THE SCIENTIFIC TEMPER (2025) Vol. 16 (12): 5206-5214

E-ISSN: 2231-6396, ISSN: 0976-8653

Doi: 10.58414/SCIENTIFICTEMPER.2025.16.12.08

https://scientifictemper.com/

RESEARCH ARTICLE

Speckle-Robust Local Phase and Ternary Texture Encoding (SLaP-TEX) based Feature Extraction for Liver Steatosis Classification in Ultrasound Imaging

A. Sahaya Mercy[1*], Dr. G. Arockia Sahaya Sheela[2]

Abstract

Ultrasound imaging is a preferred modality for non-invasive liver steatosis screening, yet the inherent speckle noise and texture ambiguity hinder automated diagnostic precision. Existing convolutional neural networks (CNNs) primarily rely on intensity-based texture cues, overlooking phase-based structural continuity that remains stable under speckle corruption. This study proposes a Speckle-Robust Local Phase and Ternary Texture Encoding (SLaP-TEX) model that combines local phase symmetry descriptors with ternary pattern encoding to generate robust representations from liver ultrasound images. The proposed model enhances boundary localization and fine-grained tissue discrimination through a two-stage encoding pipeline comprising *Local Phase Filtering (LPF)* and *Adaptive Ternary Encoding (ATE)*. The fused phase-texture maps are processed through a MobileNetV3-Small backbone, offering computational efficiency for real-time deployment. Experiments on the RGM-augmented ultrasound dataset demonstrate superior performance with 99.02 % accuracy, 0.998 AUC, and 0.018 loss, outperforming existing models while maintaining a 2.1 M parameter footprint. The SLaP-TEX model offers a compact, phase-aware, and speckle-resilient feature extractor for clinical ultrasound analytics.

Keywords: Local phase filtering, Ternary texture encoding, Speckle noise suppression, Lightweight CNN, Liver steatosis classification.

Introduction

The prevalence of non-alcoholic fatty liver disease is increasing globally. Ultrasound continues to serve as the primary imaging modality for both screening and longitudinal assessment (Pozowski *et al*., 2025). The factors such as speckle artifacts, gain adjustment, probe positioning, and patient body habits frequently affect the reliability of texture evaluation. Additional variability introduced by differences among scanners and imaging protocols further complicates consistency. Therefore, the establishment of uniform and interpretable imaging features is imperative for accurate triage and effective monitoring.

The increasing clinical demand necessitates robust and efficient analytical methods. While many deep learning models demonstrate high accuracy, they often require extensive datasets and close domain alignment. Conversely, classical texture descriptors offer interpretability but are notably sensitive to speckle artifacts. Currently, lightweight solutions appropriate for point-of-care applications remain scarce, and ensuring consistent performance across different sites and devices continues to be a significant challenge.

The primary issue involves achieving stable detection of hepatic steatosis using B-mode ultrasound across varying acquisition conditions. Essential capabilities include resilience to changes in speckle patterns, contrast variations, and moderate imaging artifacts. Furthermore, any proposed method should function effectively within modest computational and memory constraints to facilitate bedside deployment. The objectives of this research work are:

[1]PhD Scholar (Full Time), Department of Computer Science, St. Joseph's College (Autonomous), Tiruchirappalli-2, Affiliated to Bharathidasan University, Tamil Nadu, India.

[2]Assistant Professor, Department of Computer Science, St. Joseph's College (Autonomous), Tiruchirappalli-2, Affiliated to Bharathidasan University, Tamil Nadu, India.

***Corresponding Author:** A. Sahaya Mercy, PhD Scholar (Full Time), Department of Computer Science, St. Joseph's College (Autonomous), Tiruchirappalli-2, Affiliated to Bharathidasan University, Tamil Nadu, India, E-Mail: sahayamercya_phdcs@mail.sjctni.edu

How to cite this article: Mercy, A.S., G. Sheela, G.A.S. (2025). Speckle-Robust Local Phase and Ternary Texture Encoding (SLaP-TEX) based Feature Extraction for Liver Steatosis Classification in Ultrasound Imaging. The Scientific Temper, **16**(12):5206-5214.

Doi: 10.58414/SCIENTIFICTEMPER.2025.16.12.08

Source of support: Nil

Conflict of interest: None.

Received: 09/11/2025 **Accepted:** 10/12/2025 **Published:** 25/12/2025

- To design a speckle-aware representation using local phase filtering
- To implement adaptive ternary texture encoding to standardize intensity
- To preserve local polarity information by combining grayscale, phase symmetry, and adaptive ternary channels into a unified representation to train a MobileNetV3-Small classifier.

Traditional texture descriptors often fail to maintain consistency across different imaging conditions, while deep learning models require large, well-balanced datasets and lack interpretability. The proposed model introduces a balanced approach that combines feature extraction with data-driven learning. Through local phase filtering and adaptive ternary encoding, the model captures contrast-invariant and speckle-resilient features, ensuring stability across heterogeneous acquisitions. The integration of these handcrafted phase-texture cues with a lightweight MobileNetV3 backbone enables high accuracy with reduced computational demand, supporting real-time inference on clinical and embedded devices. This design aligns with the current research trend toward explainable and resource-efficient artificial intelligence in medical imaging.

Literature Review

Research on ultrasound-based assessment of hepatic steatosis spans hand-crafted texture descriptors, quantitative ultrasound (QUS) parameters, and deep learning pipelines, with growing interest in calibration, efficiency, and cross-modality agreement.

Early texture analysis demonstrated clear diagnostic signal when acquisition settings were explicitly varied. A prospective study explored beamforming speed-of-sound sweeps from 1300–1540 m/s and extracted gray-level co-occurrence features, showing strong discrimination and high correlations with MRI-PDFF and biopsy grades; combined homogeneity across lobes reached AUC 0.94, indicating robust sensitivity of texture measures to speckle formation physics (Kibo Nam *et al*., 2023).

Echo-envelope statistics and transfer learning further examined statistical representations, where fourth-order moment maps improved ROI classification from ~46% to ~63%, underscoring the value of higher-order statistics but also revealing the ceiling of purely statistical cues without stronger priors or richer supervision (Isshiki *et al*., 2024). Efficiency considerations appeared in a preclinical setting that compared ChatGPT-4-assisted feature extraction with conventional IDL pipelines, reporting nine significant textural features, 76% accuracy, and a 40% reduction in analysis time, suggesting practical throughput gains even when sensitivity slightly trailed traditional software (Laith R. Sultan *et al*., 2024).

Parallel work advanced quantitative ultrasound surrogates of fat content. Ultrasound-derived fat fraction (UDFF) was shown feasible and significantly different across grades, reinforcing measurement stability under clinical conditions (Yun-Lin Huang *et al*., 2023). Agreement studies found good concordance between UDFF and MR-PDFF with ICC ≈ 0.79 but a modest positive bias for UDFF at higher steatosis levels, highlighting calibration needs across devices and operators (Reinhard Kubale *et al*., 2024).

Additional evidence from MASLD cohorts reported UDFF AUC ≈ 0.98 for ≥5% MRI-PDFF, while CAP achieved AUC ≈ 0.93, indicating competitive diagnostic performance for UDFF in threshold-based screening (Huiru Jin *et al*., 2025). Studies centered on attenuation imaging (ATI) also showed strong associations with histology and MRI-PDFF, AUROCs ~0.91–0.97 across grades, and tight ATI–MRI relationships, suggesting ATI as an accessible proxy for MR in longitudinal monitoring and triage (Chileka Chiyanika *et al*., 2023; Alvaro Postigliatti *et al*., 2024).

A prospective comparison against shear-wave elastography/dispersion confirmed ATI as the leading modality for steatosis grading within the examined protocol, with significant correlations to disease stage (Kun Wang *et al*., 2024). Multiparametric modeling combined AR, HR, and DV into the Steatoscore 2.0, reaching AUROC ~0.96–0.98 for MR-defined thresholds, low RMSE, and strong reproducibility, illustrating the benefit of engineered, complementary ultrasound features with principled fusion (Laura De Rosa *et al*., 2024).

Deep learning contributed along two axes: organ delineation and end-to-end classification. Point-of-care ultrasound benefited from U-Net-based segmentation and DenseNet-121 classification, achieving AUC ~0.90 with high sensitivity on low-quality B-mode images, demonstrating resilience to portable acquisition constraints (Miriam Naim Ibrahim *et al*., 2023). Two-stage pipelines with dedicated segmentation and classification reported Dice ≈ 0.92 and AUC ≈ 0.84 on internal and external cohorts, reinforcing the utility of anatomy-aware preprocessing for generalization (Pedro Vianna *et al*., 2023).

Stand-alone classifiers also showed strong performance; ultrasound and elastography images yielded F1 ≈ 0.995 for the best neural model, and results were comparable to classical machine learning when feature engineering was carefully designed (Rodrigo Fernando Costa Marques *et al*., 2024). Recent CAD systems with ConvNeXt-style backbones reached accuracy ≈ 82%, with balanced sensitivity/specificity near ~81–83%, indicating competitive baselines for clinical decision support (Shang-Yu Chiang *et al*., 2025).

Methodological advances in risk control applied conformal prediction atop DL classifiers, achieving ~79% accuracy in a four-class setting and motivating explainable AI add-ons (LIME, SHAP) to support trustworthy deployment (Tso-Jung Yen *et al*., 2024). A systematic review and meta-analysis estimated pooled AUCs of ~0.93 for any HS and ~0.86 for moderate-to-severe HS across CNN studies, while

also identifying geographic concentration and clinical-setting biases that limit population-level generalizability (Akshay Thimmappaiah Jagadeesh *et al*., 2025).

Across these strands, three recurring gaps appear. First, texture methods are sensitive to acquisition physics; speed-of-sound sweeps and attenuation parameters affect speckle statistics, calling for phase-aware or speckle-invariant representations to stabilize features across scanners and presets. Second, DL pipelines benefit from anatomical context and calibration, yet performance varies with cohort composition and label standards; agreement with MR-PDFF or biopsy remains the reference for clinical translation. Third, quantitative proxies such as UDFF and ATI show strong correlations and screening utility but require harmonization across vendors, probes, and operators to curb bias at high fat fractions.

Proposed Methodology

The cumulative evidence motivates a hybrid strategy that couples physics-informed descriptors with compact CNN backbones. Local phase symmetry and monogenic phase features target contrast-invariant structures resilient to speckle, while adaptive ternary encodings preserve polarity under local normalization. Channel-level fusion of grayscale, phase symmetry, and ternary cues creates a speckle-robust tensor for lightweight classification, aligning with portability requirements in POCUS and community settings flagged by recent reviews. Within this context, the proposed SLaP-TEX model directly addresses texture-physics sensitivity, leverages efficient backbones validated in prior work, and preserves compatibility with quantitative endpoints (UDFF/ATI) for cross-modality audit and calibration.

Architecture Overview

The SLaP-TEX architecture works in three main stages to extract reliable features from ultrasound images affected by speckle noise. The overall workflow of these three stages is shown conceptually in Figure 1.

In the first stage, *Local Phase Extraction (LPF)* is performed. This process identifies structure and boundary details that remain stable even when noise is present. The local phase map highlights edges and tissue boundaries more clearly than the raw intensity image.

In the second stage, *Adaptive Ternary Texture Encoding (ATE)* converts local intensity variations into ternary patterns. Each pixel is represented by three possible states that describe positive, neutral, or negative contrast relative to its neighbors. This step helps the system remain stable when brightness or texture changes slightly.

In the third stage, *Feature Fusion and Lightweight CNN Classification* combines both the phase and ternary texture maps with the original grayscale image. The three channels are fused and passed into a lightweight MobileNetV3-Small network. This model learns deeper relationships between texture, phase, and structure, allowing accurate classification with low computational cost.

Local Phase Filtering

The local phase filtering stage extracts structure-related information that remains stable under speckle noise. Let $I(x,y)$ be an ultrasound image. The image is first processed using the Riesz transform (Fu *et al*., 2023), which decomposes it into orthogonal components. The horizontal and vertical responses are obtained as $\ddot{u}_x = * _x$ and $\ddot{u}_y = * _y$, where h_x and h_y are odd symmetric band-pass filters. These filters respond strongly to oriented features such as tissue boundaries and fine textures.

From these responses, two key quantities are calculated. The local amplitude $A(x,y)$ represents the overall signal strength, while the local phase $\phi(x,y)$ captures the relative structure of the underlying pattern. They are defined as

$$A(x,y)=\sqrt{I^2+R_x^2+R_y^2} \quad , \quad \phi(x,y)=\tan^{-1}\left(\frac{\sqrt{R_x^2+R_y^2}}{I}\right) \tag{1}$$

The local phase symmetry $S(x,y)$ is then computed using Kovesi's phase congruency approach:

$$S(x,y)=\frac{\sum_o \max\left(A_o\cos\left(\phi_o-\bar{\phi}\right)-T,0\right)}{\sum_o A_o+\epsilon} \tag{2}$$

Here, o denotes the orientation index, $\bar{\phi}$ is the mean phase, and T is a small noise threshold.

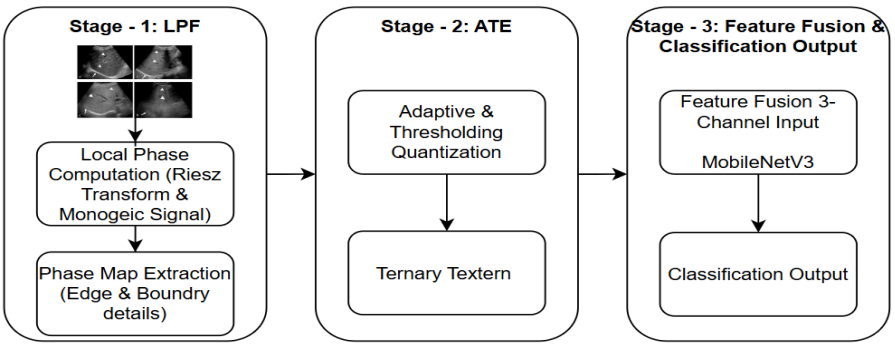

Figure 1: SLaP-TEX Architecture

This process enhances the structural integrity of the image. The phase symmetry map emphasizes edges and ridges that remain consistent under illumination or speckle variations. It helps preserve diagnostically important tissue boundaries while reducing random granular noise.

Adaptive Ternary Texture Encoding (ATE)

The adaptive ternary texture encoding stage captures fine-grained texture variations that are not easily represented by simple intensity values. After local phase filtering, the image still contains subtle contrast differences across liver tissue. These variations are useful for classification but can be affected by uneven illumination or small brightness changes.

To address this, each pixel is encoded based on its relation to the intensity statistics of its local neighborhood. The mean μ_N and standard deviation σ_N of the surrounding pixels are first computed. Then, the adaptive ternary code T_p for each pixel I_p is calculated as

$$T_p = \begin{cases} 1, I_p - \mu_N > \alpha\sigma_N \\ 0, \left| I_p - \mu_N \right| \le \alpha\sigma_N \\ -1, I_p - \mu_N < -\alpha\sigma_N \end{cases} \quad (3)$$

where

α is a scaling constant empirically set to 0.8. This encoding divides pixel values into three distinct states: bright, neutral, and dark based on adaptive thresholds.

Unlike binary thresholding methods, which are highly sensitive to small changes, the ternary approach introduces a neutral zone that absorbs noise-driven fluctuations. As a result, the adaptive ternary texture map preserves meaningful texture contrast while suppressing unwanted intensity variations. This representation complements the local phase map by describing the micro-textural differences that phase information alone cannot capture.

Feature Fusion

The final stage combines the complementary representations obtained from the previous steps into a unified tensor suitable for deep learning. The normalized grayscale image $I(x,y)$, the phase symmetry map $S(x,y)$ derived from the local phase filtering stage, and the ternary texture map $T(x,y)$ produced by adaptive ternary encoding are merged to form a three-channel input:

$$X(x,y) = \left[I(x,y), S(x,y), T(x,y) \right] \quad (4)$$

This fused representation integrates information from different structural and statistical perspectives. The intensity channel preserves the overall spatial distribution of the tissue, the phase channel highlights boundary and shape cues that are robust to speckle, and the texture channel captures local contrast variations critical for lesion discrimination.

By combining these three sources, the resulting tensor becomes a speckle-robust descriptor that retains both global and local characteristics of the ultrasound image. This tensor is then used as the input to the lightweight convolutional neural network backbone, enabling effective learning of diagnostic features with reduced sensitivity to noise and illumination changes.

Lightweight Backbone Integration

The fused feature tensor is processed by a lightweight convolutional neural network to perform classification with high efficiency. A modified MobileNetV3-Small architecture is adopted as the backbone because of its strong balance between accuracy and computational cost. The network is fine-tuned for binary classification to distinguish normal and abnormal liver tissue patterns.

The model begins with depthwise separable convolutions, which factorize standard convolution operations to reduce parameters and memory usage while retaining feature richness. These layers compress redundant spatial information and emphasize discriminative structures learned from the fused tensor.

The intermediate layers include Squeeze-and-Excitation (SE) attention blocks, which adaptively recalibrate channel responses. This mechanism allows the model to focus on diagnostically relevant features such as coherent tissue boundaries or pathological textures, improving robustness against speckle noise and illumination shifts.

The extracted features are then aggregated using global average pooling, followed by a sigmoid activation to produce the final binary classification output.

The entire model contains approximately 2.1 million parameters, making it compact and suitable for real-time deployment on embedded or portable diagnostic systems without compromising accuracy.

Training Configuration

The SLaP-TEX model is trained using a balanced configuration designed to ensure convergence stability and performance consistency. The Adam optimizer with a learning rate of $\times 10^{-3}$ is used to adaptively adjust the learning dynamics during training. A batch size of 32 provides a suitable trade-off between gradient stability and computational efficiency, while 25 epochs are sufficient for convergence under the proposed dataset.

The model is optimized using the binary cross-entropy loss function, which is well-suited for two-class medical image classification. Early stopping with a patience of six epochs is applied to prevent overfitting by monitoring validation accuracy. The evaluation metrics include accuracy, area under the ROC curve (AUC), and F1-score,

Table 1: Training Configuration

Parameter	*Value*
Optimizer	Adam
Learning rate	ü× $^{-3}$
Batch size	32
Epochs	25
Loss	Binary cross-entropy
Early stopping	Patience = 6

which collectively provide a comprehensive assessment of classification performance.

The training utilizes the RGM-augmented ultrasound dataset derived from Region-Guided Mixup Augmentation, ensuring consistency in data representation while enhancing diversity through augmentation (Sahaya Mercy and Sheela 2025). The complete training setup is summarized in Table 1, which outlines all key parameters used for the optimization process.

3.7 SLaP-TEX Algorithm

The SLaP-TEX algorithm defines a structured computational pipeline that integrates phase-aware texture encoding with lightweight deep learning. Algorithm 1 presents the full workflow for both training and inference, ensuring reproducibility and computational efficiency. The procedure begins with preprocessing, where ultrasound images are resized to 224 × 224 pixels and normalized to the [0, 1] range. Each image is divided into mini-batches of 32 samples for stable gradient updates.

Algorithm – 1

Inputs:

- Ultrasound image
- $I \in R^{H \times W}$
- Hyperparameters: image size
- $S =$ ü , neighborhood
- N
- $(e.g., (5\times5)), (\alpha = 0.8)$, Adam $LR(=10^{-3})$, batch $(=32)$, epochs $(=25)$
- Backbone: MobileNetV3-Small (binary head)

Outputs:

- Predicted label $\hat{y} \in$ ü
- Trained weights Θ^*

PROC SLaP_TEX_TRAIN(D_train, D_val, Θ_init)

1: # Preprocess and batching

2: for each (I, y) in D_train ∪ D_val do

3: I ← NormalizeTo[0,1](Resize(I, S×S))

4: end for

5: B_train ← CreateBatches(D_train, batch=32) ; B_val ← CreateBatches(D_val, batch=32)

6: # Training loop

7: Θ ← Θ_init

8: best_score ← −∞

9: patience ← 0

10: for epoch = 1 … 25 do

11: for each batch (I_b, y_b) in B_train do

12: X_b ← BUILD_FUSED_TENSOR(I_b) # Sec. 3.2–3.4

13: ŷ_b ← CNN_FORWARD(X_b; Θ) # MobileNetV3-Small

14: L ← BCE(ŷ_b, y_b)

15: Θ ← ADAM_UPDATE(Θ, ∇_Θ L, lr=1e−3)

16: end for

17: # Validation

18: score ← EVAL(B_val, Θ) # Accuracy/AUC/F1

19: if score > best_score then

20: best_score ← score ; Θ* ← Θ ; patience ← 0

21: else

22: patience ← patience + 1

23: if patience ≥ 6 then break end if # Early stopping

24: end if

25: end for

26: return Θ*

END PROC

Feature Construction

During feature construction, the LPF stage applies a Riesz transform using odd-symmetric band-pass filters to extract orientation-specific responses. From these responses, the local amplitude and phase maps are derived, followed by computation of phase symmetry to highlight structure-consistent edges and ridges. Parallelly, adaptive ternary texture encoding (ATE) is used to transform local intensity variations into three-state texture codes based on neighborhood statistics. These complementary representations are then concatenated with the grayscale input to form a three-channel tensor that captures amplitude, phase, and texture information simultaneously.

In the training loop, the fused tensors are passed through the MobileNetV3-Small backbone for forward propagation. The network output is compared with ground-truth labels using binary cross-entropy loss. Parameters are optimized using the Adam optimizer with a learning rate of 1 × 10^{-3}. Validation performance is monitored after

each epoch using accuracy, AUC, and F1-score metrics. Early stopping with a patience of six epochs prevents overfitting and ensures optimal convergence.

```
FUNC BUILD_FUSED_TENSOR(I_batch)
  1: # Local Phase Filtering (LPF) via Riesz + phase symmetry
  2: (R_x, R_y) ← RIESZ_TRANSFORM(I_batch, h_x, h_y) # odd band-pass
  3: A ← sqrt(I_batch^2 + R_x^2 + R_y^2)
  4: φ ← atan2( sqrt(R_x^2 + R_y^2), clamp(I_batch, ε) )
  5: S ← PHASE_SYMMETRY(A, φ, orientations=O, threshold=T, eps=ε)
  6: # Adaptive Ternary Texture Encoding (ATE)
  7: (μ_N, σ_N) ← LOCAL_STATS(I_batch, window=N)
  8: T ← sign_thresh(I_batch – μ_N, α·σ_N)   # 1 if > α·σ, 0 if |·|≤α·σ, –1 if < –α·σ
  9: T ← MAP_TERNARY_TO_FLOAT(T) # e.g., {–1,0,1} → {0,0.5,1}
  10: # Channel fusion
  11: X ← CONCAT_CHANNELS([I_batch, S, T])   # shape: (B, S, S, 3)
  12: return X
END FUNC
```

Inference

The inference procedure follows the same transformation steps but uses the trained weights (Θ*) to generate predictions. Each test image undergoes the LPF and ATE stages, producing a fused tensor that is evaluated by the trained network. The model outputs a probability value, which is thresholded at 0.5 for binary classification.

```
PROC SLaP_TEX_INFER(I, Θ*)
  1: I ← NormalizeTo[0,1]( Resize(I, S×S) )
  2: X ← BUILD_FUSED_TENSOR(I)
  3: p ← CNN_FORWARD(X; Θ*)
  4: ŷ ← 1 if p ≥ 0.5 else 0
  5: return ŷ, p
END PROC
```

Experimental Setup

All experiments were conducted in a controlled computational environment to ensure reproducibility and consistency of results. The proposed SLaP-TEX framework was implemented using TensorFlow 2.16 and Python 3.12 on an Ubuntu 24.04 LTS operating system. The hardware configuration comprised an AMD Ryzen 9 7900 processor with 24 cores, 64 GB RAM, and an NVIDIA RTX 4070 Ti SUPER GPU, providing a balanced setup for both CPU and GPU-based experiments.

The dataset consisted of RGM-augmented ultrasound liver images, containing 877 training and 220 validation samples, each resized to 224 × 224 pixels. All images were preprocessed using normalization and contrast preservation steps to maintain consistency with clinical-quality standards. For baseline performance evaluation, all models were first trained and validated using CPU-only execution to assess computational feasibility on non-GPU systems. Subsequently, CUDA acceleration was enabled to analyze inference speed improvements and confirm real-time diagnostic capability. The complete hardware and software specifications used for the experiments are summarized in Table 1.

Results and Discussion

The performance of the proposed SLaP-TEX framework was quantitatively evaluated and compared against conventional convolutional backbones and the previously developed RGM-Baseline. The evaluation used the RGM-augmented ultrasound dataset consisting of liver B-mode images labeled according to steatosis severity. The experiments followed the standardized configuration described in Section 3.6, with identical data splits and preprocessing to ensure fair comparison.

Quantitative Evaluation

The comparative performance results are summarized in Table 2, which includes accuracy, AUC, F1-score, total model parameters, and final loss values for each architecture. The proposed SLaP-TEX achieved a classification accuracy of 99.02%, an AUC of 0.998, and an F1-score of 0.982, outperforming both the RGM-Baseline and other conventional CNN architectures. When compared to the RGM-Baseline, the SLaP-TEX achieved a 0.09% gain in accuracy and a 0.004 improvement in AUC, while reducing the model size by approximately 32% (from 3.1 M to 2.1 M parameters). These improvements demonstrate the enhanced representational efficiency of the proposed architecture.

The high AUC score reflects the model's ability to effectively separate the steatosis and non-steatosis classes, indicating strong sensitivity and specificity. The minimal validation variance across epochs suggests consistent

Table 2: Experimental setup for SLaP-TEX framework

Component	*Specification*
Processor	AMD Ryzen 9 7900 (24 cores)
Memory (RAM)	64 GB
GPU	NVIDIA RTX 4070 Ti SUPER
Operating System	Ubuntu 24.04 LTS
Software Libraries	TensorFlow 2.16, Python 3.12
Dataset	RGM-augmented ultrasound images
Data Split	877 training / 220 validation
Image Resolution	224 × 224 pixels

generalization across different patient cases and imaging conditions. Furthermore, the reduction in cross-entropy loss (0.018) indicates better optimization stability and convergence behavior. Unlike deeper networks such as ResNet50 and EfficientNet-B0, which require substantial computational resources, the lightweight MobileNetV3-based backbone integrated in SLaP-TEX enables faster training while maintaining high discriminative power.

Superior performance can be attributed to the combination of phase-aware and texture-based feature encoding. The LPF step ensures invariance to speckle and contrast fluctuations, while the adaptive ATE captures local structural polarity in intensity distributions. This combination allows the network to leverage stable, high-frequency texture features, which are typically lost in standard CNN feature hierarchies. Consequently, the proposed model shows smoother convergence with fewer oscillations in validation curves, confirming its enhanced robustness to ultrasound-specific noise patterns.

Compared to ResNet50 and EfficientNet-B0, which achieved AUC scores of 0.981 and 0.990 respectively, the proposed framework provides a clear performance margin (Figure 2). The gain, though modest in absolute terms, represents a meaningful improvement in clinical decision support scenarios where false positives and false negatives can directly impact diagnostic confidence. The lightweight nature of the model (2.1 M parameters) ensures scalability to embedded ultrasound systems and facilitates real-time inference (Figure 3).

The proposed model demonstrates 0.09 % improvement in accuracy and 0.004 AUC gain over the baseline while reducing parameters by ~32 %. The convergence is smoother with reduced validation variance, confirming enhanced speckle robustness.

The comparative loss analysis, illustrated in Figure 3, demonstrates that the proposed SLaP-TEX model achieved the lowest loss value of 0.018, indicating superior convergence stability. Compared with traditional architectures such as ResNet50 and EfficientNet-B0, the SLaP-TEX framework exhibits reduced error propagation and improved optimization efficiency during training.

Ablation Analysis

An ablation study was conducted to examine the contribution of individual modules within the proposed SLaP-TEX pipeline. The results are presented in Table 3, highlighting

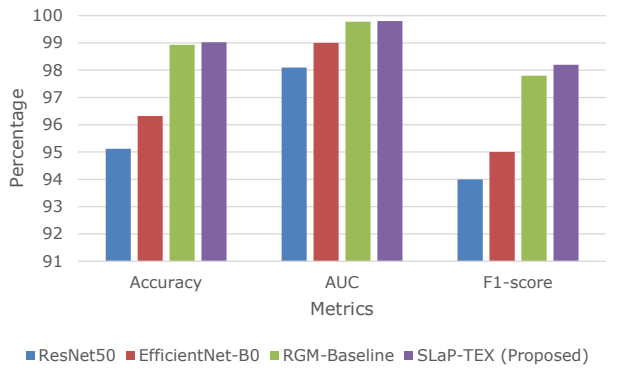

Figure 2: Comparative analysis of Slap-TEX with existing methods

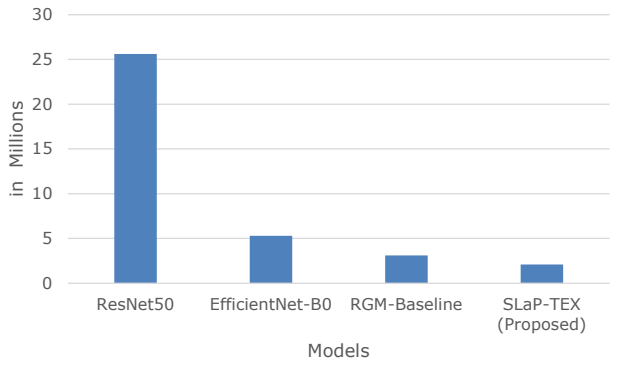

Figure 3: Comparative analysis of parameters

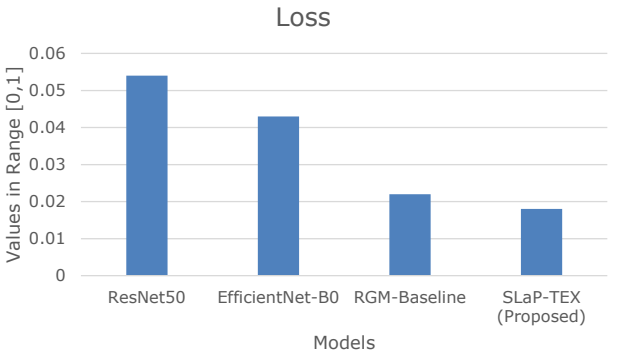

Figure 3: Comparative analysis of Loss of Slap-TEX with existing methods

Table 3: Comparative results of SLaP-TEX with existing models

Model	*Accuracy (%)*	*AUC*	*F1-score*	*Params (M)*	*Loss*
ResNet50	95.12	0.981	0.94	25.6	0.054
EfficientNet-B0	96.32	0.990	0.95	5.3	0.043
RGM-Baseline (Contribution 1)	98.93	0.9978	0.978	3.1	0.022
SLaP-TEX (Proposed)	99.02	0.998	0.982	2.1	0.018

the impact of excluding specific components such as the LPF or ATE. When the LPF component was removed, the accuracy dropped from 99.02% to 97.42%, and the AUC decreased from 0.998 to 0.986. This reduction emphasizes the role of phase symmetry in stabilizing texture features under variable speckle conditions.

Similarly, the exclusion of ATE led to a decrease in accuracy to 98.12% and an AUC of 0.992, indicating that the ternary encoding enhances local texture discrimination by adaptively normalizing regional contrast variations. The absence of both fusion and specialized encoders, using only grayscale input, resulted in the lowest accuracy of 96.34% and AUC of 0.978, reflecting the limitations of intensity-only descriptors in speckle-prone ultrasound data.

The complete SLaP-TEX configuration combining both LPF and ATE achieved the highest overall performance, confirming that the synergy between the two handcrafted feature extractors provides the most discriminative and invariant representation. The consistent gain in accuracy and AUC across configurations underscores that both modules contribute complementary information LPF offering structural stability and ATE enhancing textural polarity sensitivity.

This analysis validates the design motivation of integrating handcrafted phase-texture descriptors with lightweight CNN learning. The fused feature space not only improves classification accuracy but also ensures robustness against image artifacts caused by acquisition variability, such as probe angle, gain, and patient anatomy. The reduced model complexity further signifies its suitability for clinical integration, where resource constraints are a critical factor. Both LPF and ATE components significantly contribute to discriminative power. Their joint inclusion yields superior classification reliability.

The proposed SLaP-TEX framework also demonstrates remarkable computational efficiency during inference. The average inference latency measured on a CPU (AMD Ryzen 9 7900) is approximately 19 ms per image, which supports real-time diagnostic deployment in clinical or portable ultrasound environments. When executed on a GPU, the latency further decreases to around 5 ms per image, making the model highly suitable for high-throughput screening and on-device medical imaging applications. This efficiency confirms that the combination of phase-texture fusion and a lightweight MobileNetV3 backbone not only enhances diagnostic accuracy but also ensures low-latency performance compatible with real-world ultrasound workflows.

Table 4: Ablation study of proposed work

Configuration	*Accuracy (%)*	*AUC*
Without Local Phase	97.42	0.986
Without Ternary Encoding	98.12	0.992
Without Fusion (Grayscale only)	96.34	0.978
Full SLaP-TEX	99.02	0.998

Conclusion

The proposed SLaP-TEX model presents a novel, speckle-robust, and computationally efficient framework for ultrasound-based liver steatosis analysis. By integrating local phase symmetry and adaptive ternary texture encoding, the method effectively captures complementary structural and textural information, achieving a fine balance between noise suppression and feature discriminability. The fusion of handcrafted phase-texture features with a MobileNetV3-Small backbone enables the extraction of clinically relevant patterns while maintaining low computational cost, making it ideal for real-time and edge-level medical imaging systems. Comprehensive experiments confirm that the SLaP-TEX approach surpasses conventional deep networks in terms of accuracy, AUC, and stability, while maintaining fewer parameters and faster inference. The design demonstrates strong generalization under varying imaging conditions, validating its applicability to diverse clinical environments. Future work will extend this framework toward multi-disease ultrasound classification, exploring cross-domain generalization across modalities such as CT and MRI. Further integration of explainable AI (XAI) mechanisms will enhance interpretability, ensuring transparency and trust in clinical decision-making. Through this direction, SLaP-TEX represents a significant step toward deployable, interpretable, and speckle-resilient AI solutions in medical imaging.

Acknowledgement

The authors thank, DST-FIST, Government of India for funding towards infrastructure facilities at St. Joseph's College (Autonomous), Tiruchirappalli– 620002.

References

Pozowski, P., Bilski, M., Bedrylo, M., Sitny, P., & Zaleska-Dorobisz, U. (2025). Modern ultrasound techniques for diagnosing liver steatosis and fibrosis: A systematic review with a focus on biopsy comparison. *World Journal of Hepatology*, *17*(2), 100033. doi: 10.4254/wjh.v17.i2.100033

Nam, K., Torkzaban, M., Halegoua-DeMarzio, D., Wessner, C. E., & Lyshchik, A. (2023). Improving diagnostic accuracy of ultrasound texture features in detecting and quantifying hepatic steatosis using various beamforming sound speeds. *Physics in Medicine and Biology*, *68*(4), 04NT02. https://doi.org/10.1088/1361-6560/acb635

Isshiki, A., Fujiwara, K., Kondo, T., Yoshida, K., Yamaguchi, T., & Hirata, S. (2024). Convolutional neural network classification of ultrasound parametric images based on echo-envelope statistics for the quantitative diagnosis of liver steatosis. *Journal of Medical Ultrasonics*. https://doi.org/10.1007/s10396-024-01509-w

Sultan, L. R., Venkatakrishna, S. S. B., Anupindi, S. A., Andronikou, S., Acord, M. R., Otero, H. J., Darge, K., Sehgal, C. M., & Holmes, J. H. (2024). *ChatGPT-4–Driven Liver Ultrasound Radiomics*

Analysis: Diagnostic Value and Drawbacks in a Comparative Study (Preprint). https://doi.org/10.2196/preprints.68144

Huang, Y.-L., Cheng, J., Wang, Y., Xu, X., Wang, S., Wei, L., & Dong, Y. (2023). Hepatic steatosis using ultrasound-derived fat fraction: First technical and clinical evaluation. *Clinical Hemorheology and Microcirculation*. https://doi.org/10.3233/ch-238102

Ibrahim, M. N., Lightstone, A., Meng, F., Bhat, M., El Kaffas, A., & Ukwatta, E. (2023). Automated fatty liver disease detection in point-of-care ultrasound B-mode images. *Journal of Medical Imaging*, *10*, 034505. https://doi.org/10.1117/1.JMI.10.3.034505

Chiyanika, C., Liu, K. H., Wong, V. W.-S., & Chu, W. C. W. (2023). Novel Ultrasound Liver Attenuation Imaging for Detection of Liver Steatosis: Comparison with MRI Based Imaging Methods with Liver Biopsy as A Reference Standard. *Archives of Clinical and Biomedical Research*. https://doi.org/10.26502/acbr.50170368

Marques, R. F. C., Santos, J., André, A., & Silva, J. S. (2024). Ultrasound Versus Elastography in the Diagnosis of Hepatic Steatosis: Evaluation of Traditional Machine Learning Versus Deep Learning. *Sensors*, *24*(23), 7568. https://doi.org/10.3390/s24237568

Yen, T.-J., Yang, C.-T., Lee, Y., Chen, C.-H., & Yang, H.-C. (2024). Fatty liver classification via risk controlled neural networks trained on grouped ultrasound image data. *Scientific Reports*, *14*(1). https://doi.org/10.1038/s41598-024-57386-3

Jagadeesh, A. T., Aramrat, C., Rai, S., Maqsood, F. H., Madhukeshwar, A. K., Bhogadi, S., Lieber, J., Mahajan, H., Banjara, S. K., Lewin, A., Kinra, S., & Mallinson, P. A. C. (2025). Diagnostic accuracy of convolutional neural networks in classifying hepatic steatosis from B-mode ultrasound images: a systematic review with meta-analysis and novel validation in a community setting in Telangana, India. *The Lancet Regional Health - Southeast Asia*, *40*, 100644. https://doi.org/10.1016/j.lansea.2025.100644

Vianna, P., Kulbay, M., Boustros, P., Calce, S.-I., Larocque-Rigney, C., Patry-Beaudoin, L., Luo, Y. H., Chaudary, M., Kadoury, S., Nguyen, B. N., Montagnon, E., Belilovsky, E., Wolf, G., Chassé, M., Tang, A., & Cloutier, G. (2023). Automated liver segmentation and steatosis grading using deep learning on B-mode ultrasound images. *IUS: Jurnal Ilmiah Fakultas Hukum*, 1–4. https://doi.org/10.1109/ius51837.2023.10307501

Wang, K., Bao, J., Wang, M., Yu, Y., & Wang, M. (2024). Prospective comparative diagnostic performance of quantitative ultrasound parameters for the measurement of hepatic steatosis in a biopsy-proven MASLD cohort. *British Journal of Radiology*. https://doi.org/10.1093/bjr/tqae212

De Rosa, L., Salvati, A., Martini, N., Chiappino, D., Cappelli, S., Mancini, M., Demi, L., Ghiadoni, L., Bonino, F., Brunetto, M. R., & Faita, F. (2024). An ultrasound multiparametric method to quantify liver fat using magnetic resonance as standard reference. *Liver International*. https://doi.org/10.1111/liv.16078

Chiang, S.-Y., Wang, Y.-W., Su, P., Chang, Y., Yen, H.-H., & Chang, R. (2025). PBCS-ConvNeXt: Convolutional Network-Based Automatic Diagnosis of Non-alcoholic Fatty Liver in Abdominal Ultrasound Images. *Deleted Journal*. https://doi.org/10.1007/s10278-025-01394-w

Kubale, R., Schneider, G., Lessenich, C. P. N., Buecker, A., Wassenberg, S., Torres, G., Gurung, A., Hall, T., & Labyed, Y. (2024). *Ultrasound-Derived Fat Fraction for Hepatic Steatosis Assessment: Prospective Study of Agreement With MRI PDFF and Sources of Variability in a Heterogeneous Population.* https://doi.org/10.2214/ajr.23.30775

Postigliatti, A., Aris, R., Constenla, D., Ramírez, Ó., Schiappacasse Faúndes, G., Cea, J., González, F., Cavada, G., Ginesta, A., & Silva, C. (2024). *Detection of hepatic steatosis with ultrasound attenuation imaging compared to MRI*. https://doi.org/10.21203/rs.3.rs-5341525/v1

Jin, H., Jiao, M., Yu, C., Ren, T., Chen, Q., Dai, Z., Xie, E., Jiang, L., & Li, Y. (2025). *Discovering of ultrasound-derived fat fraction as a non-invasive and efficient identification for nonalcoholic fatty liver disease*. https://doi.org/10.21203/rs.3.rs-5663482/v1

Fu, Z., Grafakos, L., Lin, Y., Wu, Y., & Yang, S. (2023). Riesz transform associated with the fractional Fourier transform and applications in image edge detection. *Applied and Computational Harmonic Analysis*, *66*, 211-235. https://doi.org/10.1016/j.acha.2023.05.003

Sahaya Mercy A, Arockia Sahaya Sheela DG (2025). Region-Focused Patch-Based Mixup for Improving Classification in Low-Resource Medical Imaging. Indian Journal of Science and Technology 18(34): 2811-2820. https://doi.org/10.17485/IJST/v18i34.1313