Learning Continuous Spectral Dynamics for Robust and Interpretable Brain Tumor Analysis

Ujjwal Jain[1,2][0009−0005−8268−7532], Santosh Prakash Chouhan[1,2][0009−0002−0817−1798], Roshni Chakraborty[1][0000−0003−4476−403X], and Mahua Bhattacharya[1,2][0000−0002−3996−9818]

[1] ABV-Indian Institute of Information Technology and Management, Gwalior, Madhya Pradesh, India
[2] Centre for Biomedical Research (CBR), ABV-IIITM Gwalior, India
{imt_2021105,santoshc,roshni,mb}@iiitm.ac.in

Abstract. Brain tumor analysis remains a critical challenge due to the heterogeneity of tumor appearance, complex tissue morphology, and treatment-induced variations. Most existing approaches rely heavily on large pretrained vision models, which are inherently susceptible to spectral bias: their dominant low-pass inductive behaviour suppresses informative high-frequency components. As a result, essential spectral-dynamic textural cues and cognitively relevant patterns required for reliable tumor grading are overlooked. To address these limitations, we propose PhysiQ-Net, a physics-informed framework for brain tumor diagnosis operating in the continuous frequency domain. The model is based on Continuous Spectral Dynamics Learning (CSDL), which unifies classical signal-processing principles with modern deep learning to enable joint modeling of the spatial morphology and continuous spectral characteristics of anatomical structures. Our approach integrates spatial and spectral representations to learn discriminative, grade-aware features. A Neural Ordinary Differential Equation (NODE) module further refines these representations by enabling adaptive computational depth aligned with sample complexity. Finally, a complex-valued interference layer serves as the classification head, explicitly quantifying epistemic uncertainty by mapping ambiguous features to destructive phase interference. Experimental results on the BRISC 2025 dataset show that PhysiQ-Net achieves a mean accuracy of 98.50% with a macro F1-score of 0.986. By balancing diagnostic precision with a lightweight parameter footprint (16.99 M) and a latency score of 3.63 milliseconds, this framework establishes a mathematically grounded proof-of-concept for robust, continuous-domain diagnostic modeling. .

Keywords: Spectral bias, Continuous Spectral Dynamics Learning (CSDL), Neural Ordinary Differential Equation (NODE), Epistemic uncertainty

1 Introduction

Continuous Spectral Dynamics Learning (CSDL) paradigm bridges classical signal-processing theory and modern deep learning, enabling classifiers to jointly model

the spatial morphology and spectral characteristics of anatomical structures within a unified, mathematically grounded framework [1]. Deep learning paradigms for medical imaging, particularly Convolutional Neural Networks (CNNs) and Vision Foundation Models (VFMs), have traditionally operated in the spatial domain [20, 22]. While CNNs encode local morphology and VFMs capture global context, both suffer from spectral bias, a tendency to prioritize low-frequency structural features while attenuating high-frequency textural gradients crucial for identifying diffuse tumor margins [2, 18].

Brain tumors, particularly gliomas, exhibit high intra- and inter-class heterogeneity. Because MRI data are complex biomedical signals governed by frequency responses rather than mere pixel grids, tumor sub-regions often exhibit overlapping intensities yet distinct spectral signatures [5, 16]. Spatial appearance alone is frequently insufficient for reliable discrimination, and models relying on discrete spatial sampling struggle to model continuous frequency interactions essential for identifying subtle pathological variations [19, 21]. This limitation motivates a shift toward frequency-domain formulations that natively capture multi-scale anatomical patterns and enhance robustness to noise, domain shifts, and acquisition artifacts [17, 23].

To explicitly mitigate this spectral bias, we introduce PhysiQ-Net, a physics-informed brain tumor grading framework. Our central contribution is the Continuous Spectral Dynamics Learning (CSDL) paradigm. To correctly formulate medical image diagnosis in the frequency domain, CSDL operates in two synergistic phases: first, it extracts multi-scale spectral features using discrete, orthogonal wavelet bases to perfectly preserve high-frequency textural cues; subsequently, it models the evolution of these frequency representations as a continuous dynamical system via a Neural Ordinary Differential Equation (NODE) module.

By treating the feature trajectory as a continuous biological dynamic rather than a fixed sequence of discrete layers, PhysiQ-Net enables adaptive computational reasoning depth that scales with sample complexity. Furthermore, to address the critical need for diagnostic reliability, a phase interference layer is employed as the classification head to explicitly quantify epistemic uncertainty, natively mapping ambiguous features to destructive interference. Consequently, PhysiQ-Net achieves highly robust and interpretable performance on the BRISC 2025 dataset [16]. By balancing diagnostic precision with a lightweight parameter footprint and low inference latency, this framework establishes a mathematically grounded pathway toward trustworthy clinical decision support across heterogeneous clinical environments.

2 Methodology

2.1 Problem Formulation

We define the brain tumor classification task as learning a mapping from an input MRI volume $X \in \mathbb{R}^{H \times W \times C}$ to a categorical diagnosis $y \in \{1, \ldots, K\}$.

Unlike standard deep learning models that approximate this via a static function $y = f(X)$, we aim to learn a dynamic mapping coupled with explicitly quantified uncertainty. Our goal is to capture the variable complexity of tumor heterogeneity by treating the inference process as a continuous evolution of spectral features [1], rather than a fixed sequence of discrete spatial operations.

2.2 Architectural Overview

To realize this continuous dynamic formulation, PhysiQ-Net is a hybrid neurophysical architecture designed to mimic diagnostic reasoning through six tightly coupled, interdependent stages: (I) Physics-Informed Perception, (II) Continuous Trajectory Modeling, (III) Prototype-Based Semantic Grounding, (IV) Phase-Interference Uncertainty Estimation, (V) Topology-Preserving Segmentation, and (VI) Optimization, as shown in Figure 1.

Formally, PhysiQ-Net acts as a composite dynamic operator $\mathcal{F} : \mathcal{X} \to \mathcal{Y} \times [0, 1]$, mapping the input tensor X to a diagnostic probability distribution and an associated epistemic uncertainty score. Rather than operating as isolated modules, each stage explicitly prepares the representational manifold for the next, ensuring logical consistency from signal extraction to final classification.

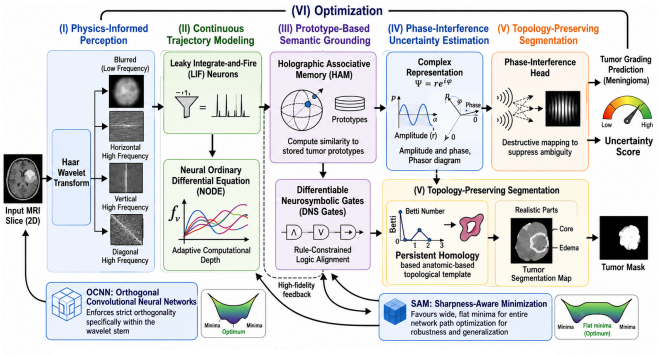

Fig. 1: Comprehensive architectural workflow of the proposed PhysiQ-Net.

2.3 Stage I: Physics-Informed Perception

The foundation of our architecture relies on extracting highly discriminative signals from the raw MRI. Standard CNNs suffer from spectral bias, effectively filtering out the high-frequency textural details (e.g., infiltrative margins) crucial for tumor grading [2]. To prevent this initial diagnostic data loss, we must first decompose the image into a multi-scale spectrum of frequencies.

Haar Wavelet Frequency Stem: We introduce a parallel frequency branch utilizing the 2D Discrete Wavelet Transform (DWT). Unlike learnable convolutions that are biased toward low frequencies, we fix the kernels to orthogonal

Haar bases to guarantee lossless spectral decomposition. For an input x, the decomposition is performed via separable convolution with low-pass and high-pass filters:

$$L = \frac{1}{\sqrt{2}}\begin{bmatrix}1\\1\end{bmatrix}, \quad H = \frac{1}{\sqrt{2}}\begin{bmatrix}1\\-1\end{bmatrix}$$

This yields four sub-bands at each scale j:

$$\begin{aligned} x^j_{LL} &= (x^{j-1} * L) * L^T \quad \text{(Approx.)} \\ x^j_{LH} &= (x^{j-1} * L) * H^T \quad \text{(Horiz. Detail)} \\ x^j_{HL} &= (x^{j-1} * H) * L^T \quad \text{(Vert. Detail)} \\ x^j_{HH} &= (x^{j-1} * H) * H^T \quad \text{(Diag. Detail)} \end{aligned} \tag{1}$$

The detail coefficients $\{x_{LH}, x_{HL}, x_{HH}\}$ capture directional gradients and microtexture. These are normalized via Instance Normalization and fused with the spatial features $F_{spatial}$ using a channel-wise attention gate $\mathcal{A}$:

$$F_{fused} = F_{spatial} \oplus \mathcal{A}\left(\text{Concat}(x_{LH}, x_{HL}, x_{HH})\right) \tag{2}$$

Fractal Pooling: Once the spatial-spectral features are fused, they must be downsampled for computational efficiency before passing to the dynamic stages. Because tumor boundaries exhibit fractal roughness distinct from healthy tissue, we employ Fractal Pooling [5] to selectively compress the data. We define the pooling operator P over a local window Ω_k as:

$$y_k = \mu_k + \gamma \log\left(1 + \frac{\sigma_k^2}{\mu_k + \epsilon}\right) \cdot (\max(\Omega_k) - \mu_k) \tag{3}$$

This pooling operation reduces the high-dimensional fused tensor F_{fused} to a compact spatial-spectral latent grid $y_k \in \mathbb{R}^{H' \times W' \times d_{model}}$, which serves as the morphology-aware initial state $h(0)$ for the continuous dynamics.

2.4 Stage II: Continuous Trajectory Modeling

With the rich spectral features extracted as $h(0)$, standard architectures would pass these through discrete network layers. Instead, we filter these signals for acquisition noise and model their subsequent transformation as the continuous trajectory of a dynamical system.

Spiking Leaky Integrate-and-Fire (LIF) Neurons: Before tracking the continuous evolution of these features, we must suppress Rician acquisition noise inherent in MRI, which often masquerades as high-frequency detail. We employ a layer of LIF neurons [11]. The membrane potential $u_i[t]$ evolves dynamically, imposing a temporal sparsity constraint:

$$u_i[t] = \underbrace{\beta u_i[t-1]}_{\text{Leak}} + \underbrace{\sum_j w_{ij} x_j[t]}_{\text{Input}} - \underbrace{S_i[t-1]\theta_{th}}_{\text{Reset}} \tag{4}$$

This acts as a non-linear low-pass filter, passing only the most salient, noise-filtered features forward to the ODE solver.

Neural Ordinary Differential Equations (ODEs) & LTC Gates: The noise-filtered signals now undergo continuous refinement via a Neural ODE [1]. We define the hidden state $h(t)$ as a continuous function of network depth t, governed by a neural vector field f_θ:

$$\frac{dh(t)}{dt} = f_\theta(h(t), t) = \sigma(\text{GN}(\mathbf{W}_2\phi(\mathbf{W}_1 h(t) + b_1))) \tag{5}$$

where GN is Group Normalization and ϕ is the Tanh activation. The output state is the solution to the IVP: $h(T) = h(0) + \int_0^T f_\theta(h(t), t)dt$. We utilize the Dormand-Prince (Dopri5) solver, an adaptive Runge-Kutta method [6]. The solver's error tolerance ϵ implicitly controls the computational depth, i.e., complex inputs induce stiff dynamics requiring more integration steps (N_{steps}), allowing the network to adapt its reasoning depth to sample difficulty ($N_{steps} \propto$ Complexity).

To account for varying tumor morphologies, we integrate Liquid Time-Constant (LTC) Gates [10]. The integration time constant τ is modulated by a learnable gating network $\tau_\psi(x)$, allowing the solver to dynamically adjust its integration steps (N_{steps}). This forces the model to accelerate over healthy tissue and decelerate to finely integrate complex pathological regions. Following global average pooling, the ODE solver outputs the final integrated latent state $h(T) \in \mathbb{R}^{d_{model}}$, representing the dynamically scaled feature vector for the entire MRI slice.

2.5 Stage III: Prototype-Based Semantic Grounding

While Stage II continuously refines the latent trajectory, we explicitly ground the refined feature vector in a learned semantic space before classification. We map the evolved features to a complex hypersphere using Holographic Reduced Representations [14] to compute similarities against stored key-value prototype pairs. This mechanism explicitly aligns the continuous features with a learned manifold of historical tumor prototypes. To further constrain these representations, we apply Differentiable Neurosymbolic Gates using real-valued T-norms [3], ensuring the latent vectors conform to predefined anatomical and textural rules prior to classification.

Holographic Associative Memory: We map the evolved features to a complex hypersphere $\mathbb{C}^d$ using Holographic Reduced Representations. By performing circular correlation ($\star$) between a query vector q and stored key-value prototype pairs $\{k_i, v_i\}_{i=1}^N$:

$$\text{Sim}(q, k_i) = \text{Re}(q \star k_i) = \text{Re}(\mathcal{F}^{-1}(\mathcal{F}(q) \odot \overline{\mathcal{F}(k_i)})) \tag{6}$$

This mechanism explicitly grounds the continuous features in a learned, interpretable manifold of historical tumor prototypes.

Differentiable Neurosymbolic Gates: To further constrain these grounded prototypes, we enforce logical consistency using real-valued logic. Differentiable

T-norms (e.g., Product for AND, Probabilistic Sum for OR) allow the network to evaluate explicit rules (e.g., $Texture \wedge Location \implies Class$). The neurosymbolic gates output a grounded prototype vector $z_{cog} \in \mathbb{R}^{d_{model}}$, ensuring the subsequent classifier receives features that are statistically likely and logically consistent.

2.6 Stage IV: Phase-Interference Uncertainty Estimation

As Stage III outputs the semantically grounded feature set z_{cog}. However, medical diagnosis carries inherent ambiguity. Standard Softmax layers force probability distributions that frequently result in overconfident predictions on ambiguous data.

Interference Layer: To decouple semantic ambiguity from feature confidence, we model the latent state as a complex wave function $\Psi = re^{i\varphi}$ [4]. The classification logit y_c is derived from the interference intensity:

$$y_c = \langle \Psi | M_c | \Psi \rangle = r(x)^T M_c r(x) \cdot \cos(\varphi(x)) \tag{7}$$

Divergent or conflicting evidence from the logic gates directly induces destructive interference ($\varphi \to \pi/2$, $\cos\varphi \to 0$), natively suppressing confidence for ambiguous inputs without requiring Monte Carlo sampling.

Kolmogorov-Arnold Networks (KAN): Once the interference layer suppresses ambiguous signals, the final deterministic classification is performed by a KAN [7], which represents multivariate functions as sums of univariate B-splines. This yields a highly expressive and parameter-efficient decision boundary compared to standard MLPs.

2.7 Stage V: Topology-Preserving (Biological) Segmentation

While the classification head provides the diagnosis, bounding the network's spatial attention ensures clinical interpretability. We introduce an auxiliary topological segmentation branch utilizing the unpooled spatial grid y_k from Stage I.

Neural Cellular Automata & Topological Persistence: We formulate segmentation as a local morphogenetic process [8], where the mask S_t evolves via local gradient sensing. To strictly enforce that the tumor mask grows as a single connected biological component, we apply a persistent homology loss [9], minimizing the Wasserstein distance between the persistence diagrams $\mathcal{D}$ of the prediction and ground truth:

$$\mathcal{L}_{topo} = \inf_{\gamma \in \Gamma(\mathcal{D}_{pred}, \mathcal{D}_{gt})} \sum_{(p,q)\in\gamma} \|p - q\|_2^2 \tag{8}$$

2.8 Stage VI: Optimization & Robustness

To stabilize training of this complex hybrid system, we utilize: **FOTON (Orthogonal Constraints)** which enforce soft orthogonality ($W^T W \approx I$) via the

Bjorck iteration [13]: $W_{k+1} = W_k(1.5I - 0.5W_k^T W_k)$, preserving gradient norms. **Sharpness-Aware Minimization (SAM)** minimizes the worst-case loss in a neighborhood ρ: $\min_w \max_{\|\epsilon\| \leq \rho} L(w+\epsilon)$ [12] and **Causal Invariance** penalizes the Hilbert-Schmidt Independence Criterion (HSIC) [15] between features Z and scanner style metadata S: $\mathcal{L}_{indep} = \text{HSIC}(Z, S)$.

3 Experiments & Results

In this Section, we briefly discuss our experimental setup followed by detailed results and discussions.

3.1 Experimental Setup

We evaluate PhysiQ-Net on the BRISC 2025 dataset [16], comprising 6,000 T1-weighted MRI slices balanced across four classes (Glioma, Meningioma, Pituitary, No Tumor) and three anatomical planes, with expert-verified segmentation masks.

To rigorously prevent data leakage and artificially inflated performance metrics commonly associated with slice-level splitting, we implement strict **patient-level dataset partitioning**. All slices extracted from a single patient's MRI volume are explicitly isolated to a single fold. We conduct a **5-fold patient-level cross-validation**, dynamically splitting the patient cohorts into 80% training and 20% testing per fold.

Input images are resized to 224×224 and normalized using standard ImageNet statistics. The model is optimized for 50 epochs per fold (batch size 32) using Sharpness-Aware Minimization (SAM) with an AdamW base (learning rate $= 10^{-4}, \rho = 0.05$) to promote generalization. We employ a joint loss function combining Categorical Cross-Entropy for classification, MSE for segmentation, and our topological persistence loss, utilizing random horizontal flipping for augmentation.

3.2 Performance Comparison

Experimental results consistently surpasses a broad range of CNN and transformer benchmarks on the BRISC 2025 dataset. Our model achieves the highest classification accuracy(98.50%) and F1-score(0.986) exceeding that of large capacity models like EfficientNet-V2, DenseNet-161 and Vision Transformers while using significantly less number of parameters as shown in Table 1. The results clearly demonstrate that our model achieves strong class-balanced performance, shown by a Matthews correlation coefficient (MCC) of 0.9796 and Cohen's Kappa of 0.9795. The model yields well-calibrated predictions, demonstrated by a small log loss (0.0629), while it attains a high level of spatial agreement (IoU =0.9733) as shown in Table 2. Lightweight architectures such as MobileNetV2 and ShuffleNetV2 reduce computational overhead to under 5M parameters utilizing mechanisms like depth-wise separable convolutions and channel shuffling. However,

this efficiency comes at the cost of representational capacity, these models typically plateau at approximately 78–87% accuracy on granular tumor classification tasks, failing to capture the subtle textural heterogeneity required for precision diagnosis. Conversely, while heavy architectures like Vision Transformers (ViT-B) achieve state-of-the-art performance (93%), they do so with a prohibitive computational footprint (>86M parameters), creating a critical need for methods that balance efficiency with deep feature extraction. **PhysiQ-Net** needs 16.99 M parameters and roughly 3.5 GFLOPs. The low inference latency (3.63 ms/image) and high throughput (275.4 images/s) further demonstrate its feasibility for real-time and resource-constrained clinical deployment.

Table 1: Experimental results comparison with large pretrained models in terms of accuracy, F1-score, parameters and latency on the BRISC 2025 dataset.

Category	Model (Variant)	Acc. (%)	F1	Params (M)	Latency (ms)
CNN-Based Models					
EfficientNet	V2-M	95.0	0.9497	52.86	11.18
	V2-S	84.0	0.8399	20.18	9.13
	B4	66.0	0.5934	17.56	9.25
DenseNet	161	95.0	0.9496	26.48	15.72
	201	93.0	0.9285	18.10	11.68
ResNet	34	86.0	0.8572	21.29	7.86
	50	79.0	0.7779	23.52	9.20
	152	90.0	0.9011	58.15	13.14
Wide ResNet	50	90.0	0.9001	66.84	18.47
	101-2	95.0	0.9495	124.85	16.49
MobileNet	V2	87.0	0.8637	2.23	7.32
	V3-Large	73.0	0.6780	4.21	6.97
RegNet	X-400MF	90.0	0.9002	5.10	7.29
	Y-400MF	84.0	0.8330	3.90	9.45
ConvNeXt	Tiny	87.0	0.8649	27.82	9.03
	Small	94.0	0.9395	49.46	14.37
ShuffleNet	V2	78.0	0.7637	1.26	9.36
Others	AlexNet	92.0	0.9185	57.02	6.92
	GoogLeNet	87.0	0.8688	5.60	8.13
	Inception-V3	89.0	0.8878	25.12	15.86
	VGG19-BN	89.0	0.8947	139.60	12.76
	SqueezeNet-1.1	85.0	0.8531	0.72	7.37
	MNASNet-1.0	30.0	0.1832	3.11	9.87
Transformer-Based Models					
ViT	B/16	93.0	0.9300	85.80	16.81
Swin	Tiny	81.0	0.8064	27.52	11.02
	V2-Tiny	91.0	0.9104	27.59	12.20
MaxViT	Tiny	85.0	0.8526	30.41	12.85
Proposed Model					
PhysiQ-Net	–	**98.50**	**0.9864**	**16.99**	**3.63**

Table 2: Performance and efficiency scores of **PhysiQ-Net** on BRISC 2025.

Metric	**Value**
Accuracy	0.985
Top-2 Accuracy	0.997
F1-Score (Macro)	0.9864
F1-Score (Weighted)	0.9850
Matthews Correlation Coefficient (MCC)	0.9796
Cohen's Kappa	0.9795
Log Loss	0.0629
Hamming Loss	0.0150
Jaccard Index (IoU)	0.9733
Total Parameters	16.99 M
Estimated FLOPs	∼3.5 GFLOPs
Model Size	64.9 MB
Inference Latency	3.63 ms / image
Throughput	275.4 images/s

3.3 Class Performance Analysis

As described in Table 3, **PhysiQ-Net** achieves high sensitivity and specificity for all the classes, which demonstrates that the proposed algorithm is a good detector with too much false alarm. Meningioma achieves perfect sensitivity (1.000) and near-perfect specificity (0.9988), while Glioma and Pituitary also achieve high recall (0.9869 and 0.9646, respectively) despite considerable morphological heterogeneity. The No Tumor class also retains both high sensitivity (0.9933) and specificity (0.9971), which means a reduced rate of false positive that are critical in clinical applications. Error-related quantities provide similar reassurance, with uniformly low FDR ($\leq$0.0351) and uniformly high NPV ($\geq$0.9881) for all classes. Furthermore, low Brier scores, particularly for Meningioma (0.0011) and No Tumor (0.0038), indicate well-calibrated probability predictions.

Table 3: Experimental results on per-class clinical performance on the BRISC 2025 dataset.

Class	**Sensitivity**	**Specificity**	**NPV**	**FDR**	**Brier Score**
Pituitary	0.9646	0.9987	0.9881	0.0041	0.0089
Glioma	0.9869	0.9841	0.9942	0.0351	0.0129
Meningioma	1.0000	0.9988	1.0000	0.0071	0.0011
No Tumor	0.9933	0.9971	0.9971	0.0067	0.0038

Finally, t-SNE visualization of the latent space (Fig. 2) reveals well-separated and compact class clusters with minimal overlap, validating the effectiveness of the proposed representation in capturing class-specific structure.

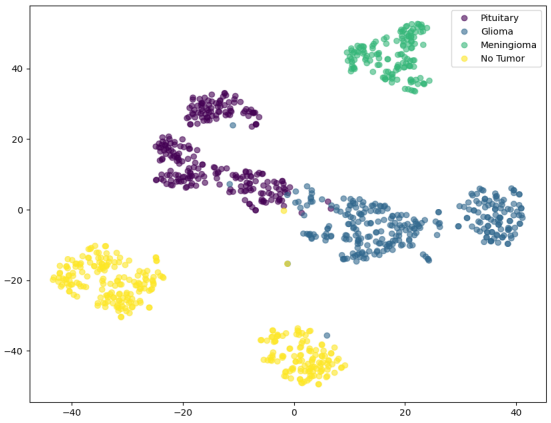

Fig. 2: t-SNE visualization of the learned latent space showing clear class separation among tumor categories.

3.4 Quantitative Classification Performance

The proposed architecture was evaluated on the BRISC dataset, demonstrating exceptional diagnostic capabilities. As illustrated in the Confusion Matrix (Fig. 3), the model achieves an overall accuracy of **98.5%** (985/1000 correct predictions). Notably, the model demonstrates perfect sensitivity for Meningioma

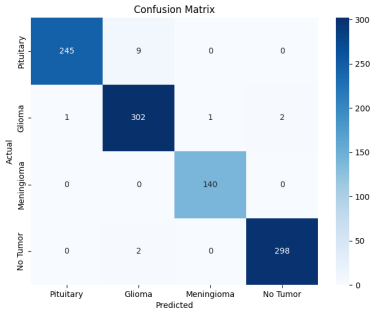

Fig. 3: Confusion matrix of the proposed architecture evaluated on the BRISC dataset.

(140/140) and near-perfect detection for healthy, non-tumor scans (298/300). The most prominent, albeit minor, source of confusion occurs between Pituitary tumors and Gliomas, with 9 actual Pituitary cases misclassified as Gliomas. Clinically, this overlap is biologically plausible due to the close anatomical proximity

of certain suprasellar extensions and the overlapping spectral intensities these soft tissues can exhibit in MRI sequences. However, the model successfully minimizes False Negatives (predicting "No Tumor" when a tumor is present); only 2 Glioma cases were misclassified as healthy tissue, heavily validating the effectiveness of the Haar Wavelet Frequency Stem in preserving subtle infiltrative margins that standard CNNs often miss.

3.5 Discriminative Power: ROC and DET Analysis

To evaluate the model's discriminative thresholding, we analyzed the Receiver Operating Characteristic (ROC) and Detection Error Tradeoff (DET) curves. The One-vs-Rest ROC curves (Fig. 4) reveal a perfect Area Under the Curve (**AUC = 1.00**) across all four classes. The curves uniformly hug the upper-left quadrant, indicating that the model achieves near-perfect True Positive Rates before incurring any False Positives. Because ROC curves can obscure perfor-

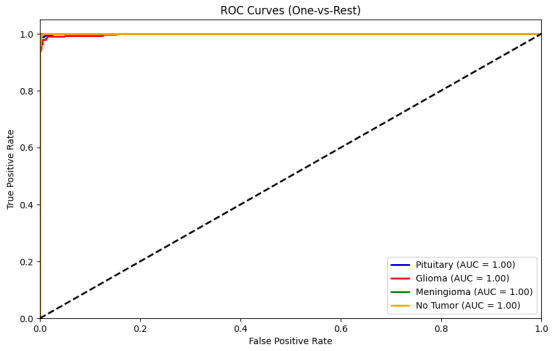

Fig. 4: One-vs-Rest Receiver Operating Characteristic (ROC) curves for the four classification categories.

mance nuances in highly imbalanced or near-perfect models, we further analyzed the DET curve plotted on a logarithmic scale (Fig. 5). The DET curve is highly relevant for medical diagnostics, as it explicitly visualizes the trade-off between the False Negative Rate (FNR, or miss rate) and False Positive Rate (FPR, or false alarm rate). The curves for all classes remain tightly constrained to the bottom-left origin. The Glioma class exhibits a slightly elevated trajectory compared to the others, accurately reflecting its highly heterogeneous morphology and the inherent difficulty in establishing its diffuse boundaries. Nevertheless, the extremely low log-scale error rates confirm that the continuous depth reasoning provided by the Neural ODEs effectively scales its computational complexity to handle these difficult cases.

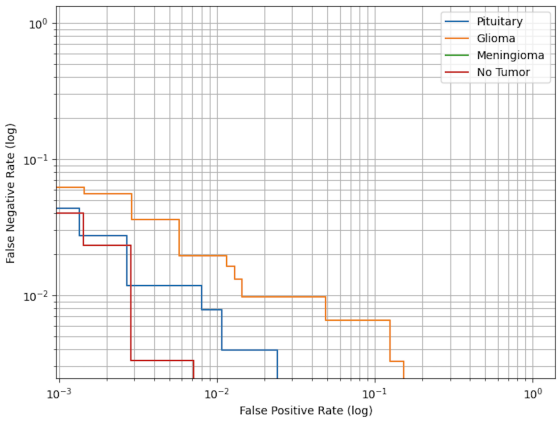

Fig. 5: Detection Error Tradeoff (DET) curve plotted on a logarithmic scale.

3.6 Ablation Study

We study the effect of each component in **PhysiQ-Net** with a fixed training and testing protocol described in Table 4. The accuracy dropped by about 1.7% when the physics-informed wavelet stem (Stage I) was removed, indicating its importance on texture-aware feature learning. Replacing the continuous-depth Neural ODE model (Stage II) with a finite-depth backbone resulted in a 1.4% performance drop, indicating a decrease in robustness. Without the phase interference uncertainty module (Stage IV) the accuracy degraded and the calibration became worse, with a drop in accuracy of 1.1%. Replacing the KAN classifier (Stage III) with a conventional MLP decreased accuracy by 0.9%, and removing the NCA (Stage V) segmentation head introduced structural inconsistency, decreasing performance by 1.3%. Lastly, elimination of SAM optimization and/or orthogonality constraints (Stage VI) had negative impact on convergence and generalization, causing the performance loss of 1.6% in terms of accuracy. In summary, the full configuration as shown in Fig. 1 yields the best score, which is an overall gain in about 2–3% over the partly ablated versions.

Table 4: Ablation Study Showing the Contribution of Individual Components in **PhysiQ-Net**

Model Variant	**Accuracy (%)**	**MCC**	**F1-score**
w/o Wavelet Stem	96.80	0.9580	0.9621
w/o Neural ODE (Fixed Depth)	97.10	0.9624	0.9678
w/o Phase Uncertainty Layer	97.40	0.9681	0.9712
KAN $\rightarrow$ MLP Classifier	97.60	0.9703	0.9735
w/o NCA Segmentation Head	97.20	0.9642	0.9689
w/o SAM + Orthogonality	96.90	0.9601	0.9654
Proposed Model	**98.50**	**0.9796**	**0.9864**

3.7 Qualitative Analysis

The qualitative results (Fig. 6) demonstrate that **PhysiQ-Net** can produce interpretable and decision-transparent results for all tumor types. Grad-CAM uniformly highlights anatomically meaningful regions around tumor sites while physics-aware wavelet energy maps highlight frequency-based structural patterns that are not observed with spatial attention alone. Integrated gradients map also unveils the fine-grained attribution signals, assisting that the predictions are led by clinically relevant features rather than artifacts in the background. In all cases, even for the no-tumor samples, the fused interpretability views provide unified and mutually reinforcing explanations, strengthening the reliability and confidence of the model's prediction.

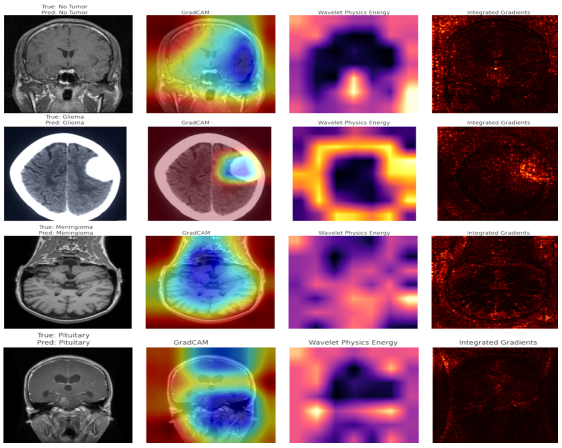

Fig. 6: Qualitative interpretability analysis using Grad-CAM, wavelet energy maps, and integrated gradients across tumor classes.

4 Conclusion

PhysiQ-Net rethinks brain tumor grading by moving beyond pixel-centric learning toward a continuous spectral perspective. By unifying spatial morphology with spectral dynamics through the CSDL paradigm, the model demonstrates that incorporating physically grounded representations leads to highly discriminative and interpretable tumor grading. The integration of Neural ODEs enables adaptive computational depth, while the phase-interference mechanism provides robust uncertainty quantification. While the strong performance on the BRISC 2025 dataset highlights the framework's potential, we acknowledge that evaluations were limited to a single dataset of 2D T1-weighted slices. Extensive external validation, including multi-institutional testing across diverse MRI scanner

protocols and site shifts, is required before clinical translation can be fully justified. Future work will extend this continuous-depth framework to 3D volumetric analysis and evaluate its robustness on heterogeneous, multi-site clinical cohorts.

Acknowledgements The presented work was undertaken with the assistance of resources and services provide by the Center for Biomedical Research which is supported by ABV-IIITM, Gwalior, Madhya Pradesh, India.

References

1. Chen, Ricky T. Q., Yulia Rubanova, Jesse Bettencourt, and David K. Duvenaud. "Neural Ordinary Differential Equations." In *Advances in Neural Information Processing Systems* 31 (2018).
2. Rahaman, Nasim, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. "On the Spectral Bias of Neural Networks." In *International Conference on Machine Learning*, 5301–5310. PMLR, 2019.
3. H. Shindo, V. Pfanschilling, D. S. Dhami, and K. Kersting, "Learning differentiable logic programs for abstract visual reasoning," *Machine Learning*, vol. 113, no. 11, pp. 8533–8584, 2024.
4. Q. Li, D. Gkoumas, A. Sordoni, J.-Y. Nie, and M. Melucci, "Quantum-inspired neural network for conversational emotion recognition," in *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI-21)*, vol. 35, no. 15, pp. 13270–13278, 2021.
5. Florindo, Joao B. "Fractal Pooling: A New Strategy for Texture Recognition Using Convolutional Neural Networks." *Expert Systems with Applications* 243 (2024): 122978.
6. Naranjo-Noda, Frank Sadan, and Juan Carlos Jimenez. "Locally Linearized Runge-Kutta Method of Dormand and Prince for Large Systems of Initial Value Problems." *Journal of Computational Physics* 426 (2021): 109946.
7. Z. Liu, Y. Wang, S. Vaidya, F. Ruehle, J. Halverson, M. Soljačić, T. Y. Hou, and M. Tegmark, "KAN: Kolmogorov–Arnold Networks," in *Proceedings of the International Conference on Learning Representations (ICLR)*, 2025.
8. Mordvintsev, Alexander, Ettore Randazzo, Eyvind Niklasson, and Michael Levin. "Growing Neural Cellular Automata." *Distill* 5, no. 2 (2020): e23.
9. Hu, Xiaoling, Fuxin Li, Dimitris Samaras, and Chao Chen. "Topology-Preserving Deep Image Segmentation." *Advances in Neural Information Processing Systems* 32 (2019).
10. Hasani, Ramin, Mathias Lechner, Alexander Amini, Daniela Rus, and Radu Grosu. "Liquid Time-Constant Networks." In *Proceedings of the AAAI Conference on Artificial Intelligence* 35, no. 9, pp. 7657–7666. 2021.
11. Roy, Kaushik, Akhilesh Jaiswal, and Priyadarshini Panda. "Towards Spike-Based Machine Intelligence with Neuromorphic Computing." *Nature* 575, no. 7784 (2019): 607–617.
12. Foret, Pierre, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. "Sharpness-Aware Minimization for Efficiently Improving Generalization." *arXiv preprint arXiv:2010.01412* (2020).

13. Wang, Jiayun, Yubei Chen, Rudrasis Chakraborty, and Stella X. Yu. "Orthogonal Convolutional Neural Networks." In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11505–11515. 2020.
14. Alam, Mohammad Mahmudul, Edward Raff, Stella Biderman, Tim Oates, and James Holt. "Recasting Self-Attention with Holographic Reduced Representations." In *International Conference on Machine Learning*, 490–507. PMLR, 2023.
15. Wang, Tinghua, Xiaolu Dai, and Yuze Liu. "Learning with Hilbert–Schmidt Independence Criterion: A Review and New Perspectives." *Knowledge-Based Systems* 234 (2021): 107567.
16. Fateh, Amirreza, Yasin Rezvani, Sara Moayedi, Sadjad Rezvani, Fatemeh Fateh, and Mansoor Fateh. "BRISC: Annotated Dataset for Brain Tumor Segmentation and Classification with Swin-HAFNet." *arXiv preprint arXiv:2506.14318* (2025).
17. Beneke, Lisa-Maria, Michell Boerger, Philipp Lämmel, Helene Knof, Andrei Aleksandrov, and Nikolay Tcholtchev. "Leveraging Liquid Time-Constant Neural Networks for ECG Classification: A Focus on Pre-Processing Techniques." In *International Conference on Data Science, Technology and Applications*, 2025.
18. Zhou, Lifang, Yu Jiang, Weisheng Li, Jun Hu, and Shenhai Zheng. "Shape-Scale Co-Awareness Network for 3D Brain Tumor Segmentation." *IEEE Transactions on Medical Imaging* 43, no. 7 (2024): 2495–2508.
19. Karim, Shahid, Geng Tong, Yiting Yu, Asif Ali Laghari, Abdullah Ayub Khan, Muhammad Ibrar, and Faisal Mehmood. "Developments in Brain Tumor Segmentation Using MRI: Deep Learning Insights and Future Perspectives." *IEEE Access* 12 (2024): 26876–26895.
20. Badža, Milica M., and Marko Č. Barjaktarović. "Classification of Brain Tumors from MRI Images Using a Convolutional Neural Network." *Applied Sciences* 10, no. 6 (2020): 1999.
21. Asiri, Abdullah A., Toufique Ahmed Soomro, Ahmed Ali Shah, Ganna Pogrebna, Muhammad Irfan, and Saeed Alqahtani. "Optimized Brain Tumor Detection: A Dual-Module Approach for MRI Image Enhancement and Tumor Classification." *IEEE Access* 12 (2024): 42868–42880.
22. Shamshad, Nadia, Danish Sarwr, Ahmad Almogren, Kiran Saleem, Alia Munawar, Ateeq Ur Rehman, and Salil Bharany. "Enhancing Brain Tumor Classification by a Comprehensive Study on Transfer Learning Techniques and Model Efficiency Using MRI Datasets." *IEEE Access* 12 (2024): 106007–106019.
23. Irfan, Muhammad, Anum Nawaz, Riku Klen, Abdulhamit Subasi, Tomi Westerlund, and Wei Chen. "Improved Brain Tumor Detection in MRI: Fuzzy Sigmoid Convolution in Deep Learning." in *2025 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8, IEEE, 2025.