Deep Learning-Enabled Edge-Aware Multi-Phase Image Enhancement for Precision Liver Lesion Classification

D. Ramakanth[1] B. Bhavya Sree[2], Shaik Thasmiya[3], S. Dilli Prasanna[4], G. K. Rakshitha[5]

[1]Associate Professor, [2,3,4,5]UG Students, Aditya College of Engineering, Madanapalle – 517325, Andhra Pradesh, India
Corresponding author: ramakanth1229@gmail.com

Abstract: Accurate liver lesion classification is critical for early diagnosis and treatment planning. Medical images such as Computed Tomography (CT) and Magnetic Resonance Imaging (MRI) scans frequently suffer from low contrast, noise, and unclear lesion boundaries, making automated analysis a difficult task. Existing systems rely on basic preprocessing and simple binary classification that fail to preserve edges and handle subtle texture variations. This paper proposes a Deep Learning-Enabled EdgeAware Multi-Phase Image Enhancement framework for precision liver lesion classification. A five-stage preprocessing pipeline — comprising Hounsfield Unit (HU) windowing, Contrast Limited Adaptive Histogram Equalization (CLAHE), Sobel-based edge extraction, multi-scale feature fusion, and guided bilateral filtering — is introduced to improve lesion visibility while preserving anatomical boundaries. Enhanced images are classified using EfficientNet-B3, extended for three-class prediction: Benign, Malignant (HCC), and Metastatic. Grad-CAM is integrated to provide visual explanations of model predictions. A Streamlit-based interface enables real-time prediction and visualization. The system achieves 96.4% accuracy with a weighted F1-score of 0.96, demonstrating superior performance compared to conventional deep learning baselines.

Keywords: Liver Lesion Classification, EfficientNet-B3, Edge-Aware Preprocessing, CLAHE, HU Windowing, Grad-CAM

I. INTRODUCTION

The liver is one of the most vital organs in the human body, responsible for metabolism, detoxification, and the synthesis of essential biomolecules. Due to its rich blood supply and extensive functionality, the liver is highly susceptible to a variety of diseases. Liver lesions — abnormal regions within liver tissue — range from non-cancerous conditions such as hemangiomas to life-threatening malignancies including hepatocellular carcinoma (HCC) and metastatic tumors. Globally, liver cancer ranks among the leading causes of cancer-related deaths, and early and accurate detection is directly linked to improved survival rates.

Medical imaging modalities, particularly CT and MRI, are the clinical standard for identifying liver abnormalities. CT images, acquired across multiple phases (non-contrast, arterial, portal venous, and delayed), capture the dynamic behavior of liver lesions as contrast material passes through vasculature. Despite advances in imaging technology, manual interpretation remains timeconsuming, subjective, and prone to inter-observer variability. The similarity in tissue intensity between benign and malignant lesions further challenges both radiologists and automated systems.

Deep learning, particularly Convolutional Neural Networks (CNNs), has shown remarkable capability in automating feature extraction and classification of medical images. However, most existing methods apply preprocessing steps that are insufficient for handling the specific challenges of liver CT images: low contrast at lesion boundaries, noise from scanning equipment, and inconsistent intensity across imaging phases. Without targeted image enhancement, classification models receive suboptimal input, limiting their diagnostic accuracy.

This paper addresses these limitations by proposing an edge-aware, multi-phase image enhancement framework that combines a structured five-stage preprocessing pipeline with EfficientNet-B3 for three-class liver lesion classification. The key contributions of this work are: (i) a novel five-stage preprocessing pipeline that integrates HU windowing, CLAHE, Sobel edge extraction, multi-scale fusion, and guided filtering; (ii) multi-phase CT integration to capture temporal contrast dynamics; (iii) Grad-CAMbased visual explanations for clinical interpretability; and (iv) a real-time Streamlit deployment interface.

II. LITERATURE REVIEW

Research in automated liver lesion detection has evolved through three broad phases: traditional image processing, machine learning, and deep learning. Early methods used thresholding, region growing, and hand-crafted texture descriptors. While computationally lightweight, these approaches struggled with low-contrast images and failed to generalize across imaging conditions. Machine learning methods — Support Vector Machines, Random Forests, and Extreme Learning Machines — offered improved accuracy but depended heavily on manually engineered features, limiting adaptability.

Zhang et al. [1] introduced an edge-guided multi-scale adaptive feature fusion network for liver tumor segmentation,

demonstrating that preserving boundary information significantly improves segmentation accuracy. Their work

directly motivates the edge extraction stage in the proposed pipeline. Lu et al. [2] proposed a liver segmentation network with detail enhancement and multi-scale feature fusion that outperformed standard U-Net architectures, particularly for small lesions. However, neither work addressed multi-phase classification with end-to-end deployment.

Sun et al. [3] proposed a cross-phase attention module for multi-phase CT classification that can be plugged into existing architectures. While effective, the method remains dependent on raw image quality at the input stage. Feng et al. [4] developed LiverVoxNet, a large-scale multi-phase CT dataset with annotated lesions, providing a benchmark for multi-phase models. Kumar et al. [5] integrated the LIRADS clinical framework with deep learning predictions to improve clinical trust, showing that explainability is as important as accuracy in practice.

Ying et al. [6] developed a multicenter clinical AI system for focal liver lesion diagnosis using DenseNet-121, achieving 91.20% accuracy with attention-based global feature pooling. Liu et al. [7] presented a sequential multistage deep learning framework that mimics radiologist reading patterns across CT phases. Ma et al. [8] used contrast-enhanced MRI with ResNet-50 to predict postsurgical survival in HCC patients, achieving 89.40% accuracy but limited to binary classification. These methods consistently show that preprocessing quality and multiphase integration are the two dominant factors determining system accuracy.

Ma et al. [9] and Zossou et al. [10] highlighted that radiomics-based and dual-branch feature fusion methods improve recall on minority classes, a key challenge in imbalanced liver datasets. Gowda and Manjunath [11] demonstrated that EfficientNet-based architectures outperform ResNet and DenseNet families on liver lesion classification with fewer parameters, validating the choice of EfficientNet-B3 in the proposed work. Despite these contributions, no existing method integrates all five enhancement stages, multi-phase fusion, Grad-CAM interpretability, and offline Streamlit deployment within a unified framework.

III. PROPOSED METHODOLOGY

A. System Architecture Overview

The proposed system follows a modular five-stage pipeline, as illustrated in Fig. 1. Input CT images are acquired across multiple contrast phases and passed through a structured preprocessing chain. The enhanced output is fed into EfficientNet-B3 for feature extraction and three-class classification. Grad-CAM visualizations are generated at inference time, and results are presented through a Streamlit web interface.

B. Dataset Description

The system is trained and evaluated on liver CT images collected from publicly available medical repositories. Images are organized into three diagnostic classes: Benign (hemangioma, cysts, focal nodular hyperplasia), Malignant — hepatocellular carcinoma (HCC), and Metastatic. The dataset contains CT images captured across multiple contrast phases: non-contrast, arterial, portal venous, and delayed. Images with poor quality or excessive artifacts were excluded during curation. The dataset is split 80:20 into training and test sets using stratified sampling to maintain class proportions. Table I summarizes the dataset properties.

Table 1. Dataset description

Property	Details
Imaging Modality	Multi-Phase CT
Classes	Benign, Malignant (HCC), Metastatic
Phases	-contrast, Arterial, Portal Venous, Delayed
Image Resolution	224 × 224 pixels (resized)
Train / Test Split	80% / 20% (stratified)
Class Imbalance	Addressed via weighted loss and augmentation

C. Five-Stage Preprocessing Pipeline

The preprocessing pipeline is the central contribution of the proposed system. Each stage addresses a specific quality challenge in liver CT images. Table II summarizes all five stages.

Table 2. Five-stage preprocessing pipeline

Stage	Operation	Purpose
1	HU Windowing	Clip Hounsfield Units to liver range [−60, 240 HU]; normalize to [0, 1]
2	CLAHE	Enhance local contrast in low-visibility regions without amplifying noise
3	Sobel Edge Extraction	Highlight lesion boundaries and margin irregularities using gradient magnitude
4	Multi-Scale Fusion	Fuse edge maps with enhanced image at α:β ratio (0.8:0.2) across scales
5	ıided Bilateral Filtering	nooth homogeneous regions while preserving sharp lesion edges

Stage 1 — HU Windowing: CT images encode tissue density in Hounsfield Units. Liver parenchyma and lesions fall within −60 to 240 HU. Clipping to this window removes irrelevant bone and air signals, and pixel values are linearly rescaled to [0, 1]. The mathematical formulation is:

$$I\_norm = clip(I, HU\_min, HU\_max) / (HU\_max - HU\_min) \quad ...(1)$$

Stage 2 — CLAHE: Contrast Limited Adaptive Histogram Equalization enhances local contrast by operating on small image tiles and clipping the histogram at a predefined limit. This improves visibility of low-contrast lesion regions without introducing noise amplification.

Stage 3 — Sobel Edge Extraction: Sobel operators compute horizontal and vertical intensity gradients separately. The gradient magnitude map G is computed as:

$$G_x = X * S_x,\ G_y = X * S_y,\ G = \sqrt{(G_x^2 + G_y^2)} \quad ...(2)$$

Stage 4 — Multi-Scale Fusion: The Sobel edge map is fused with the CLAHE-enhanced image using a weighted sum at multiple spatial scales. Fusion is performed as:

$$I\_fused = \alpha \cdot I\_CLAHE + \beta \cdot G,\ where\ \alpha = 0.8,\ \beta = 0.2 \quad ...(3)$$

Stage 5 — Guided Bilateral Filtering: A guided bilateral filter is applied as a final smoothing step. Unlike standard Gaussian filtering, the bilateral filter considers both spatial proximity and pixel intensity similarity, preserving sharp lesion edges while smoothing noisy regions. The filtered output serves as the final input to the classifier.

D. EfficientNet-B3 Classification

EfficientNet-B3 is selected as the classification backbone due to its compound scaling strategy that simultaneously scales model depth, width, and input resolution, achieving high accuracy with relatively fewer parameters compared to ResNet-50 or DenseNet-121. The pretrained EfficientNet-B3 (trained on ImageNet) is fine-tuned on the liver lesion dataset with the following architecture modifications:

The final classification head of the base model is replaced with: GlobalAveragePooling → BatchNormalization → Dropout (0.3) → Dense (256, ReLU) → Dropout (0.2) → Dense (3, Softmax). This produces a probability distribution across three classes: Benign, Malignant (HCC), and Metastatic. The model is compiled with Adam optimizer (learning rate = 1×10^{-4}), categorical cross-entropy loss, and accuracy as the primary metric. Class weights are computed using inverse frequency to address dataset imbalance. Training runs for 30 epochs with early stopping (patience = 5) based on validation loss.

E. Grad-CAM Interpretability

Gradient-weighted Class Activation Mapping (Grad-CAM) is integrated to improve clinical transparency. Grad-CAM computes the gradient of the predicted class score with respect to the final convolutional feature map to generate a spatial heatmap highlighting the image regions that most influenced the prediction. The importance weight for feature map k is:

$$\alpha^c = (1/Z)\ \Sigma_{ij}(\partial y^c / \partial A^k_{ij}) \quad ...(4)$$

$$L\_GradCAM = ReLU(\Sigma\ \alpha^c \cdot A^k) \quad ...(5)$$

The heatmap is upsampled and overlaid on the original CT image, allowing clinicians to verify that the model is attending to lesion regions rather than background tissue. This overlay is displayed alongside the prediction and confidence score in the Streamlit interface.

F. Streamlit Deployment

The complete pipeline is deployed as an offline, browseraccessible application using Streamlit. Users can upload a CT image, select the imaging phase, and receive: (a) the five-stage preprocessing output images; (b) predicted class with confidence bar; (c) Grad-CAM overlay highlighting the discriminative region; and (d) brief clinical description of the detected lesion type. All inference runs locally — no images are transmitted to external servers — preserving patient data privacy. The system can process both singlephase and multi-phase inputs.

IV. RESULTS AND DISCUSSION

A. Preprocessing Effect

The impact of each preprocessing stage was evaluated by comparing model performance when trained on raw CT images versus images processed through the complete five-stage pipeline. Raw images yielded 84.7% accuracy with poor recall on the Metastatic class (0.61), as subtle boundary features were not captured. After applying the full pipeline, overall accuracy improved to 96.4% and Metastatic recall rose to 0.89. The most significant individual contribution came from multi-scale fusion (Stage 4), which improved accuracy by 4.2 percentage points over Stage 3 alone. Guided bilateral filtering (Stage 5) reduced false positives by 3.1% by eliminating spurious edge artifacts.

B. Classification Performance

The system was evaluated on a held-out test set. Table III presents per-class and overall classification metrics. EfficientNet-B3 with the proposed preprocessing achieves 96.4% overall accuracy and a weighted F1-score of 0.96. The Benign class achieves the highest precision (0.98), while the Malignant class shows the highest recall (0.97). The Metastatic class presents the most difficulty due to structural overlap with HCC, but achieves 0.91 F1-score significantly better than baselines that do not use multiphase data.

Table 3. Per-class classification results

Class	Precision	Recall	F1-Score	Support

Benign	0.98	0.96	0.97	180
lalignant (HCC)	0.95	0.97	0.96	210
Metastatic	0.94	0.89	0.91	95
Macro Avg	0.96	0.94	0.95	485
Weighted Avg	0.96	0.96	0.96	485

C. Comparative Analysis

Table IV compares the proposed system against existing approaches reported in the literature. The proposed method achieves the highest accuracy (96.4%) and F1-score (0.96) across all compared methods, while supporting three-class classification. Traditional ResNet-50 and DenseNet-121 baselines without edge-aware preprocessing achieve only 89.4% and 91.2% respectively. The proposed method is also computationally lighter than transformer-based approaches, making it more suitable for offline deployment on standard hardware.

Table 4. Comparative analysis with existing Methods

Reference	Backbone	Preprocessing	Classes	Accuracy (%)	F1-score
ı et al. [8]	ResNet-50	None	2 (Binary)	89.40	0.88
Ying et al. [6]	DenseNet121	Global-only	3 (Multi)	91.20	0.90
Sun et al. [3]	Swin-Transformer	Cross-phase attn.	4 (Multi)	93.50	0.92
ı et al. [7]	Multistage CNN	Multi-phase	6 (Multi)	95.10	0.94
Proposed	EfficientNet-B3	5-Stage Edge-Aware	3 (Multi)	96.40	0.96

D. Grad-CAM Analysis

Grad-CAM overlays were generated for all test images and reviewed for clinical consistency. For correctly classified Malignant and Metastatic cases, the heatmap activation was consistently concentrated within lesion boundaries, demonstrating that the model learned diagnostically meaningful features rather than background artifacts. For Benign cases, activations were distributed across the lesion interior and surrounding margins, consistent with the smooth, well-defined appearance of hemangiomas. Misclassified cases (Metastatic predicted as Malignant) showed activation in structurally similar regions, indicating that errors arise from genuine visual ambiguity rather than model failure.

E. Streamlit Interface

The Streamlit interface was tested with 15 radiologists and medical students. Users found the preprocessing output display useful for understanding model behavior (rated 4.3/5.0). The Grad-CAM overlay was rated as helpful for verifying predictions (4.5/5.0). Average prediction time per image was 1.8 seconds on a standard laptop CPU, well within clinical workflow requirements. All image processing is performed locally, with no data leaving the user's machine.

V. CONCLUSION

This paper presented a Deep Learning-Enabled EdgeAware Multi-Phase Image Enhancement framework for precision liver lesion classification. The system integrates a five-stage preprocessing pipeline — HU windowing, CLAHE, Sobel edge extraction, multi-scale fusion, and guided bilateral filtering — with EfficientNet-B3 classification and Grad-CAM interpretability. The preprocessing pipeline addresses the key challenges of low CT image contrast, noise, and poor lesion boundary visibility that limit the performance of conventional approaches. Multi-phase CT integration further improves the model's ability to distinguish between lesion types based on dynamic contrast behavior.

The system achieves 96.4% accuracy and a weighted F1score of 0.96 on a three-class classification task, outperforming ResNet-50 (89.4%), DenseNet-121 (91.2%), and a Swin-Transformer baseline (93.5%). GradCAM analysis confirms that model predictions are grounded in diagnostically relevant regions, building clinical trust. The offline Streamlit deployment preserves patient data privacy while enabling real-time use in resource-constrained healthcare settings.

Future work will focus on: (1) extending to eight-class classification including Angiosarcoma and Cholangiocarcinoma; (2) 3D volumetric CT analysis using 3D-EfficientNet or nnU-Net; (3) incorporating attention mechanisms for adaptive multi-phase feature weighting; (4) validation on larger multicenter clinical datasets; and (5) integration of an AI clinical assistant module that explains predictions in natural language for non-specialist users.

REFERENCES

1. T. Zhang, Y. Liu, et al., "Edge-guided multi-scale adaptive feature fusion network for liver tumor segmentation," Scientific Reports, vol. 14, no. 1, 2024.
2. Z. Lu et al., "Detail-enhanced multi-scale feature fusion network for liver lesion segmentation," Biomedical Signal Processing and Control, vol. 90, 2025.

3. X. Sun, H. Zhao, et al., "Plug-and-play cross-phase attention module for multi-phase CT liver lesion classification," IEEE Transactions on Medical Imaging, 2025.
4. S. Feng, L. Zhao, et al., "LiverVoxNet: A large-scale annotated multi-phase CT dataset for liver lesion detection," Scientific Data, 2025.
5. R. Kumar, A. Sharma, et al., "LI-RADS integrated deep learning for hepatocellular carcinoma diagnosis," Journal of Hepatology, 2025.
6. H. Ying et al., "A multicenter clinical AI system study for detection and diagnosis of focal liver lesions," Nature Communications, vol. 15, no. 1, p. 1131, 2024.
7. Z. Liu, Y. Chen, et al., "multi-stage sequential deep learning framework for multi-phase CT liver lesion analysis," Medical Image Analysis, vol. 92, 2024.
8. L. Ma, C. Li, et al., "Deep learning model based on contrast enhanced MRI for predicting post-surgical survival in HCC," Heliyon, vol. 10, no. 11, 2024.
9. J. Ma, B. Lin, and Z. Ma, "Dual branch multi-scale feature fusion method for liver tumor detection," Array, vol. 28, 2025.
10. [V. S. Zossou et al., "Radiomics-based classification of tumor and healthy liver on computed tomography images," Cancers, vol. 16, no. 6, 2024.
11. Y. N. Gowda and M. R. V. Manjunath, "Automatic liver tumor classification using UNet70: a deep learning model," Journal of Liver Transplantation, vol. 18, 2025.
12. P. Bilic, P. Christ, et al., "The Liver Tumor Segmentation Benchmark (LiTS)," Medical Image Analysis, vol. 84, 2023. Y. Chen, C. Zheng, et al., "MS-FANet: Multi-scale feature attention network for liver tumor segmentation," Computers in Biology and Medicine, vol. 163, 2023.