Spatial Prior-Guided Boundary and Region-Aware 2D Lesion Segmentation in Neonatal Hypoxic Ischemic Encephalopathy

Amog Rao[1,*], **Ananya Shukla[1,*]**, Jia Bhargava[1], Yangming Ou[2], Rina Bao[2,†]

[1] Plaksha University, Mohali, India

[2] Boston Children's Hospital and Harvard Medical School, Boston, USA

[*] Equal Contribution, [†] Corresponding Author

Paper Code

Introduction

Hypoxic Ischemic Encephalopathy (HIE) is a brain injury caused by a sudden decrease in oxygen or blood flow to an infant's brain. Affecting **1 to 5 per 1000 term-born infants**, HIE can lead to severe, **long-term neurocognitive deficits,** including cerebral palsy, epilepsy, or even death in **30-50%** of cases, making its accurate and early detection critical for prognosis and treatment.

Spatial and Temporal Variability of HIE Lesions

The normal range of **Apparent Diffusion Coefficient (ADC)** values varies significantly across **different brain regions** and **developmental stages** in neonates. This variability makes it challenging to use fixed ADC thresholds for lesion detection, leading to potential interpretation errors.

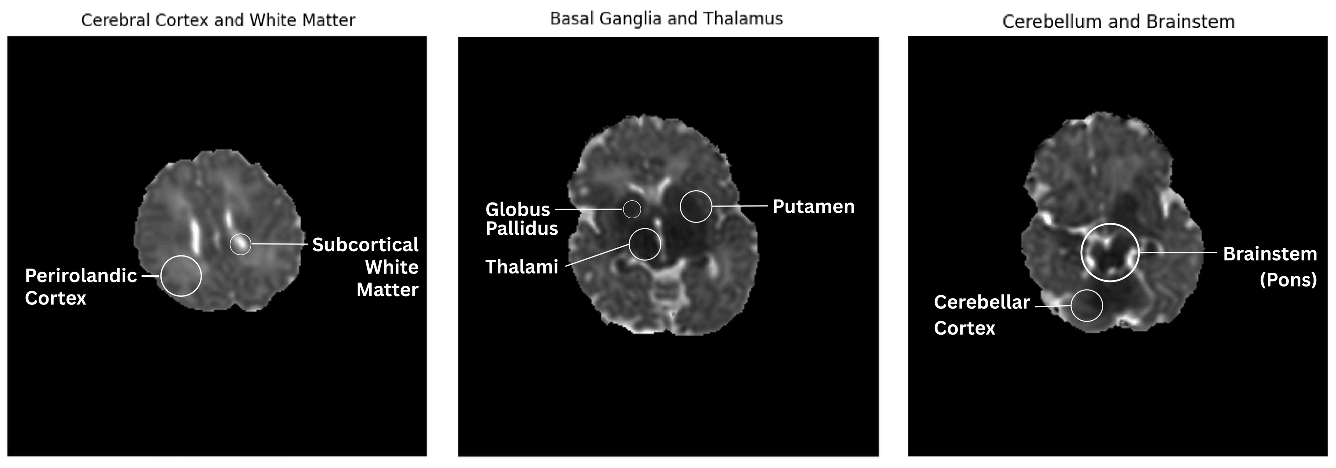

An ADC value of 800 × 10^{-6} mm²/s is normal in cortical gray matter (range: 780–1090), but the same value in white matter (range: 620–790) could indicate a lesion.

Normative Assessment: Z-Score ADC (ZADC) Maps

To overcome this, **Z-Score ADC (ZADC)** maps normalize a patient's ADC against a normative atlas of healthy neonates $Z_{\text{ADC}}(\mathbf{x}) = \frac{I(\mathbf{x}) - \mu_{\phi(\mathbf{x})}}{\sigma_{\phi(\mathbf{x})}}$. This highlights **abnormal diffusion relative to normative variability**, with a threshold of -2 (i.e. ZADC < -2) showing high diagnostic accuracy.

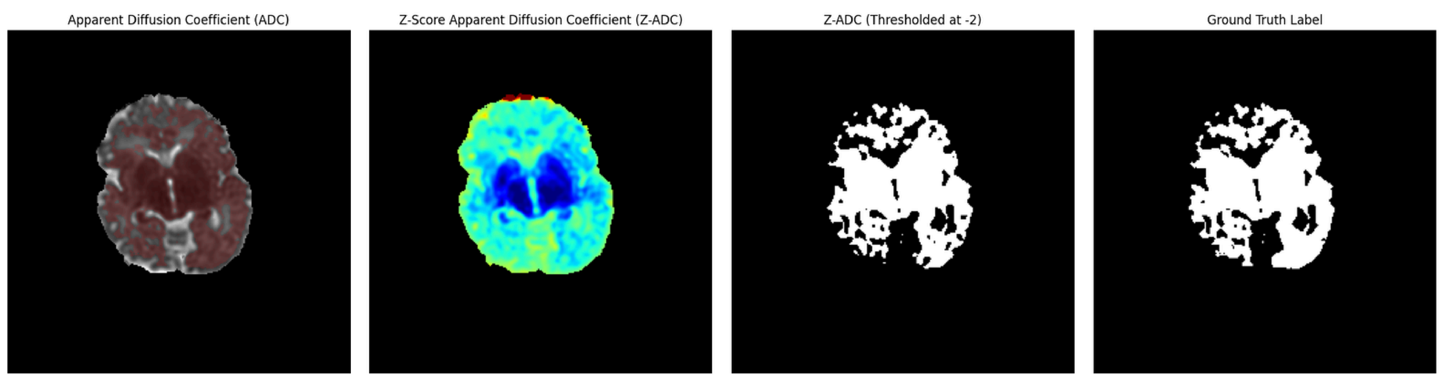

Critical Class Imbalance and Anisotropic Slicing

- The **BONBID-HIE (Bao et al., 2025)** dataset presents that over **55%** of patients have **hyper-acute lesions** affecting **less than 1% (with a median lesion volume of 0.63%)** of the total brain volume.
- Large, **uneven inter-slice gaps (2-6 mm)** in scans reduce **slice-to-slice correlation**, making a 2D slice-based framework more appropriate than 3D methods, which introduce **spurious continuity.**

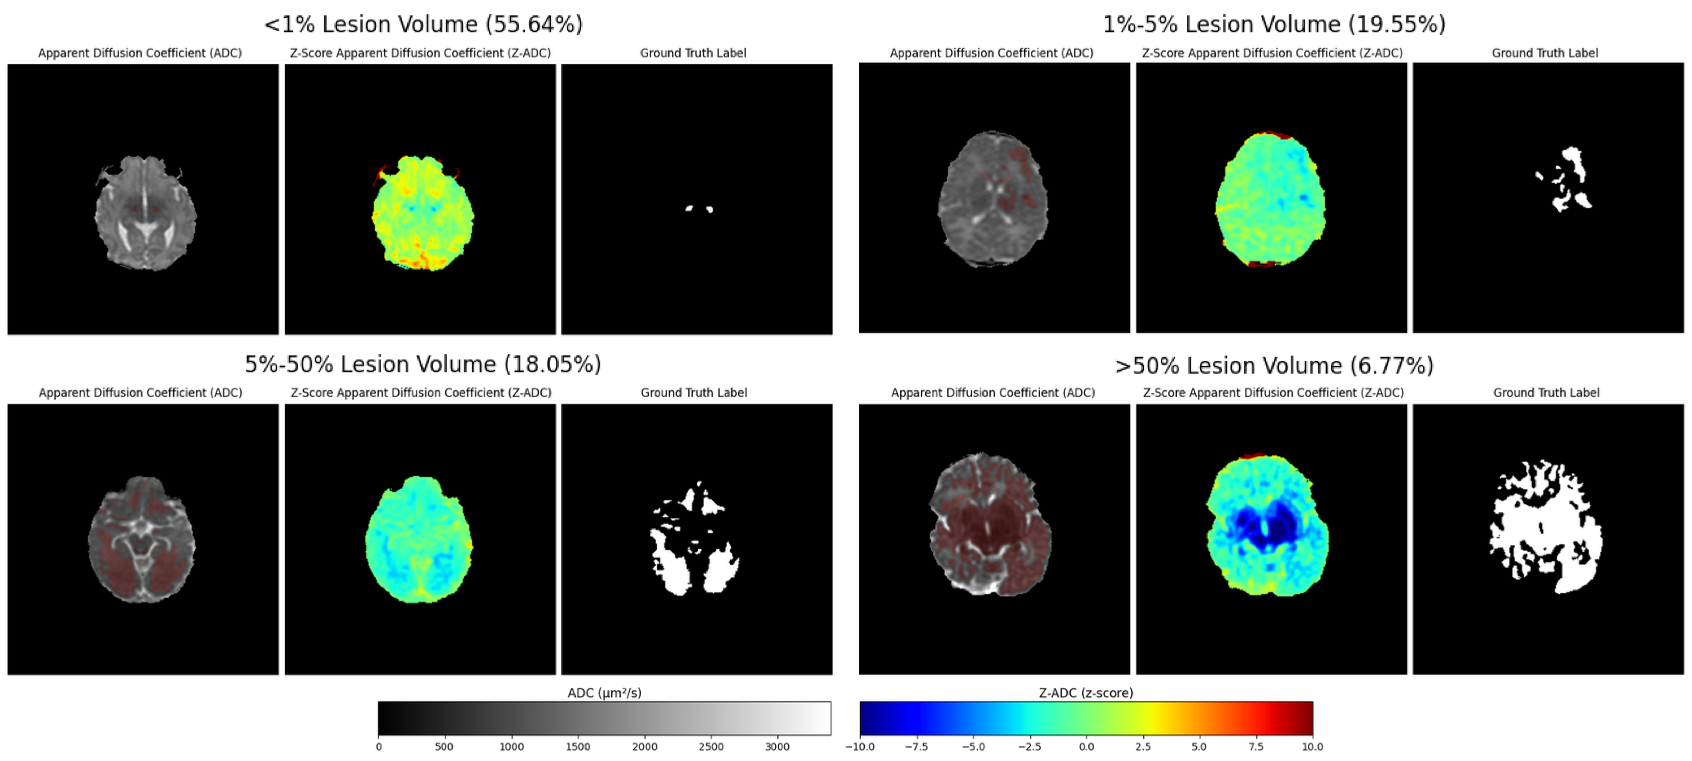

Methodology

The Spatial Prior: Input Representation and Pre-Processing

We propose a **three-channel input representation** by stacking ADC, ZADC, and a binary mask from ZADC thresholded at -2, which acts as a **non-exclusive spatial prior**, highlighting high-probability lesion regions without obscuring borderline voxels that are visible in the other channels.

Efficient & Attentive Architecture

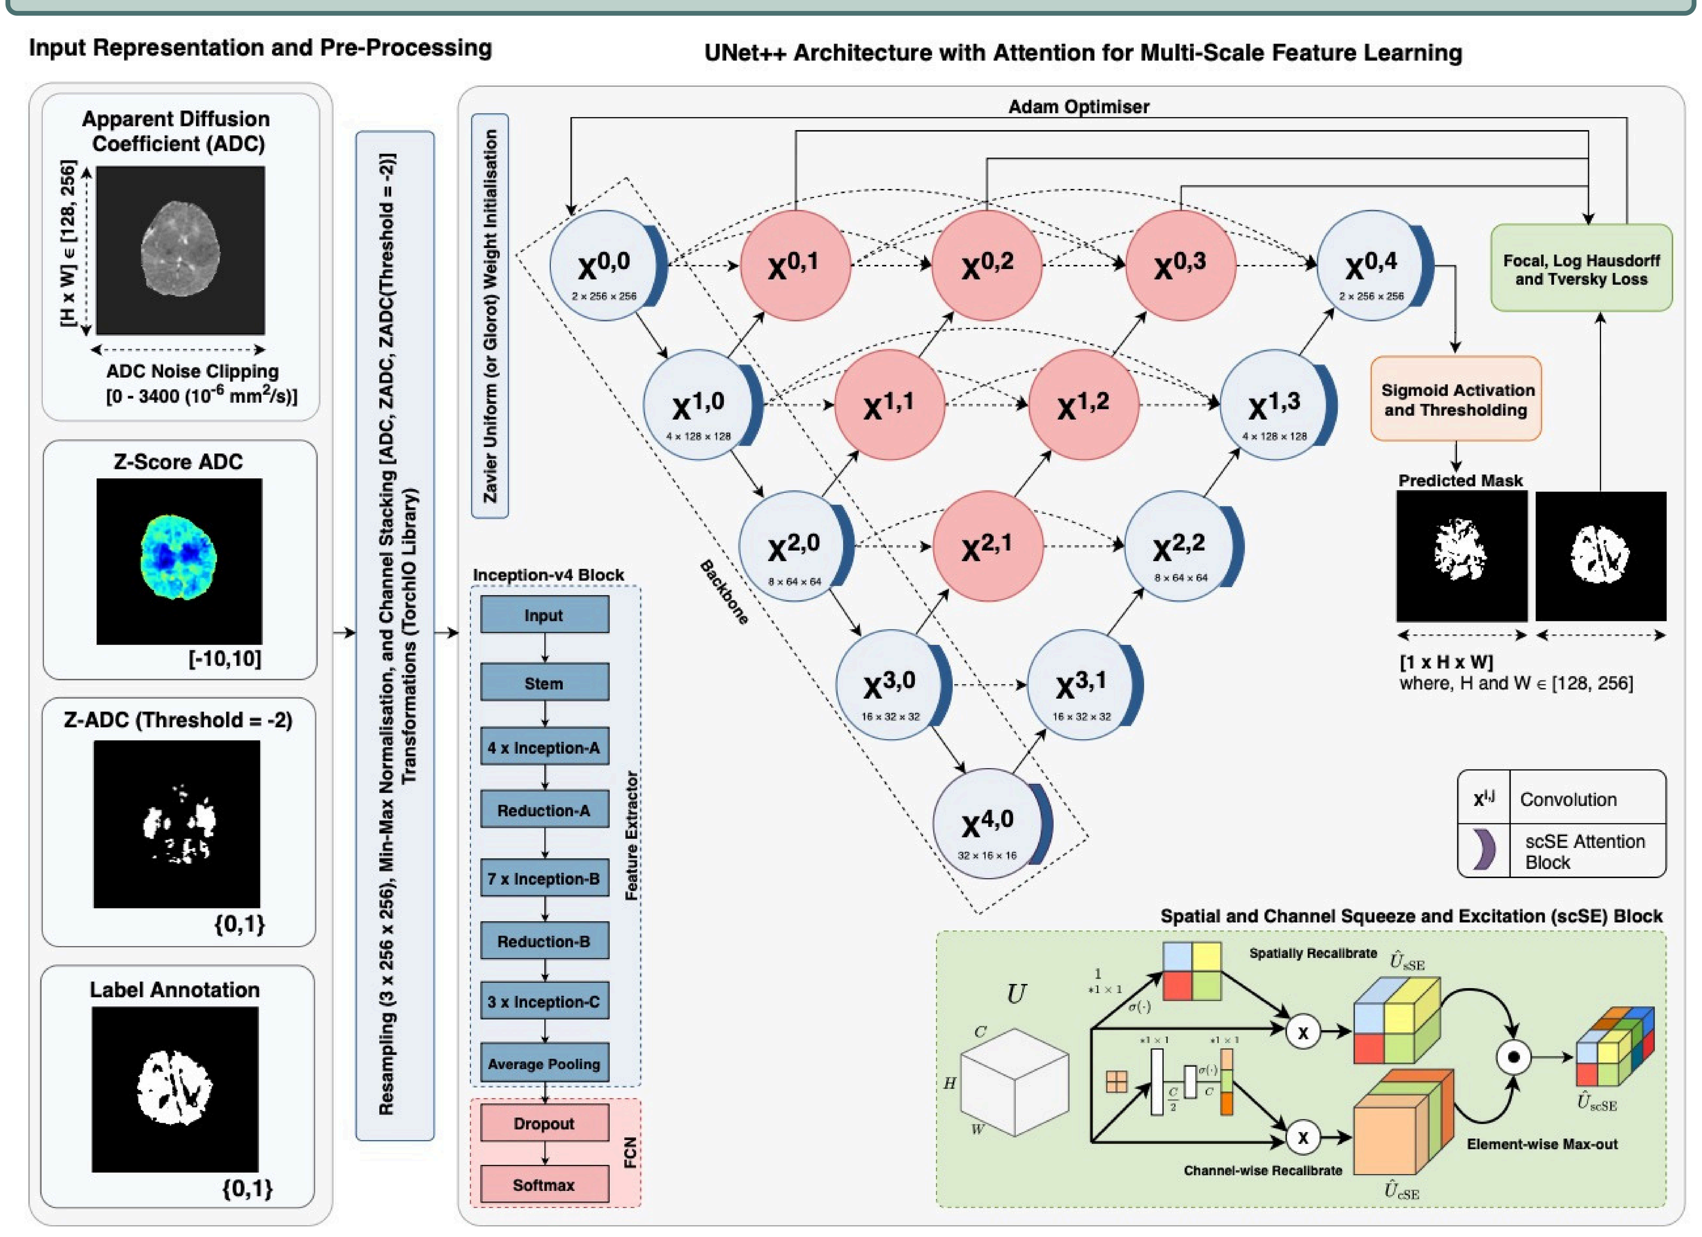

- Our framework uses the **UNet++ architecture** with an **Inception-v4 encoder** with dense nested skip connections to enhance multi-scale feature aggregation.
- **Spatial and Channel Squeeze-and-Excitation (scSE)** blocks are integrated to perform dynamic feature recalibration across spatial and channel dimensions $\hat{U}_{scSE} = \hat{U}_{cSE} + \hat{U}_{sSE}$, improving the delineation of lesion boundaries.

[1] The BONBID-HIE dataset is publicly available through the 2024 MICCAI Challenge: https://bonbid-hie2024.grand-challenge.org/

Loss Collapse

Dice-Focal Loss: Training with Dice-Focal loss leads to a **loss collapse** where the model predicted near-zero masks, causing gradients to vanish and trapping it in a local minimum. Training recovered through stochastic perturbations, but supervision for small lesions remains ineffective.

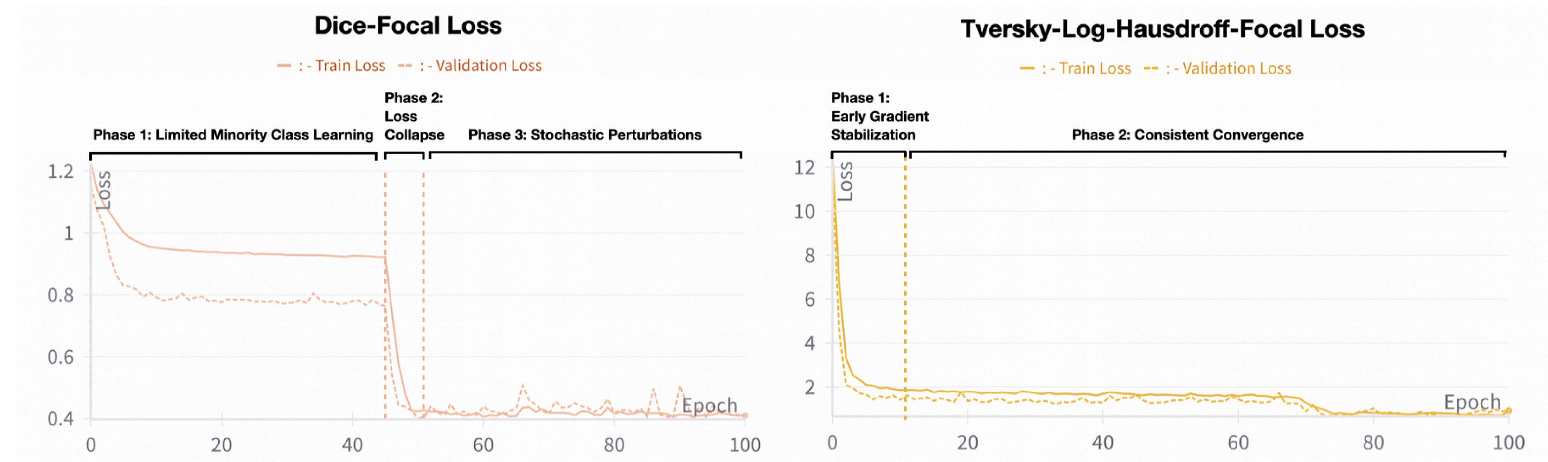

Boundary and Region Aware Loss: Tversky Log Hausdorff Focal Loss

To address loss collapse and the lack of volumetric context in 2D methods, we designed a hybrid, weighted loss function: $L_{TLHF}(p, q) = \gamma_T \cdot L_{Tversky} + \gamma_{LH} \cdot L_{LH} + \gamma_L \cdot L_{Focal}$

- **Tversky Loss:** A generalized Dice Loss that balances false positive and false negative penalties, making it robust to class imbalance.
- **Log-Hausdorff Loss:** Captures **boundary discrepancies** using distance transforms, with a logarithmic term enhancing sensitivity to small boundary errors.
- **Focal Loss:** Addresses class imbalance by focusing the model on hard, misclassified pixels.

Results and Conclusion

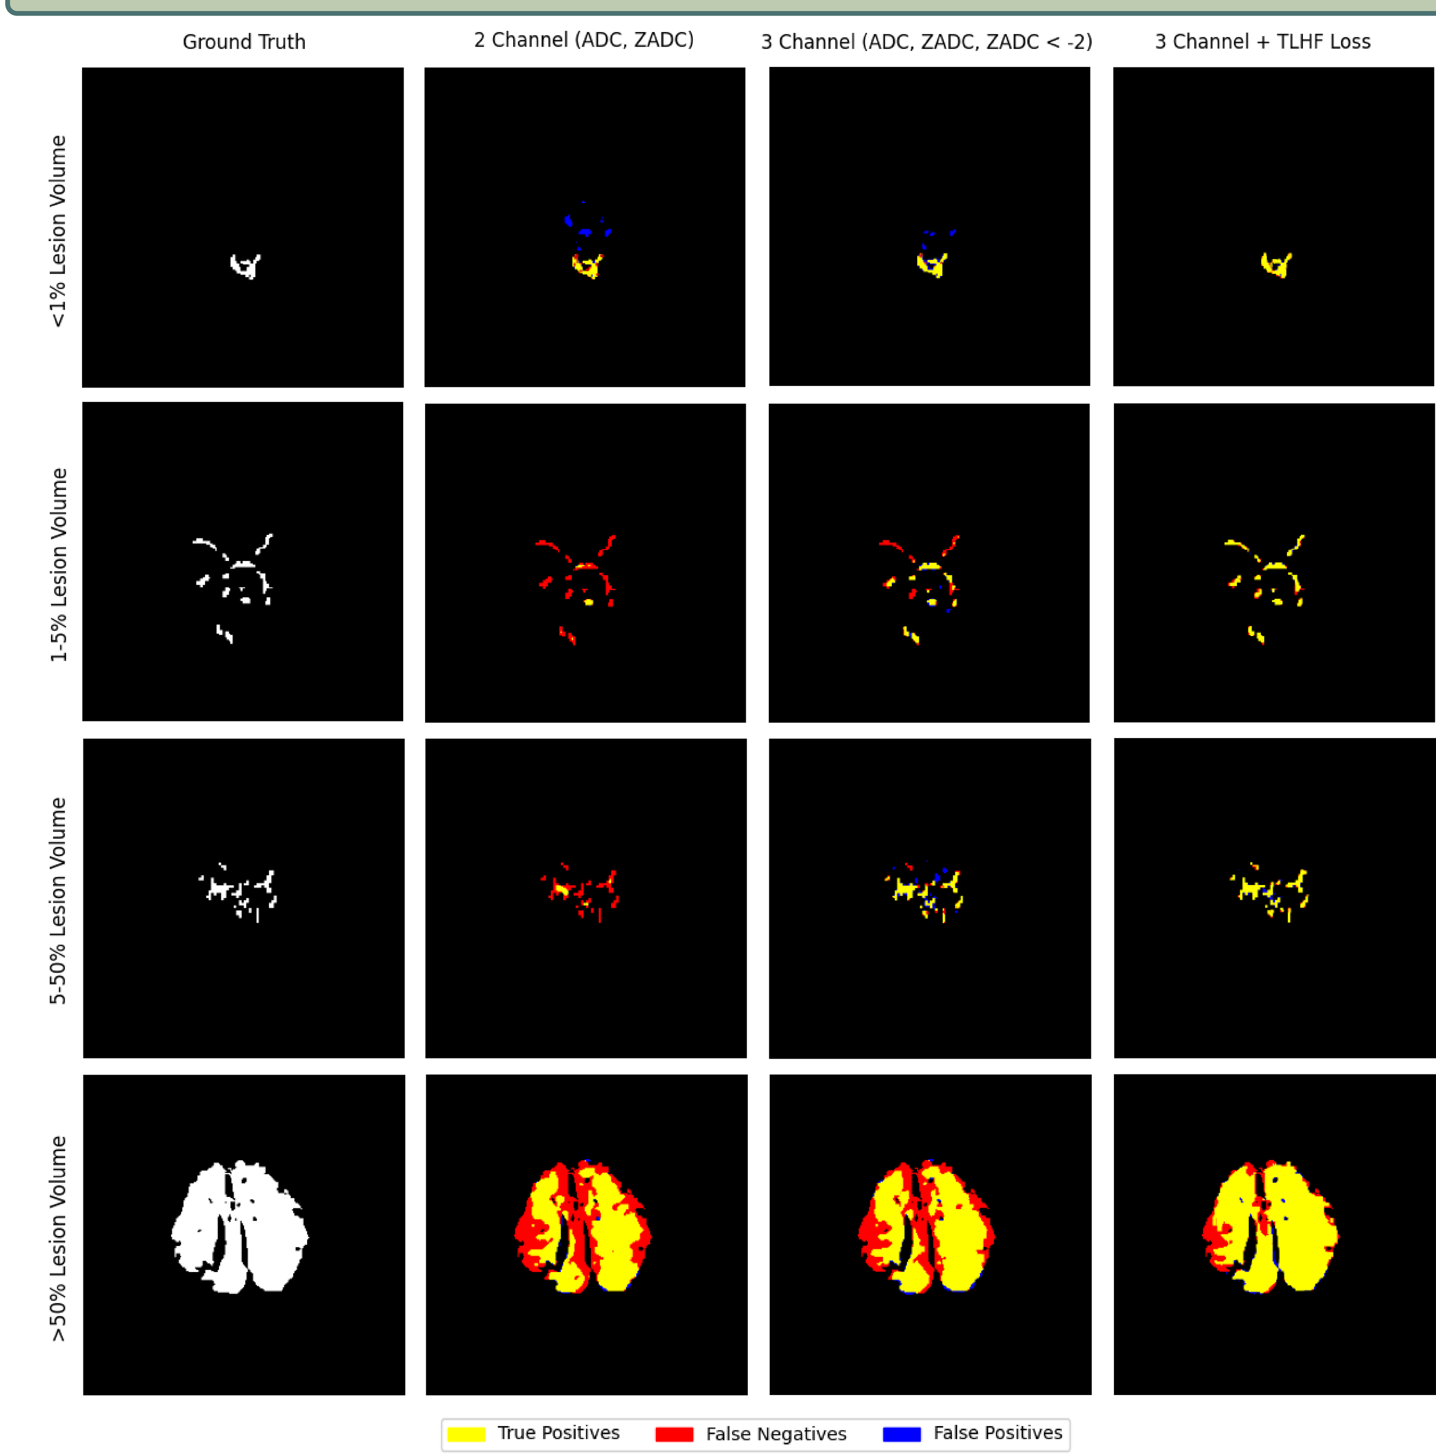

Channel Stacking

The 3-Channel input outperforms all other configurations, achieving a **5.83% DSC** improvement over the 2-channel input. This confirms the effectiveness of the thresholded ZADC map as a spatial prior.

Input	DenseNet-161			Inception-v4		
	DSC (↑)	MASD (↓)	NSD (↑)	DSC (↑)	MASD (↓)	NSD (↑)
ADC	0.2977	5.984	0.3863	0.3239	5.1795	0.4361
Z_{ADC}	0.5128	3.5919	0.6689	0.5221	4.1288	0.6641
[ADC, Z_{ADC}]	0.5438	3.2168	0.6655	0.554	3.1597	0.68
[ADC, Z_{ADC}, Z_{ADC} < -2]	0.5763	3.1023	0.6751	**0.5863**	**3.0711**	**0.6941**

Loss and Attention

The proposed TLHF loss substantially outperforms Dice-Focal loss, improving **MASD** from **3.0711 to 2.6484 (a 13.7% reduction)**. Adding scSE attention further improves boundary detection, particularly for smaller lesions, increasing the overall **NSD** from **0.7153 to 0.7477**.

Loss	Attention	DSC (↑)	MASD (↓)	NSD (↑)
Dice-Focal Loss	None	0.5863	3.0711	0.6941
TLHF Loss	None	**0.606**	**2.6484**	0.7153
TLHF Loss	**scSE**	0.5986	2.7986	**0.7477**

Design	< 1%			1-5%			> 5%		
	DSC	MASD	NSD	DSC	MASD	NSD	DSC	MASD	NSD
TLHF	0.474	**4.131**	0.5997	**0.6753**	**1.6719**	**0.8011**	**0.8132**	**0.855**	**0.8665**
TLHF + scSE	**0.4776**	4.2458	**0.6777**	0.6449	1.8769	0.7802	0.809	0.8653	0.8626

Competitive Results with 3D Architectures

Our computationally efficient 2D framework achieves performance that is highly competitive with parameter-heavy 3D architectures, securing a **top-three rank** for both **DSC (0.6060)** and **NSD (0.7477)** among all published methods on the BONBID-HIE dataset.

Participant/Team	Method	Architecture	DSC (↑)	MASD (↓)	NSD (↑)
Xu et al.	3D	UNet	0.6262	2.5017	0.7316
Toubal et al.	3D	Swin-UNETR w/ RF	0.6215	2.2556	0.7678
Ours	**2D**	**UNet++**	**0.606**	**2.6484**	**0.7153**
Ours	**2D**	**UNet++ w/ scSE**	**0.5986**	**2.7986**	**0.7477**
Koirala et al.	3D	UNet Ensemble	0.58	2.5993	0.7379
Wodzinski et al.	3D	ResUNet	0.577	2.9419	0.7379
Tahmasebi et al.	3D	nnUNet	0.4998	5.3327	0.6449
Aydin et al.	3D	Attention UNet	0.4826	3.4626	0.6176

Conclusion. We demonstrate that jointly optimizing a **parametrically efficient 2D architecture** with a **non-exclusive spatial prior** and a **boundary- and region-aware loss** effectively mitigates loss collapse, enabling robust segmentation of hyper-acute neonatal lesions from anisotropic diffusion MRI.