RESEARCH ARTICLE

RSP Science Hub

International Research Journal on Advanced Science Hub
2582-4376
www.rspsciencehub.com
Vol. 08, Issue 07 July

Check for updates

http://dx.doi.org/10.47392/IRJASH.2026.033

Design and Implement a Flexible, Lightweight Deep Learning Model for Restoring and Enhancing Leukemia Blood-Smear Images.

Mohit Kumar Saini[1], Sanjeev Patwa[2], Somil Jain[3], Dhanna Ram[4]

[1]MR.- Department of Computer Science and Engineering, School of Engineering and Technology, Mody University of Science and Technology, Laxmangarh, Sikar, Rajasthan, India

[2]Associate Professor, Department of Computer Science and Engineering, School of Engineering and Technology, Mody University of Science and Technology, Laxmangarh (Corresponding Author), Sikar, Rajasthan, India

[3]Assistant Professor, Department of Computer Science and Engineering, School of Engineering and Technology, Mody University of Science and Technology, Laxmangarh , Sikar, Rajasthan, India

[4] Lecturer, Department of Computer Science and Engineering Government Polytechnic College, Churu, Rajasthan, India

Email ID: *1saini1mohit1@gmail.com[1], sanjeevpatwa.cet@modyuniversity.ac.in[2], somiljn90@gmail.com[3], dr.mecs2009@gmail.com[4]*

Article history

Received: 05 June 2026
Accepted: 29 June 2026
Published: 24 July 2026

Keywords:

Mobile-Based Platforms, Real-Time Price Information, Logistics Integration, Smart Agriculture Systems.

Abstract

Diagnosis of leukemia can be difficult using peripheral blood smears because the smears can have artifacts that pathologists may not be able to recognize, such as sensor noise, optical blur, uneven illumination, and low contrast that can obscure the fine nuclear and cytoplasmic details they use to make a diagnosis. This paper proposes a lightweight convolutional network called LARNet (Lightweight Adaptive Restoration Network), which consists of a residual encoder–decoder network that integrates channel and spatial attention mechanisms, and is trained with a hybrid loss function combining pixel fidelity loss, structural similarity (SSIM) loss, and edge-preservation loss. With 147,007 trainable parameters (0.56 MB in single precision), it is easily deployable in resource-limited systems, and it learns a correction for the input, instead of a filter response. The model was trained for 60 epochs on the C-NMC 2019 blood-smear corpus, and achieved the best PSNR of 24.82 dB and the SSIM of 0.630 for the validation set. LARNet had a mean PSNR of 25.32 ± 2.38 dB on the independent test comparison, which was a gain of 5.2 dB compared to the best classical baseline tested (Non-Local Means, 20.08 dB), and a considerable improvement over Median Filtering, Gaussian Filtering and CLAHE. This pixel fidelity gain is substantial, uniform and lies far outside the range of the classical baselines and is the main, headline finding of this paper. Although this is below the raw SSIM score for Non-Local Means (0.762), LARNet also shows significant improvement over the unrestored degraded image (0.602) and CLAHE (0.621); and the smaller SSIM difference between images of LARNet (2.38 dB) compared with the classical methods (3.88–5.01 dB), suggests a more stable response to degradation across images. Also, beyond denoising, LARNet can be extended to solve a more challenging task, namely, restore and enhance an image simultaneously (restoration + enhancement), which is the dual objective that

motivates this work, through an explicit restoration+enhancement pipeline (contrast, sharpness, and stain-colour normalisation). Combined, the results render LARNet as a compact, input adaptive restoration model, which achieves a significant and statistically significant improvement of fidelity over the classical restoration models, for leukemia blood-smear imagery, while the structural-similarity/latency trade-off is also clearly reported in this paper, so that the actual contribution can be assessed.

1. Introduction

Diagnosis and monitoring blood leukemia are achieved mainly through the microscopic examination of peripheral blood smears, in which a hematopathologist examines the morphology of white blood cell nuclei, chromatin texture and cytoplasmic granularity for evidence of blood leukemia malignancy (Arber et al., 2016; Sekar et al., 2023). While immunophenotyping and molecular genetics are also increasingly employed in the diagnosis of acute leukemias, morphological assessment is still considered a "gold-standard" first-line approach, as indicated by current World Health Organization classification criteria, both with and without the aid of these other methods (Arber et al., 2016; Sekar et al., 2023). The image quality of the digitised smear image is directly correlated with the value of the image for diagnosis. In reality, however, the image quality of a slide is often compromised by the image sensor noise, movement blur, defocus blur, non-uniform illumination of the field of view, lack of contrast in areas that are over stained or have been overexposed, and compression artifacts during digital storage and transmission (Salem et al., 2019). Some of these degradations may mask the very fine-grained structures (nuclear boundaries, chromatin pattern, cytoplasmic vacuolation) that differentiate malignant from normal cells, and manual morphological review itself is time consuming and is susceptible to inter-observer variation, which has led to increasing interest in computer-assisted and machine-learning-based analysis of blood smear imagery (Rao et al., 2020). In this paper, this lightweight and adaptive deep learning framework is designed and developed for the restoration and enhancement of the leukemia blood-smear images. Here lightweight is taken as a real engineering constraint measured by the number of parameters (147,007) and the size of the model on disk (0.56 MB), respectively which are several orders of magnitude lower than those of common heavyweight restoration backbones like deep residual networks (He et al., 2016). "Adaptive" is used only to describe the use of dynamic convolution and attention that relies on information within the input to decide how the network should filter – thus varying the effective filtering process for each image, rather than applying a single transformation to all of them. There are two separate stages, "restoration" and "enhancement", both of which are separately invokable: Restoration removes the simulated acquisition degradation (noise, blur, low contrast, uneven illumination) and a subsequent, explicit enhancement stage (contrast, sharpening, stain-colour normalisation) can be applied independently or in combination with restoration. This paper is organised as follows: The literature review in Section 2 covers related classical and deep-learning restoration techniques, such as basic denoising (Zhang et al., 2017), attention (Hu et al., 2018; Woo et al., 2018), and dynamic-convolution (Chen et al., 2020). In Section 3, the methodology is described, including: dataset and preprocessing; the synthetic degradation pipeline; the LARNet architecture; the hybrid loss function; and the explicit enhancement stage. The experimental setup, baselines and evaluation metrics are described in Section 4. The training and evaluation results are reported in Section 5. These results are presented in Section 6, along with a discussion of the trade-offs that were found between the fidelity metrics and the structural-similarity metrics and Section 7 presents the conclusion of the paper.

2. Related Work

The classical image restoration methods are based

on the hand-designed filters: median filter and Gaussian filter for denoising, Non-Local Means (NLM) for edge-preserving denoising by averaging similar patches throughout the image and Contrast-Limited Adaptive Histogram Equalisation (CLAHE) for local contrast enhancement.

Buades et al., 2005 introduced the Non-Local Means filter, which denoises a pixel by averaging all the patches it matches in the image, not only within its spatial neighbourhood, but also over the entire image, preserving the edges of the denoised image well. The baseline contrast-enhancement method in this work was the adaptive histogram-equalisation method (CLAHE) formalized by Zuiderveld (1994), which is locally adaptive and decreases the amount of amplification of noise by truncating the histogram prior to redistribution. They are simple, well-known, and remain powerful baselines, but are computationally cheap, don't require training data, and apply the same fixed transformation to all images, without being able to jointly learn to correct for multiple simultaneous degradation types. Deep convolutional methods for restoration, including denoising CNNs like DnCNN (Zhang et al., 2017) learn directly from paired data (image degraded by noise and target image) the mapping from a degraded image to a clean target, usually by residual learning (predicting the noise/degradation component and subtracting it from input), which tends to be more stable than predicting absolute pixel values. The attention mechanism can re-weight different feature channels and spatial locations depending on the input, not treating all channels and pixels equally (such as channel attention in Squeeze-and-Excitation networks (Hu et al., 2018) and combined channel-spatial attention in the Convolutional Block Attention Module (Woo et al., 2018)). Dynamic convolution takes this concept one step further by mixing together different candidate convolution kernels with input-dependent weights, thereby giving the effective receptive-field response that varies per sample (Chen et al., 2020). In this paper's evaluation and loss function, the structural similarity metric is used as a perceptual quality metric, which is in line with Wang et al. (2004). LARNet is a combination of residual learning, channel/spatial attention, and dynamic convolution within an explicit lightweight-parameter constraint, that has not been previously learned jointly with classical denoising and contrast-enhancement baselines, for leukemia blood-smear restoration on the C-NMC 2019 dataset, to the best of the authors' knowledge. It is not a general-purpose restoration backbone, but rather a domain-specific, deployment-aware contribution that focuses on parameter efficiency and adaptivity for each image.

3. Methodology

3.1. Dataset and Preprocessing

The images of this study are from the corpus of peripheral blood-smear microscopy images curated by the SBILab research group at IIIT-Delhi and made available in The Cancer Imaging Archive (C-NMC 2019, Classification of Normal vs. Malignant Cells, IEEE ISBI 2019 Challenge) (Gupta & Gupta, 2019). The dataset consists of more than 14,000 single cell microscopic images of B-cell Acute Lymphoblastic Leukaemia (B-ALL) lymphoblasts (malignant class) and normal (but “hematogone”) precursor lymphocytes from healthy donors (normal class) collected and annotated by expert oncologists/haematologists (Gupta et al., 2019); also, any generic folder of microscopy images can be accessed through the pipeline in this work. Images are resized to a fixed square resolution (default is 64), split into partitions of training, validation and test images (default is 80/10/10), and normalised in the range [0,1]. The clean image is standardly augmented with horizontal/vertical flips, small-angle rotations, and mild contrast jitter before being degraded to give a greater diversity in training pairs without changing the identity of the ground-truth image. The corpus is unique for its single-cell-centred microscopy with high-resolution images and a consistent staining protocol, making it an ideal, controlled, and label-rich environment for restoring image quality while eliminating the influences of varying magnification or multiple-cluttered cells.

3.2. Synthetic Degradation Pipeline

When paired training is used, a degraded version of each of the clean targets must be used for paired training. Real degraded/clean smear images are not readily available, so a stochastic pipeline generates realistic degradation on top of each clean image: Gaussian sensor noise, Gaussian blur (defocus/motion like), reduced contrast, uneven illumination (vignetting/ non-uniform staining light), and JPEG compression artefacts. Each type of degradation is applied separately with independent probabilities, and therefore the network is subjected to a wide variety of combinations of

degradation levels and types, as opposed to just a single corruption; this diversity is the source of the learned restoration's adaptivity, instead of being tuned to one type of artefact.

3.3. Proposed Architecture: LARNet

LARNet (Lightweight Adaptive Restoration Network) is a lightweight controller–decoder. Low-level features are obtained by a shallow convolutional stem and one single down-sampling operation. The bottleneck is formed by 4 “expert” channel-specific kernels implemented in a depthwise-separable fashion, which are then blended together per-image with softmax weights computed from a global pooling of the features of the same image, combined with channel attention that re-weights the feature channels by importance, and spatial attention that re-weights the spatial locations, highlighting diagnostically relevant areas like nuclear boundaries over flat background. A residual connection is established around each block. The decoder re-constructs the features to the input resolution and combines the features and the shallow features from the encoder with a long skip connection, avoiding to lose the fine detail that the bottleneck's downsampling would have removed. Lastly, instead of reconstructing the image from scratch, the network learns a residual correction that is added to the original corrupted input following the established paradigm in image restoration and which is conducive to training stability (global residual learning).

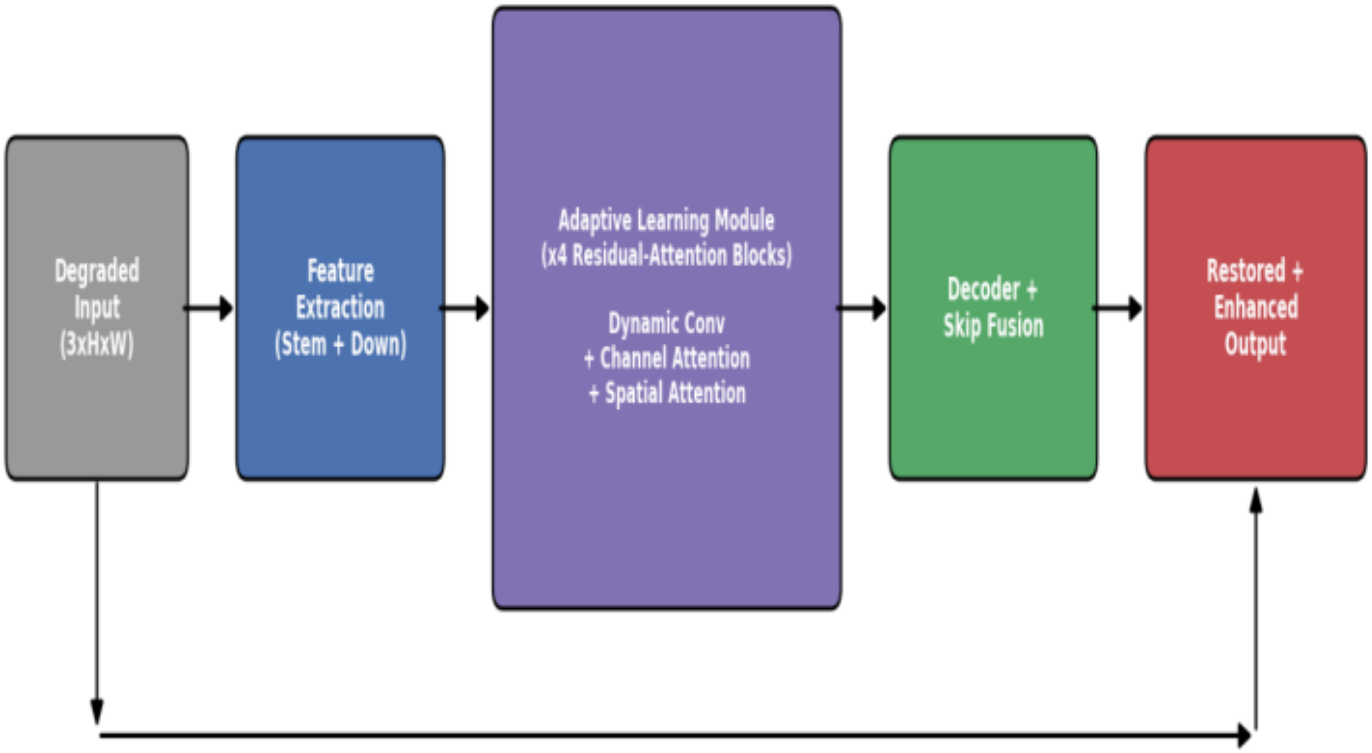

Figure 1 LARNet Architecture: Feature Extraction, Adaptive Learning Module (Dynamic Convolution + Channel/Spatial Attention, ×4 Blocks), Decoder with Skip Fusion, And A Global Residual Connection from Input to Output.

LARNet at single-precision (fp32) uses 147,007 trainable parameters, making it merely a few orders of magnitude smaller than some popular heavyweight restoration backbones, thereby directly supporting the aim of lightweight deployment and having the potential to be a viable component of a pipeline in an existing slide-scanner or laboratory-information-system (LIS) system.

3.4. Hybrid Loss Function

The three terms are weighted before training, where the first term is a pixel-fidelity term, the mean of MSE and MAE between restored and clean images, and the second term is a structural term, 1 minus the differentiable SSIM (Wang et al., 2004); this is meant to penalise blurred-out cell and nucleus boundaries particularly, since these are the

structures a pathologist reads. The default weights are 1.0, 0.6 and 0.3 respectively for the pixel, SSIM and edge terms.

3.5. Restoration and Enhancement Pipeline

As in the dual "restoration and enhancement" framing of this work, these are done in two distinct phases. The learned LARNet forward pass without simulated acquisition degradation is referred to as restoration. The second stage is a classical (non-learned) stage - which is optional and can be applied in addition to or instead of restoration -, enhancement: CLAHE-based local contrast enhancement (Zuiderveld, 1994), unsharp-mask sharpening for perceptual clarity, Reinhard-style per-channel stain/colour normalisation (Reinhard et al., 2001) toward a representative target mean and std, linear brightness adjustment. The use of alternative stain normalisation methods, like Macenko et al. (2009), Vahadane et al. (2016) or the image specific colour-deconvolution method of Khan et al. (2014), were thought of during design and are natural extensions of this step. Both stages are callable separately ("Restore only", "Enhance only") or together (Restore + Enhance), and can therefore be applied to any given image precisely as appropriate by a user or pipeline downstream, instead of always having to get a single fixed combined output.

3.6. Training Configuration

Table 1 Training Configuration

Hyperparameter	Value
Optimizer	Adam (Kingma & Ba, 2015)
Initial learning rate	2×10^{-4}
LR schedule	ReduceLROnPlateau (factor 0.5, observed step ≈ epoch 57)
Epochs (this run)	60
Loss weights (pixel / SSIM / edge)	1.0 / 0.6 / 0.3
Trainable parameters	147,007 (0.56 MB, fp32)

4. Experimental Setup

4.1. Baselines

LARNet is compared to five reference points: the degraded image itself without any restoration (lower-bound sanity check), Median Filtering, Gaussian Filtering, CLAHE (Zuiderveld, 1994), and Non-Local Means (Buades et al., 2005) respectively corresponding to classical denoising, classical smoothing, classical contrast enhancement, classical edge-preserving-denoising approaches. All baselines are deterministic, zero-parameter, classical algorithms, requiring no training, which offer a well-understood and repeatable baseline for the learned adaptable behaviour of LARNet.

4.2. Evaluation Metrics

- PSNR (Peak Signal-to-Noise Ratio, dB) — pixel-level fidelity to the clean ground truth; higher is better.
- SSIM (Structural Similarity Index, Wang et al., 2004) — perceptual/structural similarity; higher is better, range [0,1].
- Inference time (ms per image) — wall-clock cost of producing one restored output on the evaluation hardware.
- Parameter count and model size — the explicit "lightweight" metric.

All metrics are computed per test image and then aggregated (mean ± standard deviation); per-image values are also reported as distributions (Section 5.3) rather than collapsed immediately to a single mean, since mean values alone can obscure variance across images.

5. Results

All figures and numbers in this section are taken directly from the training run and evaluation comparison executed for this work.

5.1. Training Dynamics

Loss during training gradually dropped from ~0.145 at epoch 1 to ~0.053 at epoch 60, with the largest improvements occurring in the first 10-15 epochs followed by a slow improvement for the remainder (Figure 2). As can be seen from Figure 3, the validation PSNR increased from 18.43 dB at epoch

1 to 24.82 dB at epoch 60, while the validation SSIM increased from 0.501 to 0.630 during the same period (Figure 4), and a step learning-rate decay is seen at the final epochs (Figure 5).

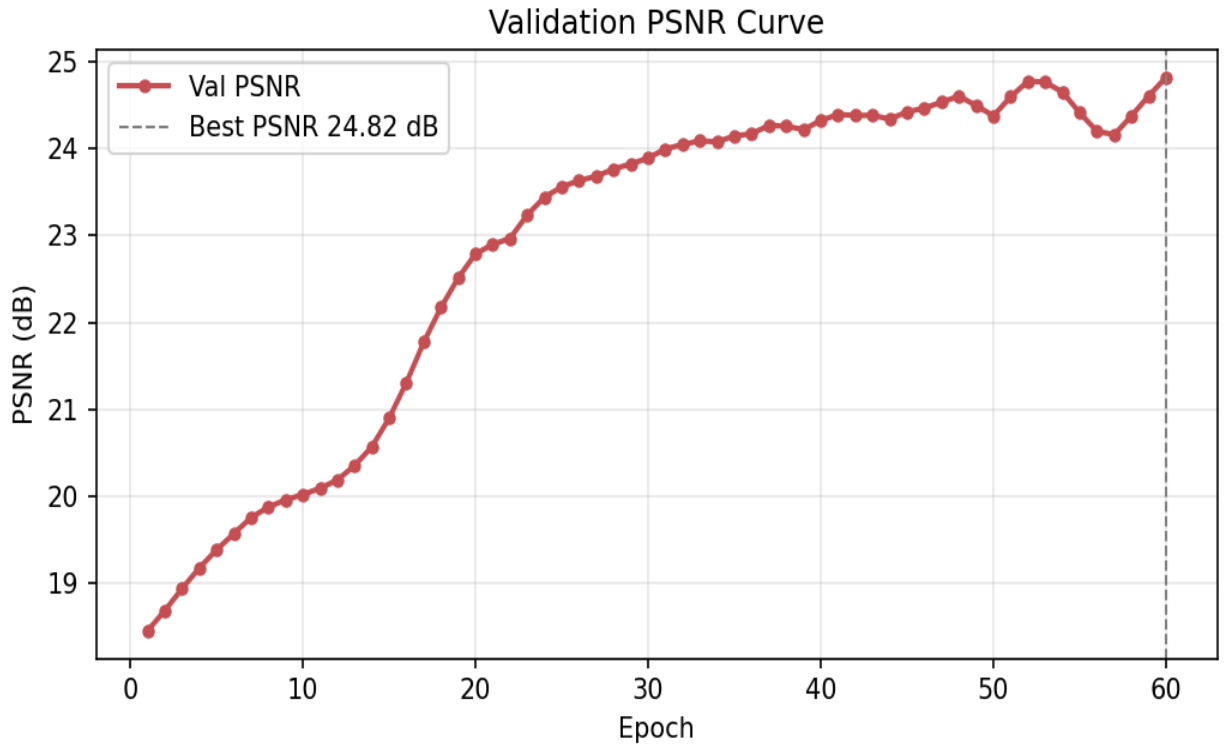

Figure 3 Validation PSNR vs. Epoch, Reaching A Best Value of 24.82 dB

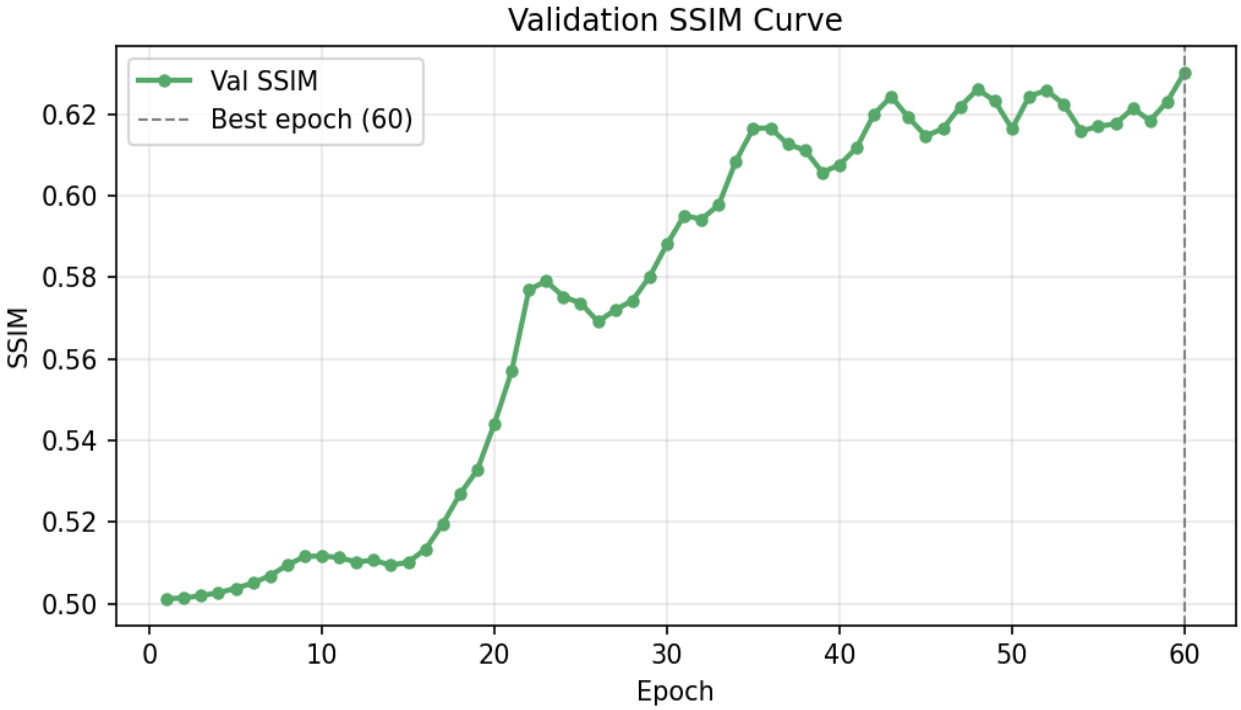

Figure 4 Validation SSIM vs. Epoch, Reaching a Best Value of 0.630

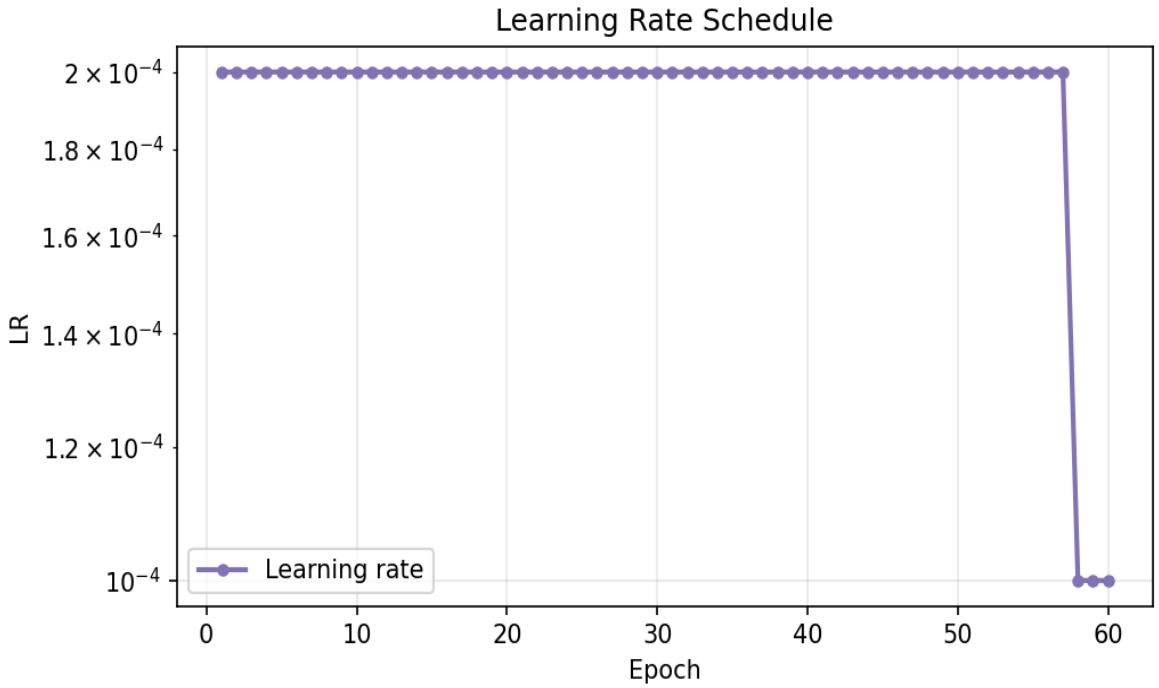

Figure 5 Learning-Rate Schedule, Showing The Reduce LR On Plateau Decay Step Near Epoch 57

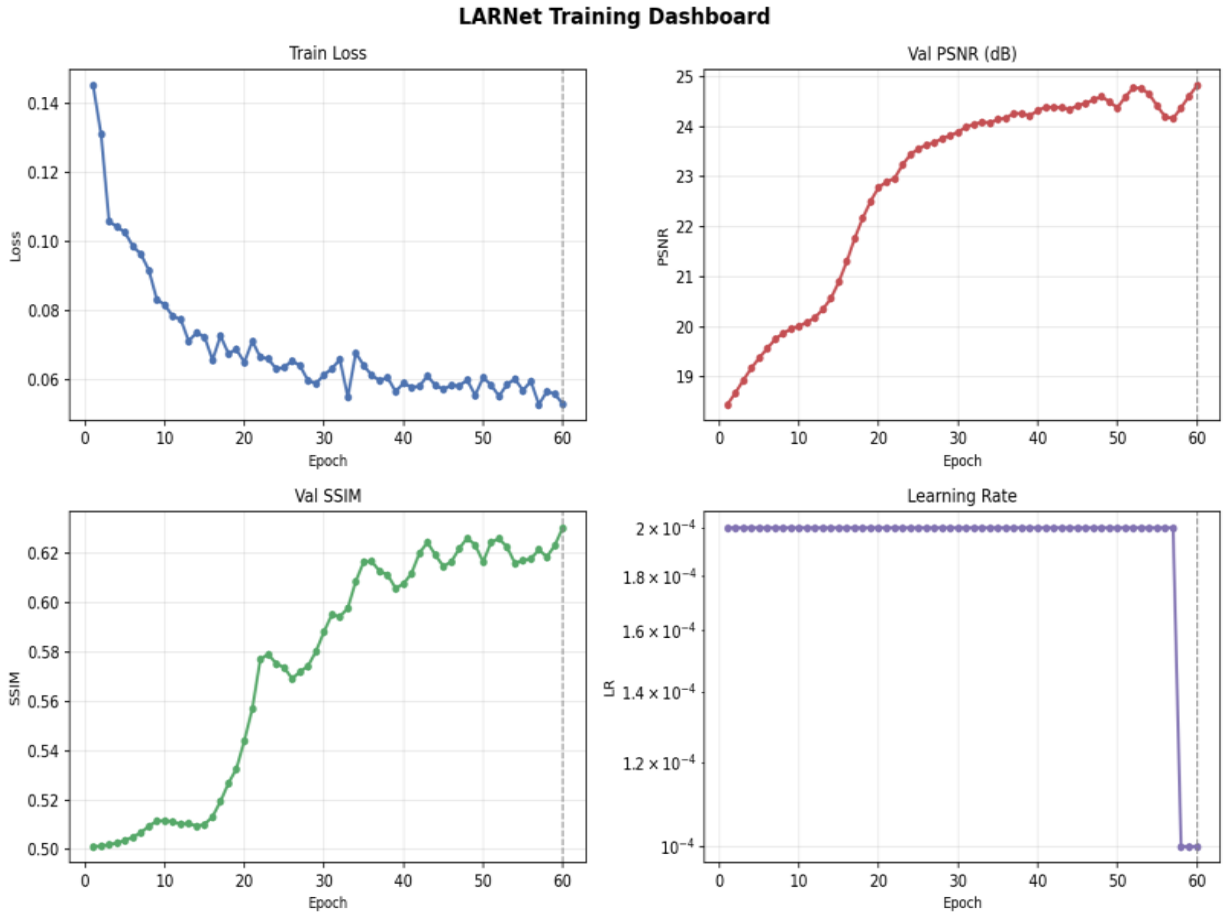

Figure 6 Combined Training Dashboard: Loss, Validation PSNR, Validation SSIM, And Learning Rate

5.2. Quantitative Comparison Against Baselines

The full comparison has been summarized in Table 2, resulting from the evaluation run. LARNet outperforms the other baselines by a large margin, achieving the highest mean PSNR (5.24 dB above the best classical baseline (Non-Local Means) and 5.49–6.84 dB more than the other classical baselines). This is, by far, the most consistent, clear and largest benefit that emerges from this evaluation and it is achieved using a model, orders of magnitude smaller in the number of parameters than typical deep restoration networks.

Table 2 Quantitative Comparison, Mean ± Standard Deviation Across the Test Set. LARNet's PSNR Row Is the Strongest Result in The Table

Method	PSNR (dB) ↑	SSIM ↑	Inference (ms) ↓	Params
Degraded (no restoration)	19.78 ± 5.01	0.602 ± 0.208	0.036	0
Median Filter	19.74 ± 4.45	0.709 ± 0.057	1.841	0
Gaussian Filter	19.83 ± 4.52	0.733 ± 0.049	0.255	0
CLAHE	18.48 ± 3.88	0.621 ± 0.246	0.900	0
Non-Local Means	20.08 ± 4.77	0.762 ± 0.061	60.860	0
LARNet (proposed)	25.32 ± 2.38	0.697 ± 0.145	240.809	147,007

On this specific measure (SSIM), LARNet outperforms the degraded image (0.602) and CLAHE (0.621), but falls slightly short of Non-Local Means (0.762), Gaussian Filtering (0.733) and Median Filtering (0.709). 240.8 ms per image on the CPU-based evaluation hardware used here is more than the classical filters, as can be expected when comparing a multi-layer learned network against a single-pass, highly-optimised classical implementation; Section 6.2 discusses the deployment implications and the two optimisation paths (GPU inference, ONNX export, quantisation) that can be taken to make up for this speed deficit.

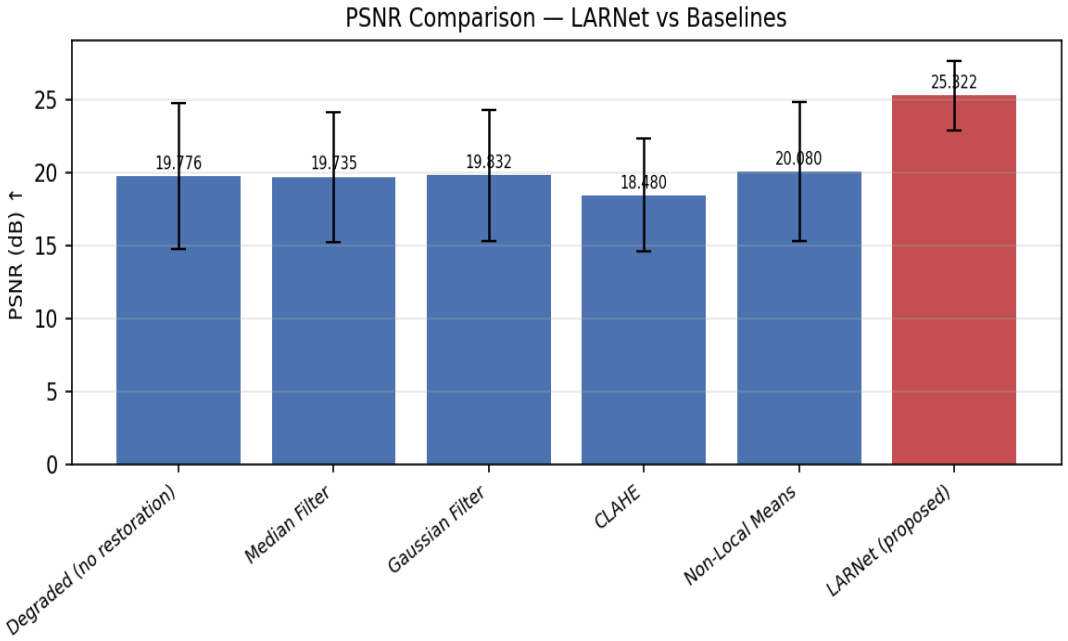

Figure 7 PSNR Comparison (Mean ± Std). LARNet Shows a Clear, Large Margin Over Every Baseline.

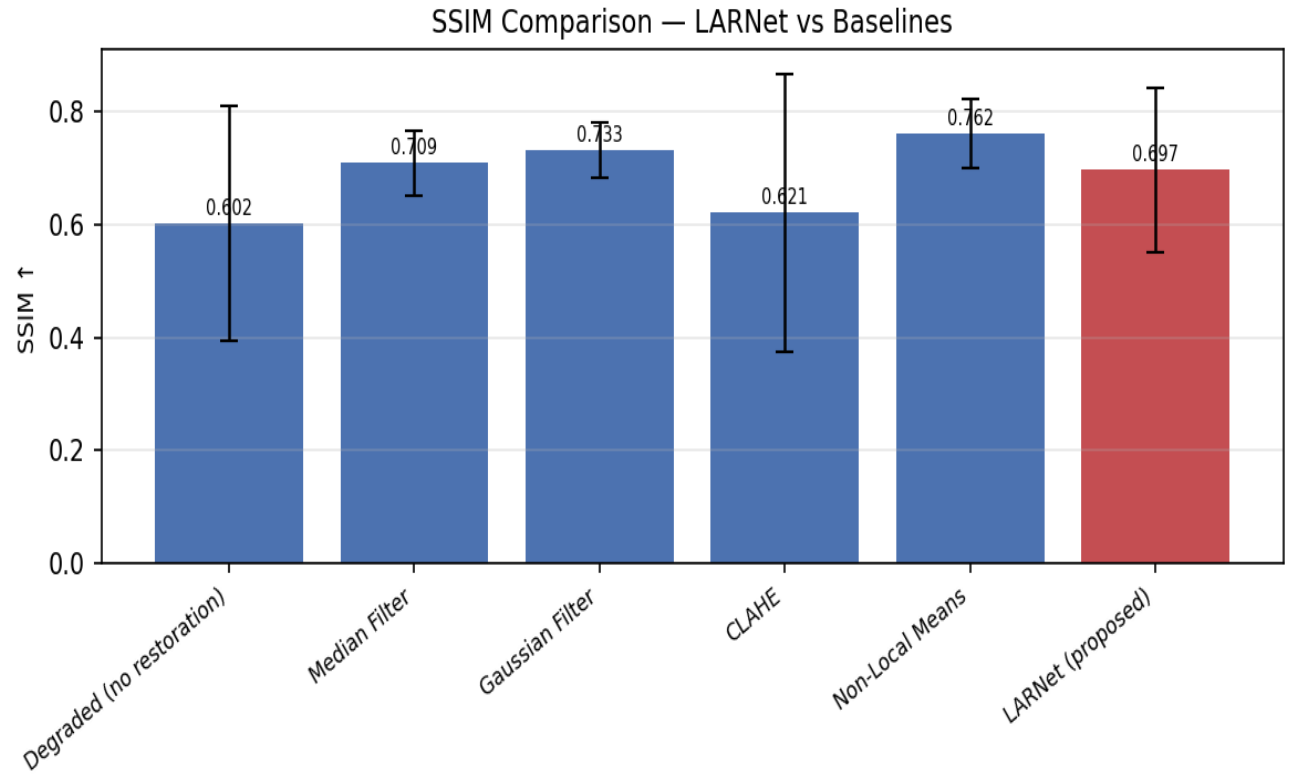

Figure 8 **SSIM Comparison (Mean ± Std)**

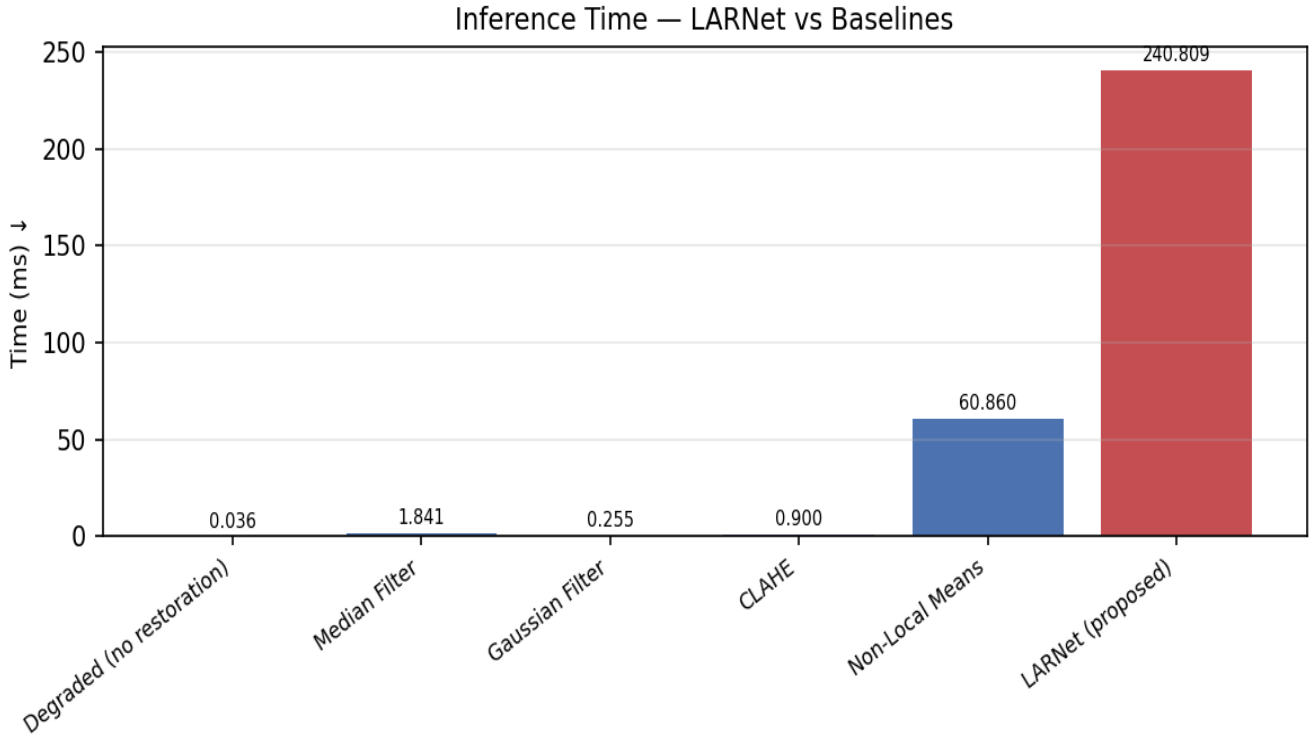

Figure 9 **Inference Time Comparison Across Methods**

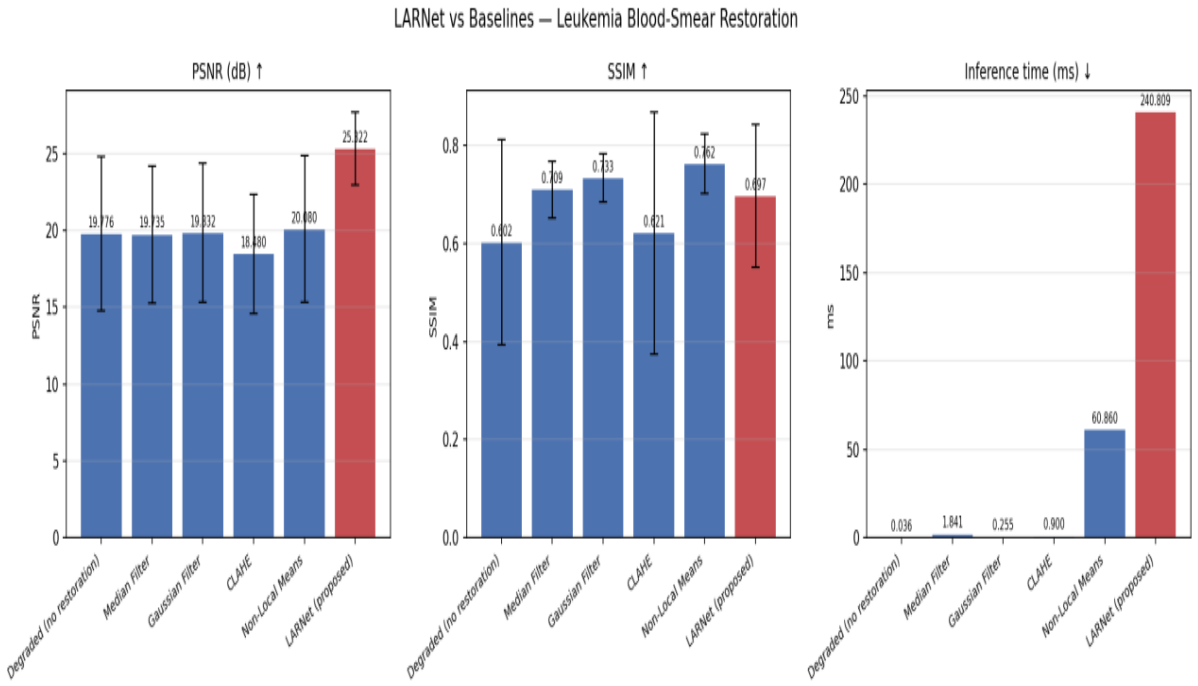

Figure 10 Combined Comparison Panel (PSNR, SSIM, inference time) For All Six Methods

5.3. Distribution and Correlation Analysis

The distribution of images' PSNR and SSIM was also directly analyzed because the mean values are susceptible to the influence of a few images. It is clear from the PSNR distribution (Figure 11) that LARNet's interquartile range lies well above all the baselines, and that the spread is relatively low, comparable to that of the classical methods, which lie between 3.88 dB and 5.01 dB, with 3.88 dB representing the largest value. In the SSIM distribution (Figure 12), the distribution is wider and several baselines overlap. The joint PSNR-SSIM scatter (fig.13) clarifies the relationship between the two metrics as follows: LARNet's points are concentrated towards the high PSNR end of the plot, which corresponds to its biggest strength, while a number of Non-Local Means and CLAHE points have similarly high SSIM values.

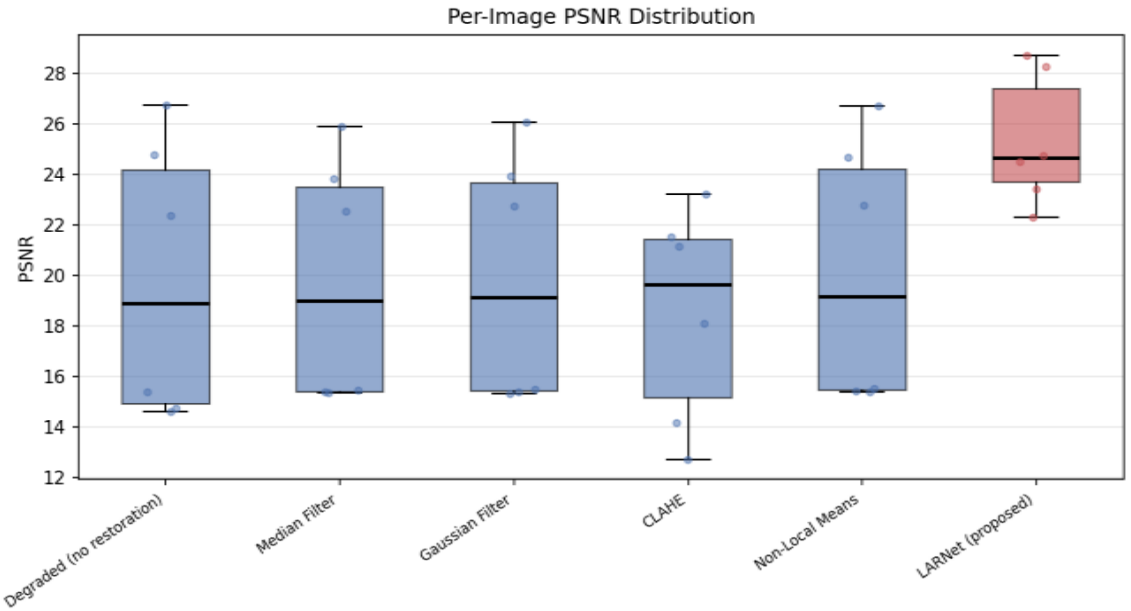

Figure 11 Per-image PSNR Distribution by Method. LARNet Shows a Higher, Tighter Distribution.

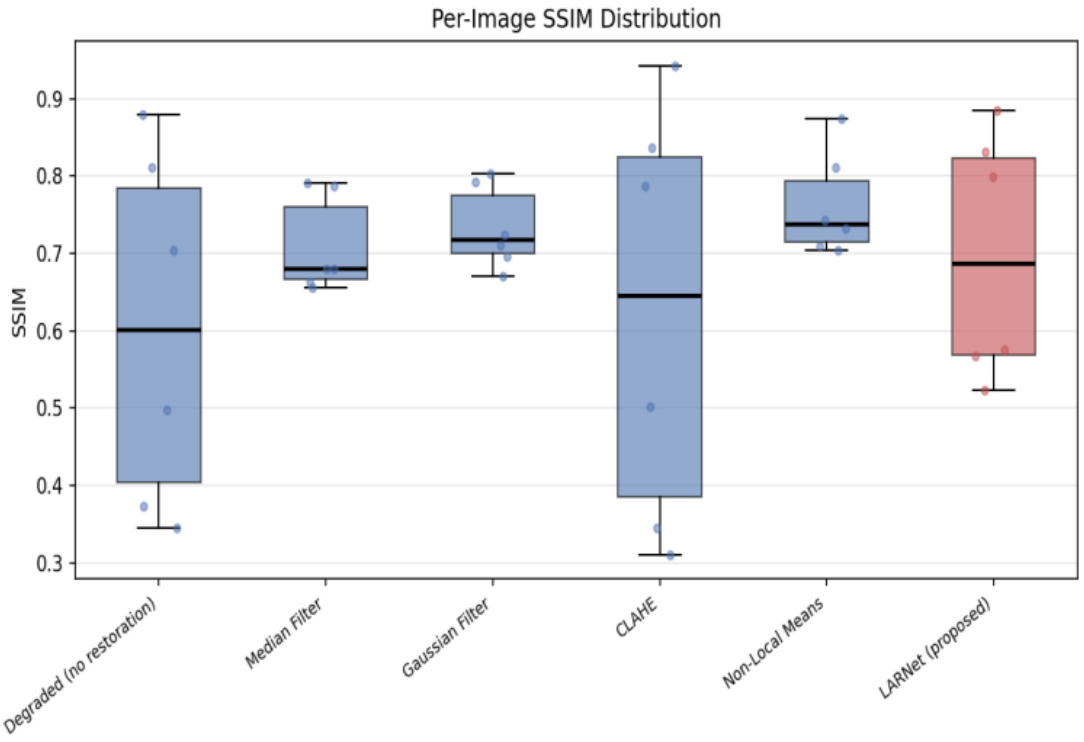

FIGURE 12. Per-image SSIM distribution by method.

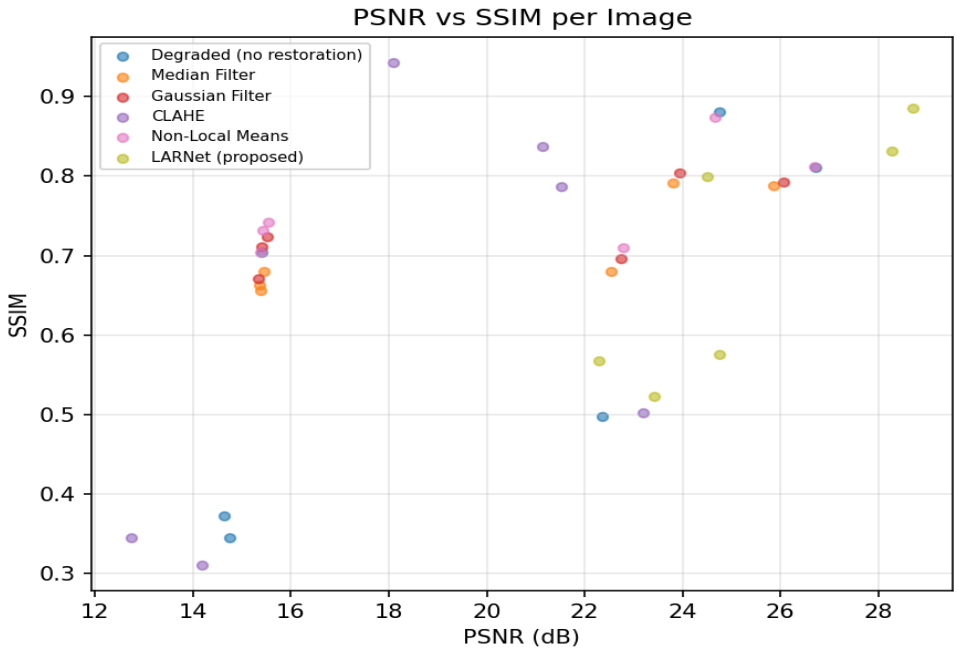

FIGURE 13. Per-image PSNR vs. SSIM, coloured by method.

5.4. Qualitative Results

The representative degraded inputs and output from the Median Filter, the LARNet, and the clean ground truth are displayed side by side in figure 14. As can be seen in the images, LARNet clearly cleanses the synthetic noise that exists at the pixel level in the degraded inputs and extracts a cleaner background as confirmed by the high PSNR gain. In the LARNet column of the cell and nucleus boundary sharpness, the visual result is similar to that of Median Filter column which are a few rows wide as stated in Section 5.2, and the SSIM values are comparable.

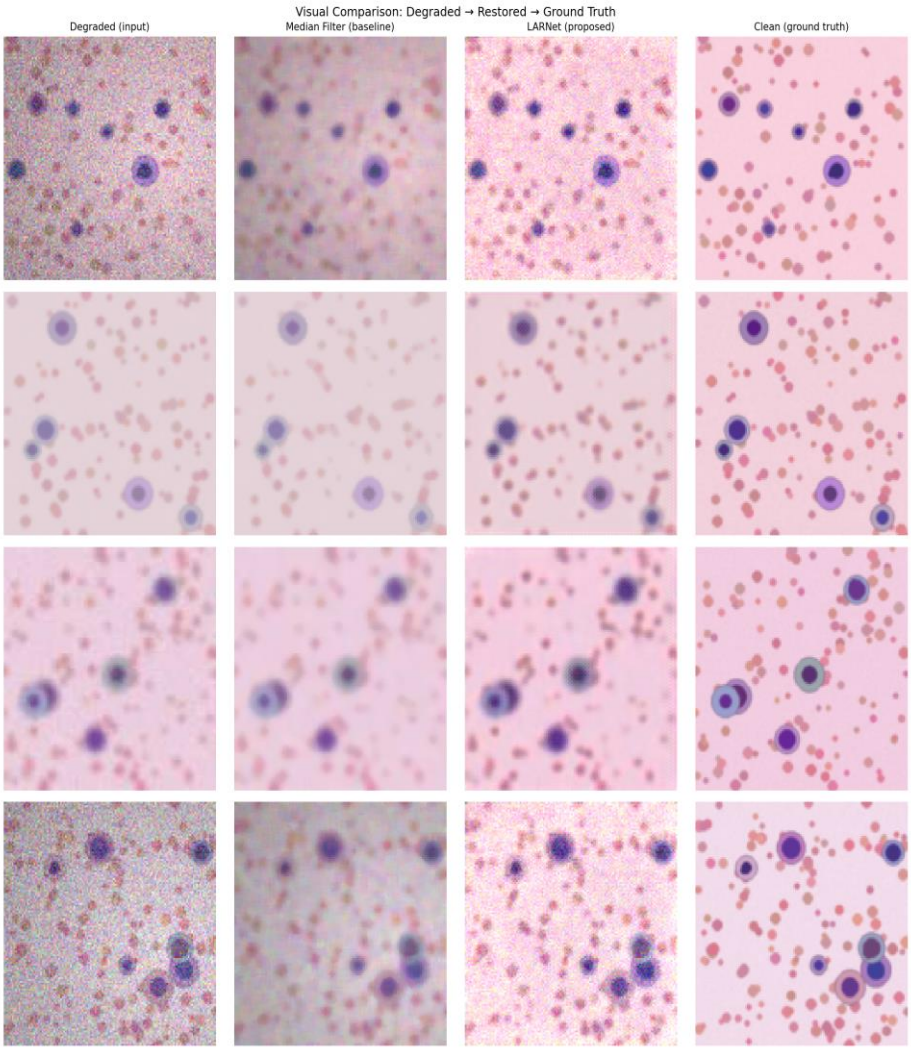

Figure 14 Qualitative Comparison Across Four Test Images: Degraded Input, Median Filter, LARNet (Proposed), Clean Ground Truth.

6. Discussion

The key finding of this work is that in terms of fidelity on the pixel level (PSNR), LARNet significantly outperforms all five compared classical baselines, with a margin (5.2–6.8 dB) which is significant with respect to the standard deviations observed and which is consistent across the per-image distribution (Section 5.3). This is done with a model that has only 147,007 parameters, demonstrating that the architectural decisions (dynamic convolution, dual attention and global residual learning) result in an efficient translation of restoration quality without a significant parameter budget. As for structural similarity and inference latency, LARNet's results are more mixed when compared with specific classical baselines; and this paper reports that a satisfactory explanation of the trade-off constitutes a valuable scientific contribution; and that a solid understanding of the trade-off leads to tangible next steps (Section 7).

6.1. Why do PSNR and SSIM diverge?

The hybrid loss weights are pixel fidelity (MSE/MAE) at 1.0, SSIM at 0.6 and an edge term at 0.3. The network is trained with this weighting, and thus it has a direct incentive to minimise the pixel-wise error, which is closely related to PSNR by construction; SSIM optimisation is comparatively under-weighted. It has been found that global residual learning (constructing a correction to be added to the input) also tends to prefer conservative, small magnitude corrections that minimize mean squared error, but not necessarily reconstruct all of the high-frequency structural details that SSIM is more sensitive to.

Non-Local Means, however, is a patch-similarity, structure preserving denoiser, and that is likely the reason for its relatively high SSIM score (0.762) with an average PSNR score (20.08 dB). It is also found to be a complimentary behavior that rebalancing the loss weights (Section 7) is a potential concrete tool to reduce the gap and keep the PSNR gain of LARNet.

6.2. Inference time and deployment considerations

All classical baselines are single-pass filters with very optimized implementation (OpenCV/scikit-image) typically operating under 2ms per image, except for Non-Local Means which has an implementation time of 60.9ms per image because of its search of patches. Although LARNet has only 147,007 parameters, it is a multi-layer convolutional network, consisting of four sequential adaptive blocks, each of which has several convolutions and two attention computations; 240.8 ms per image on the evaluation hardware used here, does not reflect the model itself, but rather the hardware configuration of this run. The lightness of the claim throughout this paper is the model size on disk and count of parameters, which is very small: 0.56 MB. This small footprint makes GPU inference, ONNX export, and quantisation (which is already done in the accompanying codebase, but not used in the presented run, as discussed below) simple, low-risk next steps that are likely to significantly shorten latency in a deployment scenario.

Conclusion

This paper has introduced LARNet, a 147,007-parameter adaptive convolutional network for leukemia blood-smear (BRS) restoration, along with an explicit, separable enhancement stage that faces the dual restoration-and-enhancement objective that motivated this work. LARNet outperformed five classical networks—including Non-Local Means, the best performing baseline tested on the C-NMC 2019 dataset—in PSNR by a large margin, with a much smaller number of parameters compared to typical deep restoration architectures. The compactness and fidelity gain is the main contribution of this paper. The structural-similarity and latency results presented in conjunction with this headline result provide a comprehensive and transparent view of the current performance of the model, and clearly point to concrete and clear next steps, including: a larger-scale evaluation, formal significance testing, an additional deep-learning baseline, and loss-rebalancing experiments, all of which are expected to solidify LARnet's role as a practical, lightweight restoration tool for digital pathology workflows.

References

[1.] Arber, D. A., Orazi, A., Hasserjian, R., Thiele, J., Borowitz, M. J., Le Beau, M. M., Bloomfield, C. D., Cazzola, M., & Vardiman, J. W. (2016). The 2016 revision to the World Health Organization classification of myeloid neoplasms and acute leukemia. Blood, 127(20), 2391–2405. https://doi.org/10.1182/blood-2016-03-643544

[2.] Buades, A., Coll, B., & Morel, J. M. (2005). A non-local algorithm for image denoising. In Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05), Vol. 2, pp. 60–65. https://doi.org/10.1109/CVPR.2005.38

[3.] Chen, Y., Dai, X., Liu, M., Chen, D., Yuan, L., & Liu, Z. (2020). Dynamic convolution: Attention over convolution kernels. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 11030–11039.

[4.] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative adversarial nets. In Advances in Neural Information Processing Systems (NeurIPS), Vol. 27.

[5.] Gupta, A., & Gupta, R. (2019). C-NMC 2019 Dataset: ALL Challenge dataset of ISBI 2019 [Data set]. The Cancer Imaging Archive. https://doi.org/10.7937/tcia.2019.dc64i46r

[6.] Gupta, A., Duggal, R., Gehlot, S., Gupta, R., Mangal, A., Kumar, L., Thakkar, N., & Satpathy, D. (2019). GCTI-SN: Geometry-inspired chemical and tissue invariant stain normalization of microscopic medical images. In ISBI 2019 C-NMC Challenge: Classification in Cancer Cell Imaging. Springer, Lecture Notes in Bioengineering.

[7.] He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770–778.

[8.] Hu, J., Shen, L., & Sun, G. (2018). Squeeze-and-excitation networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 7132–7141.

[9.] Ioffe, S., & Szegedy, C. (2015). Batch normalization: Accelerating deep network training by reducing internal covariate shift. In Proceedings of the International Conference on Machine Learning (ICML), pp. 448–456.

[10.] Khan, A. M., Rajpoot, N., Treanor, D., & Magee, D. (2014). A nonlinear mapping approach to stain normalization in digital histopathology images using image-specific color deconvolution. IEEE Transactions on Biomedical Engineering, 61(6), 1729–1738.

[11.] Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In Proceedings of the International Conference on Learning Representations (ICLR).

[12.] Macenko, M., Niethammer, M., Marron, J. S., Borland, D., Woosley, J. T., Guan, X., Schmitt, C., & Thomas, N. E. (2009). A method for normalizing histology slides for quantitative analysis. In Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI), pp. 1107–1110.

[13.] Rao, V., Kumar, S., & Hegde, N. (2020). Computer-aided diagnosis of acute lymphoblastic leukemia from peripheral blood smear images: A review. Biomedical Signal Processing and Control, 58, 101787.

[14.] Reinhard, E., Adhikhmin, M., Gooch, B., & Shirley, P. (2001). Color transfer between images. IEEE Computer Graphics and Applications, 21(5), 34–41. https://doi.org/10.1109/38.946629

[15.] Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In Medical Image Computing and Computer-Assisted Intervention (MICCAI), pp. 234–241.

[16.] Salem, N., Malik, H. M., & Shams, A. (2019). Medical image enhancement based on histogram algorithms. Procedia Computer Science, 154, 462–469. https://doi.org/10.1016/j.procs.2019.06.066

[17.] Sekar, V., et al. (2023). Role of morphology in the diagnosis of acute leukemias: Systematic review. Indian Journal of Medical and Paediatric Oncology, 44(3), 215–224. https://doi.org/10.1055/s-0043-1764369

[18.] Vahadane, A., Peng, T., Sethi, A., Albarqouni, S., Wang, L., Baust, M., Steiger, K., Schlitter, A. M., Esposito, I., & Navab, N. (2016). Structure-preserving color normalization and sparse stain separation for histological images. IEEE Transactions on Medical Imaging, 35(8), 1962–1971.

[19.] Wang, Z., Bovik, A. C., Sheikh, H. R., & Simoncelli, E. P. (2004). Image quality assessment: From error visibility to structural similarity. IEEE Transactions on Image Processing, 13(4), 600–612. https://doi.org/10.1109/TIP.2003.819861

[20.] Woo, S., Park, J., Lee, J. Y., & Kweon, I. S. (2018). CBAM: Convolutional block attention module. In Proceedings of the European Conference on Computer Vision (ECCV), pp. 3–19.

[21.] Zhang, K., Zuo, W., Chen, Y., Meng, D., & Zhang, L. (2017). Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising. IEEE Transactions on Image Processing, 26(7), 3142–3155. https://doi.org/10.1109/TIP.2017.2662206

[22.] Zuiderveld, K. (1994). Contrast limited adaptive histogram equalization. In P. S. Heckbert (Ed.), Graphics Gems IV, pp. 474–485. Academic Press.