.2022.3230943.
- Isensee, F., Jaeger, P.F., Kohl, S.A.A., Petersen, J., Maier-Hein, K.H., 2021. nnu-net: A self-configuring method for deep learning-based biomedical image segmentation. *Nat. Methods* 18, 203–211. doi:10.1038/s41592-020-01008-z.
- Jiang, C., Wang, Y., Yuan, Q., Qu, P., Li, H., 2025. A 3d medical image segmentation network based on gated attention blocks and dual-scale cross-attention mechanism. *Scientific Reports* 15. doi:10.1038/s41598-025-90339-y.
- Kuang, H., Wang, Y., Tan, X., Yang, J., Sun, J., Liu, J., Qiu, W., Zhang, J., Zhang, J., Yang, C., Wang, J., Chen, Y., 2025. Lw-trans: A lightweight hybrid network of cnn and transformer for 3d medical image segmentation. *Medical Image Analysis* 102, 103545. URL: <https://www.sciencedirect.com/science/article/pii/S1361841525000921>, doi:10.1016/j.media.2025.103545.

- Lan, C.F., Zhang, L., Wang, Y.Q., Liu, C.D., 2022. Research on improved dnn and multiresu\_net network speech enhancement effect. *Multimedia Tools and Applications* 81, 26163–26184. doi:10.1007/s11042-022-12929-6.
- Landman, B., Xu, Z., Iglesias, J.E., Styner, M., Langerak, T., Klein, A., 2015. Multi-atlas labeling beyond the cranial vault. arXiv preprint doi:10.7303/syn3193805.
- Lee-Thorp, J., Ainslie, J., Eckstein, I., Ontanon, S., 2022. Fnet: Mixing tokens with fourier transforms. URL: <https://arxiv.org/abs/2105.03824>, arXiv:2105.03824.
- Li, B., Wang, Y., Zhang, X., Chen, J., Liu, L., 2022a. Rt-unet: an advanced network based on residual network and transformer for medical image segmentation. *International Journal of Intelligent Systems* 37, 8565–8582. doi:10.1002/int.22956.
- Li, R., Wang, X., Huang, G., Yang, W., Zhang, K., Gu, X., Tran, S.N., Garg, S., Alty, J., Bai, Q., 2022b. A comprehensive review on deep supervision: Theories and applications. URL: <https://arxiv.org/abs/2207.02376>, arXiv:2207.02376.
- Lin, A., Chen, B., Xu, J., Zhang, Z., Lu, G., Zhang, D., 2022. Ds-transunet: Dual swin transformer u-net for medical image segmentation. *IEEE Transactions on Instrumentation and Measurement* 71, 1–15. doi:10.1109/TIM.2022.3178991.
- Liu, H., Huo, G., Li, Q., Guan, X., Tseng, M.L., 2023. Multiscale lightweight 3d segmentation algorithm with attention mechanism: Brain tumor image segmentation. *Expert Syst. Appl.* 214. URL: <https://doi.org/10.1016/j.eswa.2022.119166>, doi:10.1016/j.eswa.2022.119166.
- Menze, B.H., Jakab, A., Bauer, S., et al., 2015. The multimodal brain tumor image segmentation benchmark (brats). *IEEE Transactions on Medical Imaging* 34, 1993–2024. doi:10.1109/TMI.2014.2377694.
- Milletari, F., Navab, N., Ahmadi, S.A., 2016. V-net: Fully convolutional neural networks for volumetric medical image segmentation, in: 2016 Fourth International Conference on 3D Vision (3DV), pp. 565–571. doi:10.1109/3DV.2016.79.
- Pang, Y., Liang, J., Huang, T., Chen, H., Li, Y., Li, D., Huang, L., Wang, Q., 2024. Slim unet: Scale hybrid transformers to efficient 3d medical image segmentation under limited computational resources. *IEEE Transactions on Medical Imaging* 43, 994–1005. doi:10.1109/TMI.2023.3326188.
- Ronneberger, O., Fischer, P., Brox, T., 2015. U-net: Convolutional networks for biomedical image segmentation, in: *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pp. 234–241. doi:10.1007/978-3-319-24574-4\_28.
- Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C., 2018. Mobilenetv2: Inverted residuals and linear bottlenecks, in: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 4510–4520. doi:10.1109/CVPR.2018.00474.
- Shaker, A., Maaz, M., Abdul Rasheed, H., Khan, S., Yang, M.H., Khan, F., 2024. Unetr++: Delving into efficient and accurate 3d medical image segmentation. *IEEE Transactions on Medical Imaging* PP, 1–1. doi:10.1109/TMI.2024.3398728.
- Song, H.J., Wang, Y.F., Zeng, S.J., Guo, X.Y., Li, Z.H., 2023. Oau-net: Outlined attention u-net for biomedical image segmentation. *Biomedical Signal Processing and Control* 80, 104038. doi:10.1016/j.bspc.2022.104038.
- Sun, Q., Zhang, Y., Liu, H., Wang, X., Li, J., 2022. Ucr-net: U-shaped context residual network for medical image. *Computers in Biology and Medicine* 149, 106203. doi:10.1016/j.combiomed.2022.106203.
- Tai, F.W.D., Wray, N., Sidhu, R., Hopper, A., McAlindon, M., 2019. Factors associated with oesophagogastric cancers missed by gastroscopy: A case-control study. *Frontline Gastroenterol.* 11. doi:10.1136/flgastro-2019-101217.
- Valanarasu, J.M.J., Sindagi, V.A., Hacihaliloglu, I., Patel, V.M., 2022. Kiu-net: Overcomplete convolutional architectures for biomedical image and volumetric segmentation. *IEEE Transactions on Medical Imaging* 41, 965–976. doi:10.1109/TMI.2021.3130469.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I., 2017. Attention is all you need. arXiv preprint arXiv:1706.03762 URL: <https://arxiv.org/abs/1706.03762>, arXiv:1706.03762.
- Wang, S., Li, B.Z., Khabsa, M., Fang, H., Ma, H., 2020. Linformer: Self-attention with linear complexity. URL: <https://arxiv.org/abs/2006.04768>, arXiv:2006.04768.
- Wang, W., Chen, C., Ding, M., Yu, H., Zha, S., Li, J., 2021. Transbts: Multimodal brain tumor segmentation using transformer, in: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*, Springer International Publishing, Cham, pp. 109–119.
- World Health Organization, 2024. World health statistics 2024: Monitoring health for the sdgs, sustainable development goals. URL: <https://www.who.int/publications/i/item/9789240094703>. accessed: 2025-08-03.
- Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P., 2021a. Segformer: Simple and efficient design for semantic segmentation with transformers. URL: <https://arxiv.org/abs/2105.15203>, arXiv:2105.15203.
- Xie, Y., Zhang, J., Shen, C., Xia, Y., 2021b. Cotr: Efficiently bridging cnn and transformer for 3d medical image segmentation, in: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021*, Springer International Publishing, Cham, pp. 171–180.
- Zhang, J., Liu, Y., Wu, Q., Wang, Y., Liu, Y., Xu, X., Song, B., 2022. Star-shaped window transformer reinforced u-net for medical image segmentation. *Computers in Biology and Medicine* 150, 105954. doi:10.1016/j.combiomed.2022.105954.
- Zhang, X., Zhou, X., Lin, M., Sun, J., 2018. Shufflenet: An extremely efficient convolutional neural network for mobile devices, in: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 6848–6856. doi:10.1109/CVPR.2018.00716.
- Zhao, Q., Wang, H., Wang, G., 2021. Lcov-net: A lightweight neural network for covid-19 pneumonia lesion segmentation from 3d ct images, in: 2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI), pp. 42–45. doi:10.1109/ISBI48211.2021.9434023.
- Zheng, Z., Yang, M., Wang, C., Li, W., Zhou, L., Zhang, Y., Li, J., 2022. Affu-net: Attention feature fusion u-net with hybrid loss for winter jujube crack detection. *Computers and Electronics in Agriculture* 202, 107049. doi:10.1016/j.compag.2022.107049.

## Short Title of the Article

- Zhou, H.Y., Guo, J., Zhang, Y., Han, X., Yu, L., Wang, L., Yu, Y., 2023. nnformer: Volumetric medical image segmentation via a 3d transformer. *IEEE Transactions on Image Processing* 32, 4036–4045. doi:10.1109/TIP.2023.3293771.
- Zhou, Q., Wang, L., Gao, G., Kang, B., Ou, W., Lu, H., 2024. Boundary-guided lightweight semantic segmentation with multi-scale semantic context. *IEEE Transactions on Multimedia* 26, 7887–7900. doi:10.1109/TMM.2024.3372835.
- Zhou, Z., Sodha, V., Siddiquee, M.M.R., Feng, R., Tajbakhsh, N., Gotway, M.B., Liang, J., 2019. Models genesis: Generic autodidactic models for 3d medical image analysis. URL: <https://arxiv.org/abs/1908.06912>, arXiv:1908.06912.
- Zhu, Z., Sun, M., Qi, G., Li, Y., Gao, X., Liu, Y., 2024a. Sparse dynamic volume transunet with multi-level edge fusion for brain tumor segmentation. *Computers in Biology and Medicine* 172, 108284. doi:10.1016/j.combiomed.2024.108284.
- Zhu, Z., Yu, K., Qi, G., Cong, B., Li, Y., Li, Z., Gao, X., 2024b. Lightweight medical image segmentation network with multi-scale feature-guided fusion. *Computers in Biology and Medicine* 182, 109204. doi:10.1016/j.combiomed.2024.109204.