Denoising and Restoration of Images using Artificial Intelligence

Nojus Vakauskas
Electrical and Electronic Engineering
Project Engineer at Honeywell
Glasgow, United Kingdom
nojus03@icloud.com

Omar Faramawy
Electrical and Computer Engineering
Purdue Univeristy
West Lafayette, Indiana
elsayedo@purdue.edu

Ahmed Solyman
Department of Engineering
Glasgow Caledonian University
Glasgow, United Kingdom
ahmed.solyman@gcu.ac.uk

Prof. Ir. Dr. Zakaria Che Muda
Engineering and Quantity Surveying
INTI International University (INTI-IU)
Nilai, Malaysia
zakaria.chemuda@newinti.edu.my

Mohamed Hassan
Electrical and Electronics Department
Egyptian Russian University in Cairo
Cairo, Egypt
mohamed-nabil@eru.edu.eg

Ahmed Samir
Electronics and Communication Engineering
Uruk University
Baghdad, Iraq
ahmed.s.ahmed@ururk.edu.iq

***Abstract*—This paper presents a comprehensive approach to image denoising and color restoration using artificial intelligence techniques. We investigate how convolutional neural networks (CNNs) can be leveraged to enhance image quality by removing noise artifacts and restoring color to grayscale images. Our methodology employs a dual-model approach: one model dedicated to noise reduction and another focused on color restoration. Both models utilize residual connections and specialized normalization techniques to preserve image details while performing their respective tasks. Experiments conducted on the CIFAR-10 and ImageNet datasets demonstrate the effectiveness of our approach, with quantitative evaluations using Peak Signal-to-Noise Ratio (PSNR) and qualitative visual assessments. The denoising model successfully removes Gaussian and salt-and-pepper noise while preserving essential image features, achieving a PSNR of up to 22.78 dB. The color restoration model transforms grayscale images into plausible color representations, with results improving significantly over extended training periods. This research contributes to the field of image processing by providing insights into neural network architectures optimized for image enhancement tasks and demonstrating their practical applications in restoring degraded visual content.This research contributes to UN Sustainable Development Goal 9 (Industry, Innovation and Infrastructure) by advancing AI-based techniques for restoring degraded images, thereby enhancing the reliability and efficiency of digital imaging systems in modern industrial and technological applications.**

***Index Terms*—Image denoising, color restoration, image colorization, convolutional neural networks (CNNs), residual learning, normalization, Gaussian noise, salt-and-pepper noise, CIFAR-10, ImageNet, peak signal-to-noise ratio (PSNR), image enhancement, SDG 9.**

I. Introduction

Image processing and restoration have become increasingly important in various fields, from medical imaging to digital photography. The degradation of images through noise, compression artifacts, or loss of color information presents significant challenges in preserving visual data integrity. This paper explores the application of artificial intelligence techniques, specifically convolutional neural networks (CNNs), to address two fundamental image enhancement problems: denoising and color restoration.

Images, at their most basic level, are representations of voltage signals captured by photodiodes and arranged in matrices that can be interpreted by the human visual system. In digital form, these images are defined as arrays where dimensions correspond to the height and width in pixels. Each pixel is represented by a two-dimensional function $f(x, y)$, where x and y are the coordinates, and f is the intensity at those coordinates [1]. Higher values of f produce lighter pixels, while lower values result in darker appearances.

The quality of digital images can be compromised by various types of noise introduced during acquisition, transmission, or storage. Noise manifests as random variations in brightness or color that do not reflect the true image content. Common types include Gaussian noise, which follows a normal distribution, and salt-and-pepper noise, characterized by random white and black pixels scattered throughout the image [2], [3]. These distortions not only affect visual quality but can also impede subsequent image analysis tasks.

Similarly, the absence of color information in grayscale images limits their expressive capacity and utility in applications where color provides essential contextual or semantic information. The restoration of color to grayscale images presents a complex challenge, as it requires inferring plausible color distributions based on luminance patterns and contextual cues [4].

This research aims to develop effective solutions for both image denoising and color restoration using deep learning approaches. We implement and evaluate neural network architectures designed specifically for these tasks, focusing on their ability to preserve image details while removing noise or adding color information. Our methodology leverages recent advances in CNN design, including residual connections and specialized normalization techniques [5], [6], to enhance performance and robustness.

The paper is organized as follows: First, we provide background information on image processing fundamentals and relevant artificial intelligence concepts. Next, we detail our methodology, including dataset selection, preprocessing techniques, and neural network architectures. We then present experimental results and evaluate the performance of our models using both quantitative metrics and qualitative assessments. Finally, we discuss the implications of our findings and potential directions for future research in this domain.

II. Methods

A. Dataset Selection

The selection of appropriate datasets is crucial for training and evaluating image processing models. For this research, we utilized two primary datasets: CIFAR-10 [7] and the downsampled ImageNet (64×64) [8], [9]. CIFAR-10 provides 60,000 color images across 10 classes, each sized at 32×32 pixels. While this dataset offers a good starting point for model development, its limited resolution restricts the visual clarity of results. The downsampled ImageNet dataset, containing approximately 14 million images resized to 64×64 pixels, offers greater diversity and higher resolution, enabling more detailed and legible outputs. This dataset's extensive variety makes it particularly suitable for training robust models capable of generalizing across different image types.

B. Image Processing Framework

Our implementation utilizes Python with TensorFlow and Keras for neural network development, chosen for their comprehensive libraries and efficient handling of image data. The preprocessing pipeline includes several key components:

1. **Normalization**: Images are normalized to the [0,1] range to optimize compatibility with activation functions such as ReLU, sigmoid, and tanh.
2. **Color Space Conversion**: For color restoration tasks, we convert between RGB and LAB color spaces [10]. The LAB color space separates luminance (L channel) from chrominance (A and B channels), allowing our model to focus on color information independently from brightness.
3. **Noise Addition**: For denoising experiments, we apply controlled amounts of Gaussian noise [11] and salt-and-pepper noise to clean images. The noise function is defined as:

$$I_{noisy} = I_{clean} + \mathcal{N}(0, \sigma^2) \tag{1}$$

where I_{noisy} is the resulting noisy image, I_{clean} is the original clean image, and $\mathcal{N}(0, \sigma^2)$ represents Gaussian noise with mean 0 and variance σ^2.

C. Neural Network Architecture

Our neural network architecture incorporates several advanced components designed specifically for image processing tasks:

1) Residual Blocks: Residual connections are implemented to address the vanishing gradient problem and preserve low-level details during training. The residual block is defined as [5], [12]:

$$y = F(x, W_i) + x \tag{2}$$

where x is the input to the residual block, $F(x, W_i)$ represents the residual mapping to be learned, and y is the output. This structure allows information to flow directly through the network via skip connections, facilitating the training of deeper networks.

2) Normalization Techniques: We employ batch normalization to stabilize and accelerate training by normalizing layer inputs across each mini-batch:

$$\hat{x}_i = \frac{x_i - \mu_B}{\sqrt{\sigma_B^2 + \epsilon}} \tag{3}$$

$$y_i = \gamma \hat{x}_i + \beta \tag{4}$$

where μ_B and σ_B^2 are the mean and variance of the mini-batch, ϵ is a small constant for numerical stability, and γ and β are learnable parameters.

For scenarios with memory constraints necessitating smaller batch sizes, we implemented group normalization [13] as an alternative, which normalizes across groups of channels rather than across the batch dimension.

3) Network Structure: The denoising network begins with a 7×7 convolutional layer to extract initial features, followed by multiple residual blocks that progressively increase the number of filters (64, 128, 256) to capture increasingly complex features. Max pooling layers reduce spatial dimensions while preserving essential information. The final layers reconstruct the denoised image through upsampling operations.

The color restoration network follows a similar structure but includes additional components to handle the transformation from grayscale to color. It takes a grayscale image (L channel) as input and predicts the corresponding A and B color channels. The output layer uses a hyperbolic tangent activation function to constrain values to the [-1,1] range, which are later denormalized to the appropriate color space ranges.

D. Training Methodology

Both models were trained using an adaptive learning rate strategy with the Adam optimizer. The learning rate was initially set at 1e-3 and gradually decreased using an exponential decay schedule [15]:

$$\alpha_t = \alpha_0 \cdot e^{-\beta t} \tag{5}$$

where α_t is the learning rate at step t, α_0 is the initial learning rate, and β is the decay rate.

For the denoising model, we used Mean Squared Error (MSE) as the loss function:

$$L_{MSE} = \frac{1}{n} \sum_{i=1}^{n} (y_i - \hat{y}_i)^2 \tag{6}$$

where y_i represents the clean target image and $\hat{y}_i$ is the model's prediction.

For the color restoration model, we implemented a custom loss function that combines Mean Absolute Error (MAE), Mean Squared Error (MSE), and an edge-based perceptual component to better preserve color fidelity and structural details. The MAE component is defined as:

$$L_{MAE} = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i| \tag{7}$$

The edge-based component L_{edge} was calculated by applying Sobel filters [16], [17] to both the ground truth y_i and predicted $\hat{y}_i$ A and B color channels to extract edge maps. The squared difference between these edge maps then constitutes the edge loss. The total loss for the color restoration model is a weighted sum of these components:

$$L_{total} = w_1 \cdot L_{MAE} + w_2 \cdot L_{MSE} + w_3 \cdot L_{edge} \tag{8}$$

where w_1, w_2 and w_3 are weighting coefficients (e.g., 0.5 for MAE, 0.3 for MSE, and 0.2 for edge loss, as explored in developmental stages) that balance the contribution of each loss term.

To optimize computational efficiency, we implemented mixed-precision training, which uses lower precision formats (FP16) where possible while maintaining model accuracy through careful scaling of gradients.

III. Results

A. Denoising Performance

The denoising models were evaluated across multiple training epochs to assess their effectiveness in removing noise while preserving image details.

1) CIFAR-10 Denoising Model Results: For the CIFAR-10 dataset [7], we applied Gaussian noise with a factor of 0.125 (12.5% of the image) and salt-and-pepper noise with a probability of 0.01 (1%). The model's performance was measured using both Mean Squared Error (MSE) loss and Peak Signal-to-Noise Ratio (PSNR), with higher PSNR values indicating better image quality.

Table 1 summarizes the quantitative results for the CIFAR-10 denoising model:

TABLE I
QUANTITATIVE RESULTS FOR THE CIFAR-10 DENOISING MODEL.

Epochs	Loss	PSNR (dB)
5	0.00594	22.26
10	0.00536	22.71
15	0.00534	22.72
20	0.00527	22.78

The results demonstrate a consistent improvement in both loss and PSNR as training progresses, though the rate of improvement diminishes after 10 epochs. This suggests that the model approaches its optimal performance relatively early in the training process.

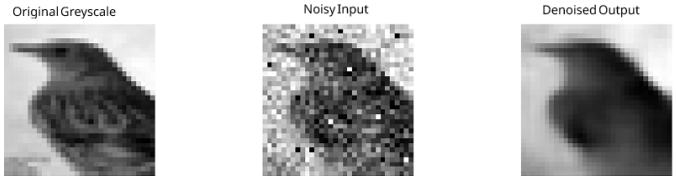

Fig. 1. CIFAR-10 denoising model, 5 epochs.

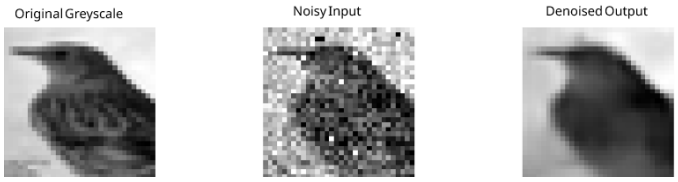

Fig. 2. CIFAR-10 denoising model, 10 epochs.

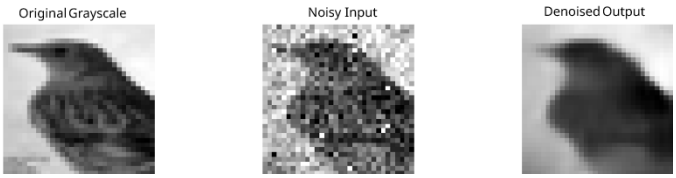

Fig. 3. CIFAR-10 denoising model, 20 epochs.

Qualitative assessment of the CIFAR-10 denoising results reveals that the model initially produces blurry versions of the original images after 5 epochs (Fig. 1), effectively removing noise but at the cost of detail preservation. This behavior is consistent with the application of a Gaussian filter, which is known to reduce noise while introducing blur. As training progresses to 10 (Fig. 2) and 20 epochs (Fig. 3), the model learns to preserve image details better while still effectively removing noise, resulting in clearer denoised images that more closely resemble the original clean versions.

2) ImageNet Denoising Model Results: For the ImageNet denoising model, we maintained the same Gaussian noise factor but increased the salt-and-pepper noise probability to 0.02 (2%) to account for the larger image size (64×64 pixels). The quantitative results are presented in Table 2.

TABLE II
QUANTITATIVE RESULTS FOR THE IMAGENET DENOISING MODEL.

Epochs	Loss	PSNR (dB)
5	0.0030	25.63
10	0.0024	26.67
15	0.0022	26.88
20	0.0022	26.98

Similar to the CIFAR-10 model, performance improvements on the ImageNet dataset [8] became less significant after extended training. However, qualitatively, the denoised images from the ImageNet model, even at 5 epochs, retained more fine details compared to CIFAR-10, likely due to the higher initial resolution. As training progressed (e.g., 10, 15, and 20

epochs), the ImageNet model produced progressively clearer and more detailed denoised outputs, effectively balancing noise removal with feature preservation.

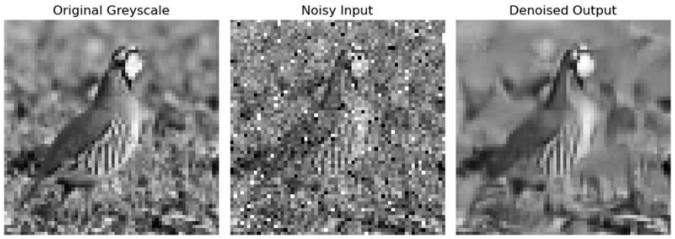

Fig. 4. ImageNet denoising model, 20 epochs.

B. Color Restoration Results

The color restoration model was evaluated on its ability to predict plausible color information (A and B channels in the LAB color space) from grayscale images (L channel). Unlike the denoising task, color restoration is inherently ambiguous, as multiple color solutions may be valid for a given grayscale input. Therefore, our evaluation focused on both quantitative metrics and qualitative visual assessment. Quantitative results for both CIFAR-10 and ImageNet datasets are summarized below.

TABLE III
QUANTITATIVE RESULTS FOR THE CIFAR-10 COLOR RESTORATION MODEL.

Epochs	Loss	PSNR (dB)
10	0.0690	24.44
25	0.0658	24.70
50	0.0571	25.02

TABLE IV
QUANTITATIVE RESULTS FOR THE IMAGENET COLOR RESTORATION MODEL.

Epochs	Loss	PSNR (dB)
10	0.0542	23.12
25	0.0528	24.02
50	0.0512	24.86
100	0.0483	26.84

Training the color restoration model, particularly on the ImageNet dataset, required significantly more epochs compared to the denoising model, with noticeable improvements continuing through 100 epochs. The early stages of training (around 5-10 epochs) for the ImageNet model produced images with a distinct sepia-like tint, indicating that the model had learned a conservative color mapping strategy that applied similar colors across different image regions.

For the ImageNet model, after initial troubleshooting of data handling and model overfitting, further training showed a clearer progression. After approximately 15 epochs, the model began to differentiate colors more effectively, with

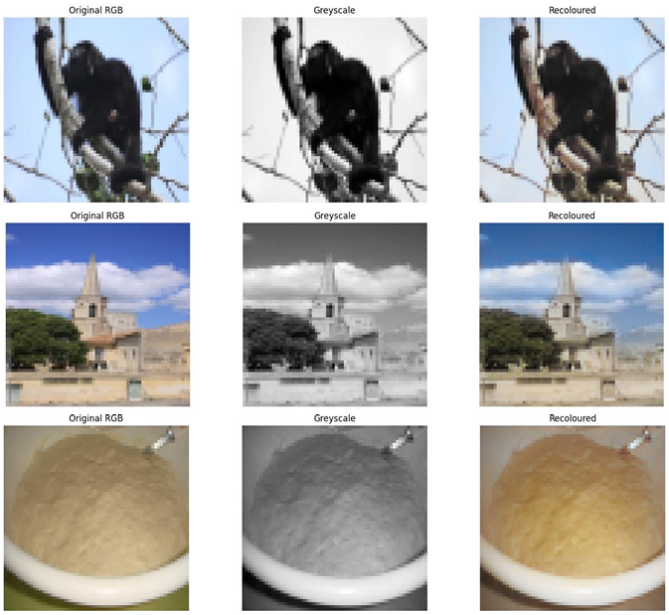

Fig. 5. ImageNet Color Restoration Model, 100 epochs.

reduced sepia effects and increased color variation. By 30-50 epochs, the restored colors became more accurate but sometimes remained somewhat dull compared to the originals. The final model, trained for 100 epochs and fine-tuned, produced significantly improved results with more vivid colors that closely matched the original images in many cases. Interestingly, the model occasionally generated plausible but different color schemes compared to the original images, particularly for backgrounds or objects with ambiguous color properties. This demonstrates the model's ability to learn general color distributions associated with different object classes and image contexts, rather than simply memorizing specific color mappings [18], [19].

The quantitative evaluation of the color restoration models showed steady improvement in loss values throughout training, though these metrics do not fully capture the perceptual quality of the colorization results. Visual inspection remains an essential component of evaluating color restoration performance, as human perception of color accuracy is influenced by factors beyond pixel-wise color differences.

IV. DISCUSSION

A. Evaluation of Denoising Results

The denoising model demonstrated effective noise reduction capabilities across both datasets, with performance improving as training progressed. The CIFAR-10 model achieved a final PSNR of 22.78 dB after 20 epochs, representing a significant improvement in image quality compared to the noisy inputs. However, the results reveal an inherent trade-off between noise removal and detail preservation that is characteristic of image denoising approaches.

At early training stages (5 epochs), the model behaved similarly to a Gaussian filter, effectively removing noise but

introducing noticeable blur. This is particularly problematic for the low-resolution CIFAR-10 images (32×32 pixels), where fine details are already limited. As training progressed, the model learned to better preserve image structures while still removing noise, though some level of detail loss remained unavoidable.

The ImageNet denoising model, working with higher resolution images (64×64 pixels), demonstrated similar behavior but with better detail preservation due to the increased pixel information available. This suggests that the effectiveness of CNN-based denoising approaches scales positively with image resolution, a finding that has important implications for practical applications.

One limitation of our approach is the use of synthetic noise patterns (Gaussian and salt-and-pepper) rather than real-world noise. While this provides a controlled environment for model evaluation, it may not fully represent the complex noise patterns encountered in practical scenarios, such as those from low-light photography or medical imaging. Future work could explore training on datasets with authentic noise patterns or developing models that can adapt to various noise types without explicit training on each.

B. Analysis of Color Restoration Performance

The color restoration task, regardless of the dataset, presented a more complex challenge due to the inherently ambiguous nature of the problem. Unlike denoising, where a clear ground truth exists, color restoration involves predicting plausible color information that may have multiple valid solutions. This ambiguity is reflected in the behavior of both developed models during training and in their final results.

The progression from sepia-toned outputs in early training epochs to more diverse and vibrant colorization after extended training was particularly evident with the ImageNet-based model, and to a degree with the CIFAR-10 model, demonstrating the model's gradual learning of color distributions associated with different image features. For instance, the ImageNet model's initial conservative approach (applying similar colors across different images) evolved into more confident and varied color predictions as the model refined its understanding of the relationship between grayscale patterns and potential color mappings over a significant number of epochs.

An interesting observation, again most notably with the ImageNet model due to the higher image detail and diversity, was its occasional deviation from the original colors while still producing visually plausible results. This suggests that the model has learned general color statistics rather than exact color mappings, allowing it to generate reasonable colorizations even when the exact ground truth is ambiguous. For instance, the ImageNet model might color a car blue instead of red, or change the background color of a landscape while maintaining natural-looking color relationships. This behavior highlights both a strength and limitation of the approach: while the model can produce aesthetically pleasing colorizations, it cannot guarantee historical accuracy for applications where exact color reproduction is critical. This limitation is inherent to the problem rather than specific to our implementation, as grayscale images fundamentally lack the chromatic information needed for perfect color recovery.

C. Computational Considerations

Both models required significant computational resources, particularly for training on the larger ImageNet dataset. The implementation of mixed-precision training and efficient batch processing helped mitigate memory constraints, but training the color restoration model to convergence (100 epochs) remained computationally intensive.

The choice of normalization technique proved crucial for balancing performance and resource utilization. While batch normalization performed well with adequate batch sizes, its effectiveness diminished with smaller batches necessitated by memory limitations. Group normalization offered a viable alternative in these scenarios, though it became redundant when sufficient computational resources were available.

These findings underscore the importance of considering hardware constraints and optimization techniques when developing deep learning models for image processing tasks, particularly when scaling to larger datasets or higher resolution images.

V. Conclusion

This research has demonstrated the effectiveness of convolutional neural networks in addressing two fundamental image enhancement challenges: denoising and color restoration. Through careful architecture design and training methodology, we have developed models capable of removing various noise types from degraded images and restoring plausible color information to grayscale inputs.

The denoising model successfully reduced both Gaussian and salt-and-pepper noise while preserving important image features, achieving a PSNR of 22.78 dB on the CIFAR-10 dataset. The results highlight the trade-off between noise removal and detail preservation, with performance improving as training progressed. The model's effectiveness scaled positively with image resolution, suggesting better results could be achieved with higher-resolution inputs.

The color restoration model demonstrated the ability to transform grayscale images into plausible color representations, with results improving significantly over extended training periods. The model's progression from conservative, sepia-toned outputs to more vibrant and diverse colorization illustrates the gradual learning of complex color distributions associated with different image features. Notably, the model occasionally produced alternative but visually plausible color schemes, reflecting the inherently ambiguous nature of the colorization task.

Several architectural components proved crucial to the success of both models. Residual connections facilitated the training of deeper networks by addressing the vanishing gradient problem and preserving low-level details. The choice of normalization technique significantly impacted performance, with batch normalization excelling in scenarios with adequate batch

sizes and group normalization offering a viable alternative when memory constraints necessitated smaller batches.

Future work could explore several promising directions. First, training on datasets with authentic noise patterns rather than synthetic noise could enhance the denoising model's applicability to real-world scenarios. Second, incorporating semantic understanding into the color restoration process could improve accuracy for objects with strong color associations. Finally, developing an end-to-end pipeline that combines both denoising and color restoration could provide a comprehensive solution for restoring heavily degraded historical images or enhancing low-quality visual content.

In conclusion, this research contributes to the field of image processing by providing insights into neural network architectures optimized for image enhancement tasks and demonstrating their practical applications in restoring degraded visual content. The results underscore the potential of deep learning approaches to address complex image restoration challenges, while also highlighting the inherent limitations and trade-offs involved in these tasks.

REFERENCES

[1] A. McAndrew, "An Introduction to Digital Image Processing with MATLAB, Notes for SCM2511 Image Processing 1". [Accessed 31 January 2025].

[2] F. Dong et al., "Salt and pepper noise removal based on an approximation of 10 norm," *Comput. Math. Appl.*, vol. 70, no. 5, pp. 789-804, 2015. DOI: 10.1016/j.camwa.2015.05.026.

[3] R. C. Gonzalez and R. E. Woods, *Digital Image Processing (Third Edition)*. Upper Saddle River, New Jersey 07458: Pearson Prentice Hall, 2008.

[4] R. Zhang, P. Isola, and A. A. Efros, "Colorful Image Colorization," *CoRR*, vol. abs/1603.08511, 2016. [Online]. Available: http://arxiv.org/abs/1603.08511.

[5] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," 2015. [Online]. Available: https://arxiv.org/abs/1512.03385.

[6] M. A. Deif, H. Attar, A. Alrosan, A. A. A. Solyman, and S. M. F. Abdelaliem, "Design and development of an intelligent neck and head support system based on eye blink recognition for cervical dystonia," *Discover Applied Sciences*, vol. 6, no. 11, p. 602, 2024.

[7] A. Krizhevsky, "Learning multiple layers of features from tiny images," techreport, 2009.

[8] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 2009, pp. 248-255. doi: 10.1109/CVPR.2009.5206848.

[9] Yadav, D. P., Deepak Kumar, Anand Singh Jalal, Ankit Kumar, Kamred Udham Singh, and Mohd Asif Shah, "Morphological diagnosis of hematologic malignancy using feature fusion-based deep convolutional neural network," *Scientific Reports*, vol. 13, no. 1, p. 16988, 2023.

[10] P. Chrabaszcz, I. Loshchilov, and F. Hutter, "A Downsampled Variant of ImageNet as an Alternative to the CIFAR datasets," 2017. [Online]. Available: https://arxiv.org/abs/1707.08819.

[11] Graphic Quality Consultancy, "Introduction to Colour Models ('Spaces')," 2025. [Online]. Available: https://www.colourphil.co.uk/lab_lch_colour_space.shtml.

[12] J. M. Parmar and S. A. Patil, "Performance evaluation and comparison of modified denoising method and the local adaptive wavelet image denoising method," in *2013 International Conference on Intelligent Systems and Signal Processing (ISSP)*, 2013. DOI: 10.1109/ISSP.2013.6526883.

[13] F. Chollet, *Deep Learning with Python, 2nd ed.* New York: Manning Publications Co., 2021.

[14] Y. Wu and K. He, "Group Normalization," 2018. [Online]. Available: https://arxiv.org/abs/1803.08494.

[15] J. Brownlee, "Gentle Introduction to the Adam Optimization Algorithm for Deep Learning," 2021. [Online]. Available: https://machinelearningmastery.com/adam-optimization-algorithm-for-deep-learning/.

[16] Keras, "ExponentialDecay," [Online]. Available: https://keras.io/api/optimizers/learning_rate_schedules/exponential_decay/.

[17] R. Fisher, "Sobel edge detector," HIPR2, School of Informatics, University of Edinburgh. [Online]. Available: https://homepages.inf.ed.ac.uk/rbf/HIPR2/sobel.htm.

[18] D. P. Yadav, D. Kumar, A. S. Jalal, A. Kumar, K. U. Singh, and M. A. Shah, "Morphological diagnosis of hematologic malignancy using feature fusion-based deep convolutional neural network," *Scientific Reports*, vol. 13, no. 1, p. 16988, 2023.

[19] Momynkulov, Z., Zhandos Dosbayev, Azizah Suliman, Bayan Abduraimova, Nurzhigit Smailov, Maigul Zhekambayeva, and Dusmat Zhamangarin, "Fast detection and classification of dangerous urban sounds using deep learning," *Computers, Materials & Continua*, vol. 75, pp. 2191–2208, 2023.