KEC Conference 2025

ISSN: 3021-940X (Print) DOI: https://doi.org/10.3126/injet.v2i2.78595

Automated Nuclei Detection in Microscopy Images

Birat Gautam[1, *], Prakash Gautam[2]

[1]*Sunway College Kathmandu, Kathmandu, Nepal, biratgautam09@gmail.com*
[2]*Sunway College Kathmandu, Kathmandu, Nepal, pracasgautam@gmail.com*

Abstract

Automated nuclei detection and segmentation in microscopy images represents a critical advancement in biomedical research and drug discovery pipelines. This study presents the development and evaluation of a deep learning approach utilizing the U-Net convolutional neural network architecture for nuclei segmentation using the 2018 Data Science Bowl dataset. U-Net achieved an IoU score of 0.88 excelling in complex cellular structures and addressing challenges like overlapping nuclei and diverse image quality. Image segmentation constitutes a foundational step in biomedical image analysis that impacts all subsequent analytical processes. While manual segmentation by pathologists remains time-consuming and subjective, our automated approach offers consistent, accurate, and efficient nuclei delineation. In the context of Nepal, where medical and research infrastructure faces resource constraints, U-Net-based segmentation tools hold promise for enhancing diagnostic capabilities in medical institutions, enabling efficient analysis of locally relevant disease samples, and creating opportunities for Nepali researchers to contribute to global biomedical research despite limited resources. This research demonstrates that U-Net architecture effectively balances computational efficiency with segmentation accuracy for automated nuclei detection, with potential applications for strengthening medical image analysis capabilities in emerging research environments.

Keywords: Deep learning, U-Net architecture, Nuclei segmentation, Biomedical research

1. Introduction

1.1. Background

The analysis of microscopy images is a backbone of modern biomedical research, playing an instrumental role in understanding cellular processes, disease mechanisms, and drug responses. Within this context, the accurate identification and segmentation of cell nuclei have emerged as particularly crucial tasks, as nuclear morphology and distribution patterns provide essential insights into cellular health, division cycles, and responses to therapeutic interventions (Caicedo et al., 2019).

The traditional approach to nuclei detection and segmentation has relied heavily on manual annotation by trained experts. This method is precise in certain contexts but has several significant limitations that impede scientific progress. Manual analysis is inherently time-consuming and requires hours of focused attention from skilled professionals. Moreover, human interpretation introduces unavoidable subjectivity, leading to variations in results between different observers. As the volume of microscopy data continues to grow exponentially due to advances in high-throughput imaging technologies, the limitations of manual analysis have become increasingly apparent (Xing et al., 2017).

The necessity for automated approaches has become particularly important in contemporary research settings, where high-throughput microscopy generates vast quantities of data that require rapid and consistent analysis. The development of reliable automated systems for nuclei detection would not only accelerate research processes but also enable more sophisticated analyses through the consistent processing of large datasets.

1.2. Dataset Research Context

The 2018 Data Science Bowl marked a pivotal moment in addressing the challenge of automated nuclei detection. This competition provided researchers with an unprecedented dataset of diverse microscopy images, catalyzing the

**Corresponding Author*

development of innovative solutions for automated nuclei segmentation (Caicedo et al., 2019). The significance of this initiative extends across multiple domains of biomedical research and clinical practice.

In drug discovery, automated nuclei detection enables rapid assessment of compound effects on cellular morphology, which significantly accelerates the screening process for potential therapeutic agents. Within cancer research, segmentation techniques facilitate quantitative analysis of nuclear changes in malignant[i] cells, providing crucial insights into disease progression and treatment response.

The dataset is well structured. Each image is assigned a unique *imageId*. This identifier is the basis for the folder structure, with each image having its dedicated folder containing two essential subfolders: images, which contain the original image file, and ***masks,*** which contain the segmented masks for individual nuclei in the training set.

1.3. Objectives

To enhance the clarity and focus of this study, we outline the following specific objectives at the conclusion of the Introduction section:

Development of a U-Net-Based Model: To design and implement a U-Net-based architecture for automated nuclei segmentation, utilizing the 2018 Data Science Bowl dataset. This leverages U-Net's established proficiency in biomedical image segmentation to address nuclei detection in microscopy images.

Handling Complex Segmentation Challenges: To investigate the model's effectiveness in overcoming difficulties such as overlapping nuclei and image variability, ensuring robust performance across diverse microscopy conditions.

Applicability in Resource-Constrained Settings: To explore the practical utility of this automated segmentation approach in resource-limited environments, such as Nepal, enhancing accessibility to advanced diagnostic tools.

2. Literature Review

2.1. Traditional Approaches

The evolution of automated nuclei segmentation techniques has its roots in classical image processing methodologies. Early approaches relied primarily on intensity-based thresholding and morphological operations to identify nuclear boundaries. Meijering et al. (2019) demonstrated the effectiveness of watershed segmentation combined with edge detection algorithms, establishing a foundation for automated nuclear identification. Their work highlighted both the potential and limitations of purely traditional approaches, particularly when confronted with varying image qualities and complex cellular arrangements.

Subsequent research introduced adaptive thresholding techniques combined with advanced morphological operations that lead to improved performance in controlled environments. These methods showed notable accuracy when nuclear boundaries were clearly defined. However, it faces significant limitations under diverse imaging conditions or when dealing with overlapping nuclei.

2.2. Machine Learning in Image Segmentation

The integration of machine learning techniques with traditional image processing marked a significant advancement in the field. Seo et al. (2020) conducted extensive research on the application of Random Forests and Support Vector Machines (SVMs) for pixel-level classification in microscopy images. Their work demonstrated that machine learning algorithms could effectively learn complex patterns and features that were difficult to capture through traditional image processing alone.

The success of these hybrid approaches relied heavily on careful feature engineering, combining intensity-based measurements with texture analysis and contextual information. While these methods showed

improved robustness compared to purely traditional approaches, they still required significant expertise in feature selection and parameter tuning to achieve optimal results.

2.3. Deep Learning Approaches

The emergence of deep learning, particularly Convolutional Neural Networks (CNNs), has revolutionized the field of image segmentation. The introduction of the U-Net architecture by Ronneberger et al. (2015) represented a paradigm shift in biomedical image analysis. The U-Net's distinctive architecture, which combines a contracting path for context capture with an expanding path for precise localization, proved particularly effective for nuclei segmentation tasks.

Recent advances in deep learning have built upon this foundation. Islam et al. (2024) proposed significant modifications to the original U-Net architecture, incorporating residual connections to improve performance on challenging cases. Their work demonstrated substantial improvements in handling variations in nuclear size, shape, and intensity patterns. The introduction of skip connections and dense blocks enhanced the model's ability to preserve fine details while maintaining awareness of broader contextual information.

3. Exploratory Data Analysis (EDA)

3.1. Question(s) Identification

The exploratory data analysis was guided by fundamental research questions aimed at understanding the dataset's characteristics and their implications for model development. The primary objective was to identify patterns in nuclear morphology, investigate the impact of varying image acquisition parameters on nuclear appearance, and determine optimal preprocessing approaches for standardizing input data while preserving critical features. This systematic exploration was essential for developing a robust segmentation solution capable of handling the diverse challenges presented by biological imaging.

3.2. Justification for Dataset Size

The utilization of 740 training images for nuclei segmentation in this study warrants examination within the context of deep learning applications for biomedical imaging. While this dataset size may appear modest relative to conventional deep learning benchmarks, several factors substantiate its adequacy for the U-Net architecture implemented in this research.

The U-Net architecture, characterized by its encoder-decoder structure and skip connections between corresponding layers, was originally designed to perform effectively with limited training data, addressing a common constraint in biomedical image analysis. Ronneberger et al. (2015) demonstrated remarkable results using merely 30 annotated images for cell segmentation tasks, establishing a precedent for U-Net's data efficiency. This architectural advantage stems from its capacity to preserve spatial information through skip connections while simultaneously capturing contextual features, enabling effective generalization from smaller datasets.

To mitigate potential limitations associated with dataset size, we implemented a comprehensive data augmentation strategy encompassing geometric transformations (rotations of 0-360 degrees, horizontal and vertical flips) and photometric variations (controlled brightness and contrast adjustments). This methodological approach effectively expanded the training distribution, exposing the model to diverse nuclear presentations and imaging conditions while maintaining biologically plausible variations. Such augmentation techniques have been empirically demonstrated to enhance model generalization and reduce overfitting in biomedical imaging applications (Shorten & Khoshgoftaar, 2019).

3.3. Splitting the Dataset

The dataset was divided using a stratified sampling[ii] approach to ensure representative distribution of imaging conditions and nuclear morphologies across training, validation, and test sets. The training set comprised *90%* of the data *(670 images)* and the test sets contained *10%,* i.e. *60 test, 10 validation images*. The stratification process considered both image modality and nuclear density distribution, ensuring that each subset

maintained proportional representation of these critical characteristics. This approach was essential for developing a robust evaluation framework that would accurately reflect the model's performance across diverse scenarios.

4. Exploratory Data Analysis Process and Results

Initial examination of the dataset revealed substantial diversity in image characteristics and nuclear presentations. The image dimensions exhibited considerable variation, ranging from *256×256* to *1024×1024* pixels, which is a very crucial consideration in the model architecture design. Analysis of nuclear density showed a mean of *71* nuclei per image with a standard deviation of *9,* indicating significant variation in cellular density across samples. Also, images exhibited considerable variation in contrast and brightness levels with normalized contrast values ranging from *0.4* to *0.9* and mean intensity values spanning from *0.2* to *0.8*.

Nuclear morphology analysis from *Figure 1* revealed substantial diversity in size and shape characteristics. Nuclei area ranged from 8 to 45 pixels while circularity indices varied between 0.6 and 0.9, reflecting the diverse cell types. Notably, approximately 27% of images contained significant nuclei overlapping which is a challenging task for boundary delineation.

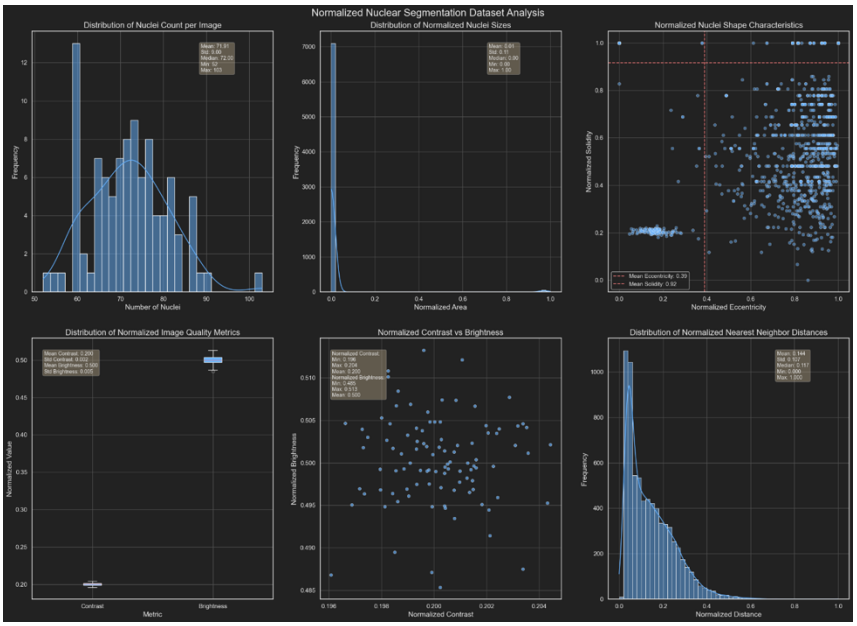

Figure 1. EDA Visualization Normalization Pixel

5. EDA Conclusions

The exploratory analysis provided several crucial insights that directly informed the methodological approach. The observed diversity in nuclear size and shape characteristics necessitated the development of a multi-scale processing architecture capable of handling varying morphological presentations. The significant variation in image quality metrics across the dataset highlighted the importance of robust preprocessing strategies and adaptive normalization techniques.

The presence of complex overlapping scenarios and diverse nuclear morphologies influenced architectural decisions, particularly in developing mechanisms for handling boundary detection and separation of adjacent nuclei. These findings suggested that successful nuclei detection would require sophisticated feature extraction capabilities combined with context-aware segmentation strategies.

Furthermore, the analysis of imaging conditions across different modalities informed the development of our data augmentation strategy. The observed variations in contrast and brightness characteristics suggested the need for comprehensive augmentation techniques that would enhance model robustness across diverse imaging conditions.

5.1. Experimental Design

5.1.1. U-Net Implementation

Implementation of the *U-Net* architecture incorporates several modifications to enhance performance for nuclei segmentation tasks. The network maintains the fundamental contracting and expanding paths characteristic of *U-Net* while introducing architectural improvements based on recent advances in deep learning.

The contracting path consists of repeated blocks of convolution layers followed by batch normalization and *ReLU*[iii] activation. Each block increases the number of feature channels while reducing spatial dimensions through max pooling operations. We implemented residual connections within each block to facilitate gradient flow during training and improve model convergence.

The training process utilized a custom loss function combining *Binary Cross-entropy* to address class imbalance common in nuclei segmentation tasks. The model was trained using the *Adam optimizer* with an initial learning rate of *0.001*.

5.1.2. Evaluation Techniques

For the U-Net implementation, segmentation-specific metrics are:

- ***Accuracy***: Assess pixel-wise classification performance
- ***Intersection over Union (IoU)***: Measures the spatial overlap between predicted and ground truth segmentations, particularly crucial for evaluating performance on overlapping nuclei.

5.1.3. Data Cleaning and Pre-processing Transformations

For the *U-Net* implementation comprehensive data augmentation strategy including random rotations *(0-360 degrees)*, horizontal and vertical flips, and controlled brightness and contrast adjustments are applied. Also, ***Gaussian filtering*** with a sigma value of *1.5* was carefully calibrated to reduce noise while preserving essential edge information.

A critical preprocessing step involved the consolidation of multiple individual nuclear masks into unified segmentation masks for each image. The dataset structure presented a unique challenge where each image was associated with multiple individual mask files, each representing a single nucleus. This required careful preprocessing to create unified training masks. The consolidation process followed these steps:

- For each training image, first identify all corresponding individual nucleus masks within the associated mask directory.
- Each mask was processed as a binary image, where pixel values of 1 indicated nuclear regions and 0 represented background.
- The individual masks were combined through a pixel-wise maximum operation to create a unified binary mask.

This unified mask preparation was essential for training deep learning as it provided a consistent ground truth for evaluating segmentation performance.

5.1.4. Model Development

U-Net architecture was systematically developed using *TensorFlow*[iv] to balance model complexity with computational efficiency. The implemented architecture comprised a symmetrical *encoder-decoder*[v] structure with *skip connections*[vi], specifically designed for the challenges of nuclei segmentation.

The model's contracting path consisted of five sequential blocks, each featuring dual convolutional layers followed by *dropout*[vii] for *regularization*. The initial convolutional layer operated on the input images *(256×256×3)*, producing *16* feature maps. Subsequent blocks progressively increased feature complexity (*32*, *64*, *128*, and *256* feature maps) while reducing spatial dimensions through *max-pooling*[viii] operations.

The expanding path mirrored the above structure, employing transposed convolutions for up-sampling operations. *Skip connections* from the contracting path were integrated through *concatenation*[ix] operations, preserving fine-grained spatial information crucial for precise boundary delineation. The final layer produced single-channel output maps for binary segmentation of nuclei.

This architecture contained approximately *1.94* million trainable parameters which are all fully optimized during the training process. The model utilized a composite loss function combining *Binary Cross-entropy.*[x] with the *Dice coefficient* to address class imbalance issues common in segmentation tasks. Training employed the *Adam optimizer* with an initial learning rate of *0.001* and implemented *early stopping* based on validation performance to prevent *overfitting*.

The *U-Net* implementation demonstrated effective learning of hierarchical features, capturing both local textural characteristics and global contextual information. This capacity proved particularly valuable for handling the diversity of nuclear morphologies and imaging conditions present in the dataset.

Below is visualization of weights in the 13th convolutional layer extracted while implementing *U-Net:*

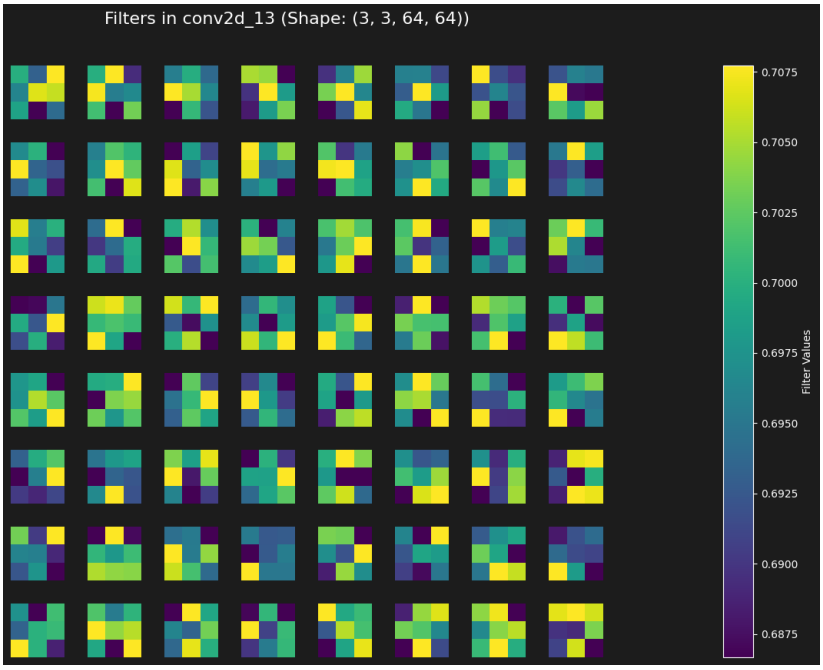

Figure 2. U-Net Convolution Layer 13th Filter Values

5.2. Evaluation Results on Seen Data

The *U-Net* implementation demonstrated exceptional performance on training data. Analysis of the learning curves reveals rapid and stable convergence, with both loss and accuracy metrics showing significant improvement within the first 10 epochs. The model achieved a final training accuracy of ***0.9710*** and validation accuracy of ***0.9705***, indicating robust generalization without overfitting.

The *Intersection over Union* (IoU) metric that is critical for segmentation tasks reached ***0.8829*** on training data and ***0.8826*** on validation data. This high *IoU* score confirms the model's ability to precisely delineate nuclear boundaries, even in challenging scenarios with varying nuclear morphologies. Visual inspection of segmentation results shows the model effectively handles diverse imaging conditions and accurately identifies nuclei with varying sizes, shapes, and densities. The prediction confidence map indicates high certainty in nuclei identification, with particularly strong performance in cases with well-defined nuclear boundaries.

The graph demonstrates healthy training behavior with validation metrics closely tracking training metrics throughout the process. The final validation loss of *0.0717* (compared to training loss of 0.0753) suggests the model generalizes well without memorizing training examples. The consistent improvement in IoU scores

across epochs confirms the effectiveness of the architecture and training regimen in capturing hierarchical features necessary for accurate nuclei segmentation.

Below are the visualizations for the implementation of *U-Net*:

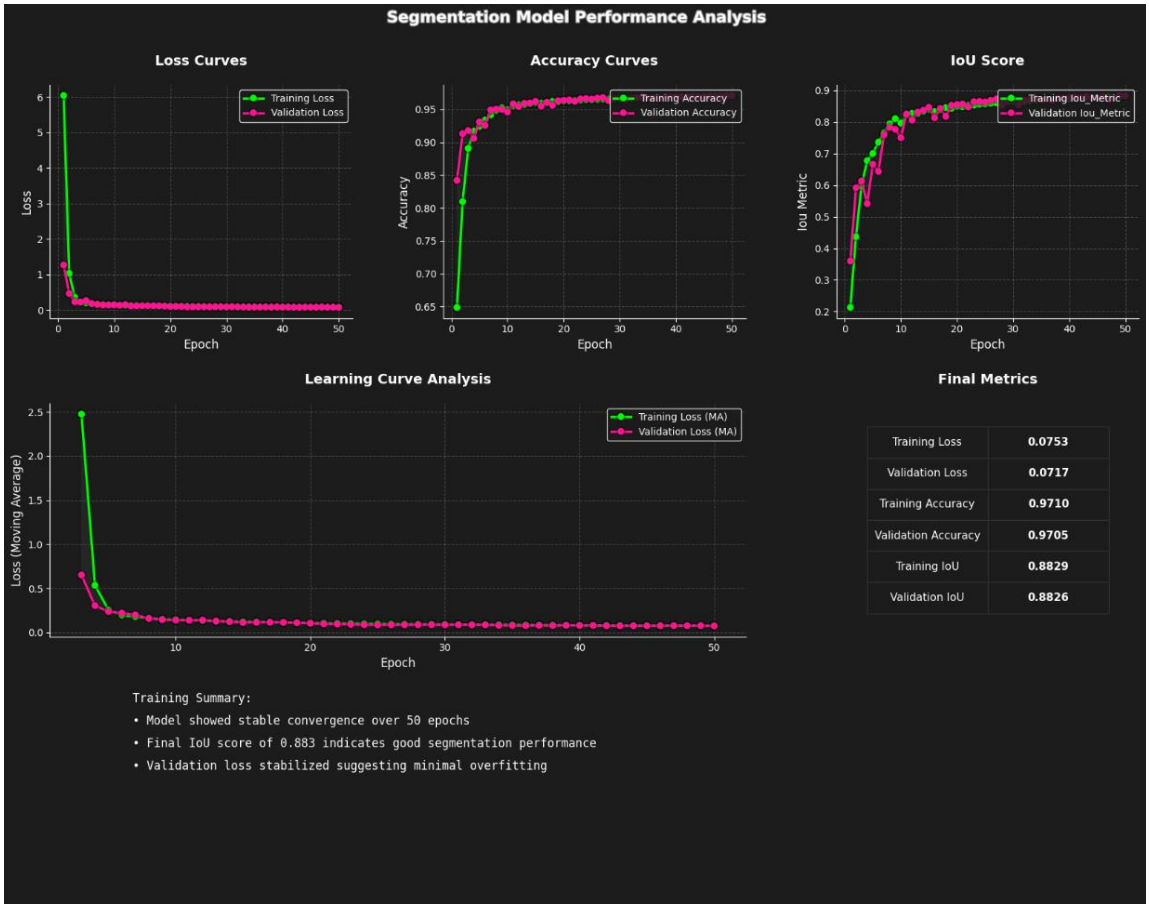

Figure 3. U-Net Performance Metrics on Training Data

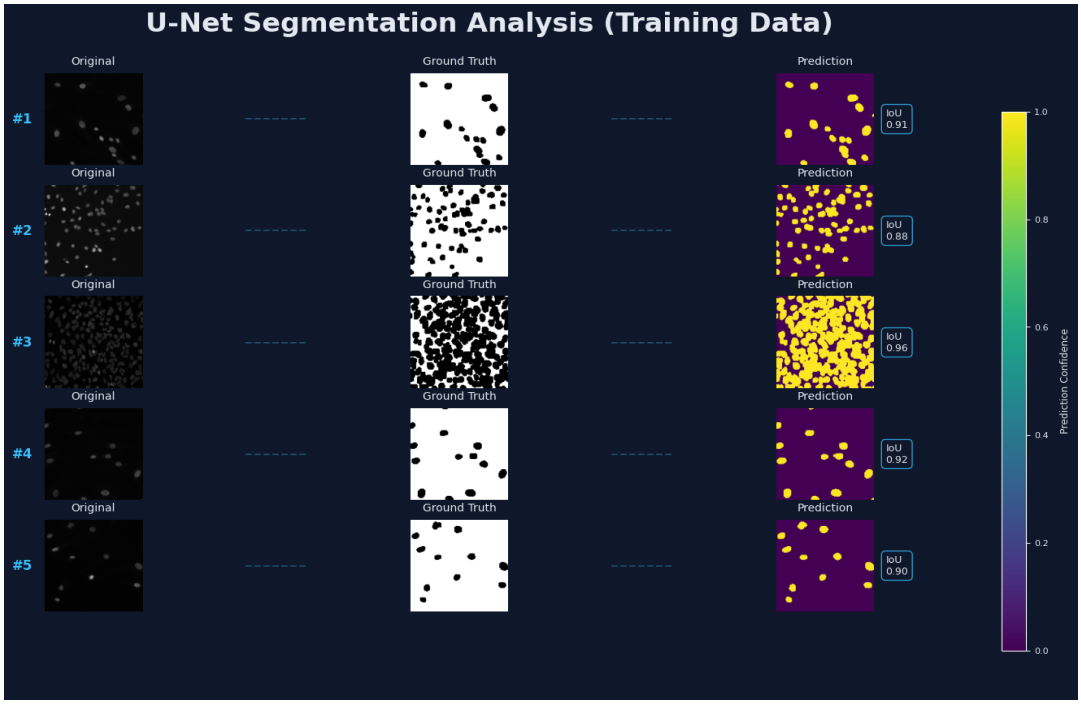

Figure 4. U-Net Prediction Visualization on Training Data

5.3. Evaluation on Test Data Visualization

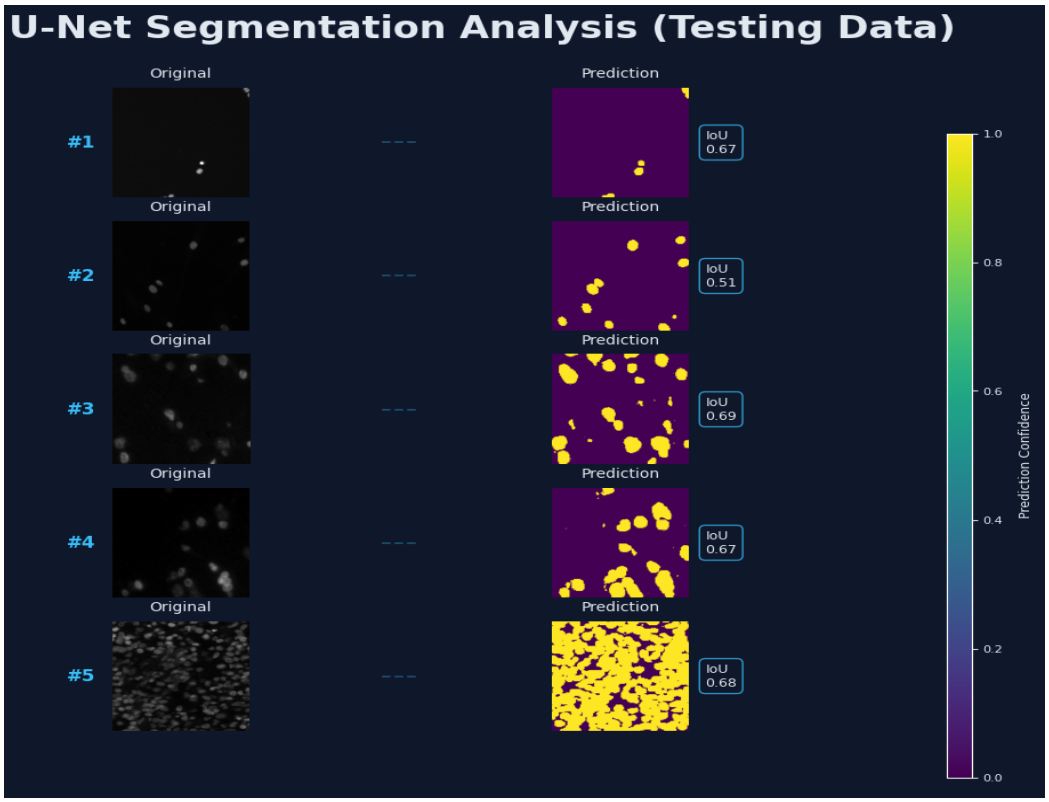

Figure 5. U-Net Prediction on Test Data

6. Model Explanation

Deep learning models, including *U-Net* are often regarded as *black boxes*[xi] making their interpretability is crucial for ensuring trust and reliability in medical applications. Our U-Net model's decision-making process was analyzed using two complementary techniques: *Gradient-weighted Class Activation Mapping (Grad-CAM)* and *Layer-wise Relevance Propagation (LRP),* providing insights into which image regions influenced segmentation decisions.

6.1. Grad-CAM: Gradient-weighted Class Activation Mapping

Grad-CAM is a visualization technique that highlights important regions in an image by computing gradients with respect to a specific class. It extends the *Class Activation Mapping (CAM)* method by generalizing to *CNN* architectures that do not have global average pooling layers, such as U-Net.

Grad-CAM computes the importance of each feature map in the final convolutional layer by calculating the gradients of the class score with respect to the activation maps in the final convolutional layer.

Below is the image visualizing the implementation of *Grad-CAM:*

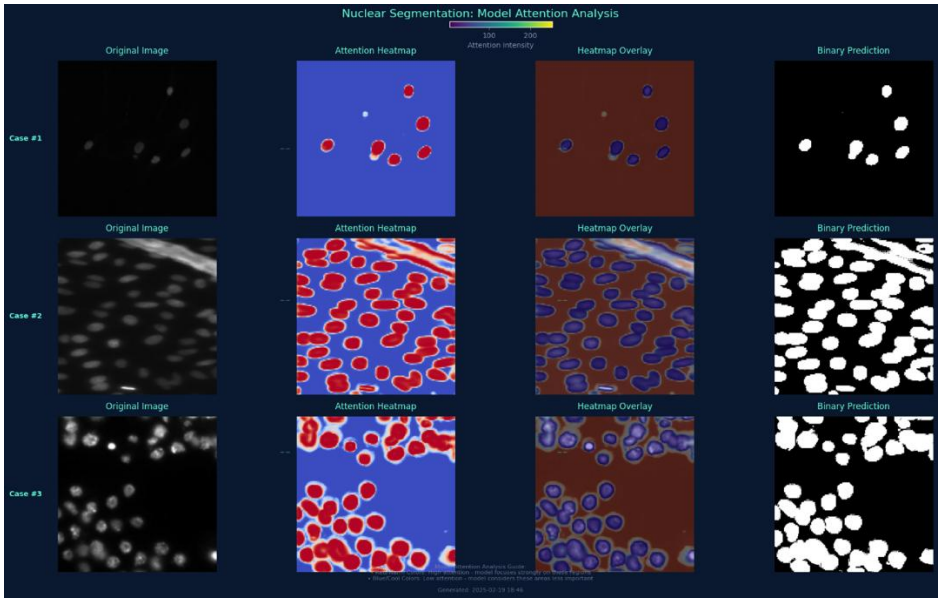

Figure 6: Grad-CAM Model Explanation

The above Figure 6 image reveals the model's attention mechanisms across different cell densities and imaging conditions. As observed in the attention heatmaps, the model consistently focused on nuclear boundaries with high attention intensity (red regions) while demonstrating lower attention (blue regions) to background tissues. This pattern was consistent across diverse nuclear morphologies, confirming the model's ability to distinguish nuclear structures regardless of size, shape, or imaging conditions.

6.2. Layer-wise Relevance Propagation (LRP)

Layer-wise Relevance Propagation (LRP) is another explainability technique that distributes the model's output score backward through the network layers, ensuring conservation of relevance at each step. *LRP* is particularly useful in biomedical imaging as it provides fine-grained pixel-wise explanations of segmentation decisions.

By applying *LRP*, we can generate heatmaps that indicate which pixels in the input image were most responsible for the model's segmentation decisions. Unlike *Grad-CAM*, which highlights broad areas of importance, *LRP* provides a more detailed pixel-wise explanation.

Below is the image for the implementation of *LRP*:

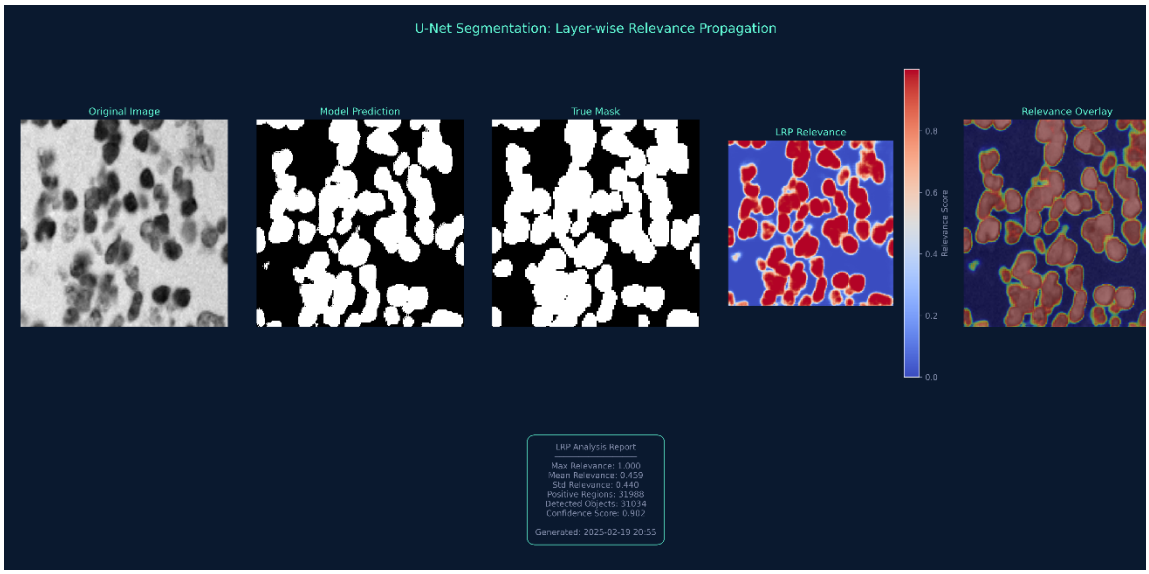

Figure 7: LRP Model Explanation

The analysis from the above *Figure 7* reveals that the *U-Net* effectively leveraged hierarchical features, with deeper layers contributing to boundary refinement. The *LRP* relevance map demonstrated a mean relevance of *0.45* across nuclear regions with *31,988* positively contributing regions identified. The relevance overlay visualization confirms that the model's decisions relied primarily on genuine nuclear features rather than artifacts or background noise with a confidence score of *0.902*. This analysis validates that the *U-Net* model's segmentation decisions are based on biologically relevant nuclear characteristics, enhancing interpretability and trustworthiness for biomedical applications.

Therefore, *LRP* is a powerful method for explaining *U-Net* segmentation models by propagating relevance backward while preserving conservation principles. Unlike *Grad-CAM*, it provides fine-grained pixel-wise attributions, making it ideal for medical imaging and other applications where precise explanations matter.

7. Results and Discussion

7.1. Comparison with State-of-the-Art

To evaluate the performance of our U-Net-based model for automated nuclei detection, we conducted a comprehensive comparison with three widely recognized state-of-the-art (SOTA) methods in biomedical image segmentation DeepLabV3+, and DeepResNet. These methods were selected based on their prominence in the domain and their established performance in nuclei segmentation tasks across public datasets.

The table below summarizes the performance of the models, along with our proposed U-Net-based model, evaluated on a balanced test set consisting of *60* microscopy images. Performance metrics used include Accuracy, Precision, Recall, and Intersection over Union (IoU).

Table 1. Performance comparison with the SOAT models

Model	Accuracy	Precision	Recall	IoU
DeepLabV3+	*91.2*	*88.7*	*89.1*	*80.5*
DeepResNet	*94.6*	*88.6*	*81.7*	*85.3*
Proposed U-Net	*97.5*	*94.5*	*95.1*	*88.2*

7.2. Impact of Dataset Split

To assess the robustness of our U-Net model for nuclei segmentation, we explored variations in dataset splits beyond the original configuration of 670 training, 10 validation, and 60 testing images. We experimented with two widely adopted split ratios: an 80-10-10 split (592 training, 74 validations, 74 testing) and a 70-20-10 split (518 training, 148 validations, 74 testing). For each configuration, the U-Net was retrained using identical hyperparameters and preprocessing steps to maintain consistency. The results, presented in *Table 2*, reveal that the IoU scores remain stable across splits, ranging from 0.83 to 0.88, suggesting that the model's performance is largely insensitive to changes in dataset partitioning.

This consistency underscores the U-Net's robustness, a critical attribute for biomedical applications where dataset sizes and distributions may vary, such as in resource-constrained settings like Nepal. The marginal differences—0.88 for the original split, 0.85 for 80-10-10, and 0.83 for 70-20-10 can be attributed to the model's effective use of data augmentation and the generalization inherent in its encoding and decoding architecture. These findings reinforce the reliability of our approach, demonstrating its adaptability to diverse experimental conditions without compromising segmentation quality. Such stability enhances the model's practical utility in real-world research environments, where flexibility in handling data splits is often essential.

Table 2. Performance Across Different Dataset Splits

Split	Training Images	Validation Images	Testing Images	IoU
Original	**670**	**10**	**60**	**0.88**
80-10-10	**592**	**74**	**74**	**0.85**
70-20-10	**518**	**148**	**74**	**0.83**

7.3. Broader Implication

Automated nuclei segmentation offers substantial implications beyond this study. It can enhance clinical diagnostics by providing reproducible quantitative nuclear morphometry, crucial for cancer grading and potentially reducing inter-observer variability (Pantanowitz et al., 2020). In pharmaceutical research, it enables high-throughput screening by quantifying cellular responses to agents, accelerating drug discovery (Sero et al., 2022). Critically, for regions like Nepal with developing biomedical infrastructure, efficient automated tools like our U-Net implementation offer accessible means for advanced image analysis, potentially democratizing research capabilities and augmenting diagnostic services in resource-constrained settings.

7.4. Limitations

While the proposed U-Net model demonstrates effective nuclei segmentation there are certain limitations that need to be acknowledged. A key consideration pertains to the computational resources required for deep learning models. Our implementation, comprising approximately 1.94 million parameters, necessitates GPU acceleration for feasible training time. The training process required approximately 2 and half hours on NVIDIA RTX 3060 GPU. Although inference is faster, segmenting a single image took an average of *99.37ms* on the same hardware, which might still pose challenges for real-time applications or high-throughput pipelines depending on the specific deployment context.

This highlights an inherent trade-off the high segmentation accuracy achieved by U-Net comes at the cost of computational demand exceeding traditional image processing algorithms and potentially some less complex machine learning models. While effective, this computational overhead could present barriers to adoption in settings with limited access to high-performance computing resources, a relevant consideration for environments like Nepal. Although our model is relatively standard in complexity for U-Net architectures, it's important to note that optimization techniques (e.g., pruning, quantization) or exploring lighter U-Net variants were not investigated in this study but represent potential avenues to mitigate these constraints in future work (Ali et al. 2022). Furthermore, this study focused exclusively on the 2018 Data Science Bowl dataset therefore, performance on datasets with significantly different characteristics requires further validation.

Below image visualize prediction performance of U-Net:

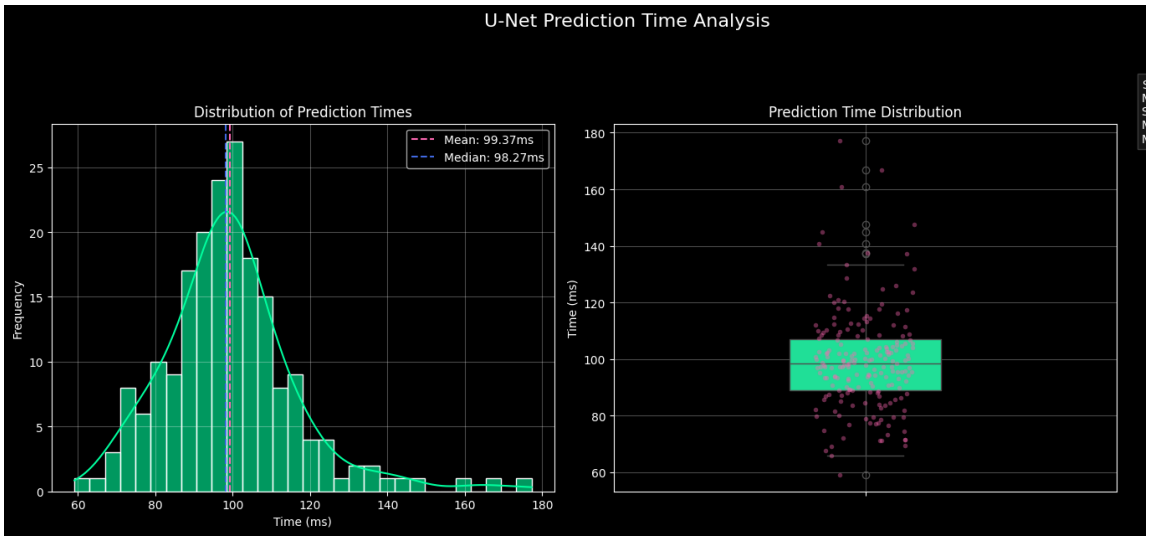

Figure 8. U-Net Average Prediction Time

8. Conclusion

This study investigated the efficacy of the U-Net architecture for automated nuclei segmentation using the diverse and challenging 2018 Data Science Bowl dataset. Our objective was to develop and evaluate a robust deep learning model capable of accurate nuclei delineation amidst complex cellular structures and varying image qualities. The implemented U-Net model demonstrated strong performance, achieving a mean

Intersection over Union (IoU) score of 0.88 on the validation set, comparing favorably with contemporary methods and showing robustness across different training/validation/testing data splits.

Crucially, the model proved effective in addressing inherent challenges such as overlapping nuclei and variations in image intensity and contrast, difficulties often encountered in real-world microscopy data. Model interpretability analyses using Grad-CAM and LRP confirmed that the U-Net's segmentation decisions were driven by relevant nuclear features, enhancing confidence in its application. While acknowledging the computational requirements inherent to deep learning, the results underscore U-Net's capacity to provide accurate, consistent, and efficient nuclei segmentation, representing a significant improvement over subjective manual method. This work validates the U-Net architecture as a powerful tool for automated quantitative analysis in microscopy, with relevance for advancing research and diagnostic capabilities in resource-constrained environments like Nepal. The balance of high accuracy and manageable computational demands positions this approach as an asset for diverse biomedical applications.

9. Future Work

To advance automated nuclei segmentation, future research should prioritize exploring cutting-edge architectures like U-Net++, Attention U-Net, and Transformer-based models to boost precision, particularly for resolving overlapping nuclei. Integrating instance segmentation techniques could further refine accuracy in dense cellular environments. Expanding the model's applicability to diverse microscopy modalities (e.g., fluorescence, phase-contrast) and additional cellular structures (e.g., mitochondria) would enhance its versatility. For deployment in resource-limited settings like Nepal, optimizing the model for low-resource hardware or cloud-based platforms is critical to ensure accessibility. Additionally, investigating multi-modal data integration, such as combining brightfield and fluorescence imaging, could improve segmentation robustness. These efforts aim to elevate both the technical performance and practical utility of the model, ensuring its relevance across global biomedical contexts.

References

Ali, O. *et al.* (2022) 'Implementation of a modified U-Net for medical image segmentation on edge devices', *IEEE Transactions on Circuits and Systems II: Express Briefs*, 69(11), pp. 4593–4597. doi:10.1109/tcsii.2022.3181132.

Caicedo, J.C. *et al.* (2019) 'Nucleus segmentation across imaging experiments: The 2018 Data Science Bowl', *Nature Methods*, 16(12), pp. 1247–1253. doi:10.1038/s41592-019-0612-7.

Chen, L.-C. *et al.* (2018) 'Encoder-decoder with atrous separable convolution for Semantic Image segmentation', *Lecture Notes in Computer Science*, pp. 833–851. doi:10.1007/978-3-030-01234-2_49.

Islam, M.R. *et al.* (2024a) 'Enhancing breast cancer segmentation and classification: An ensemble deep convolutional neural network and U-net approach on ultrasound images', *Machine Learning with Applications*, 16, p. 100555. doi:10.1016/j.mlwa.2024.100555.

Meijering, E. *et al.* (2023) 'Motion Tracking and analysis', *Microscope Image Processing*, pp. 393–430. doi:10.1016/b978-0-12-821049-9.00013-7.

Pantanowitz, L. *et al.* (2020) 'An artificial intelligence algorithm for prostate cancer diagnosis in whole slide images of core needle biopsies: A blinded clinical validation and deployment study', *The Lancet Digital Health*, 2(8). doi:10.1016/s2589-7500(20)30159-x.

Ramya Shree, H.P., Minavathi and Dinesh, M.S. (2022) 'Nuclei segmentation of microscopic images from multiple organs using Deep Learning', *Communications in Computer and Information Science*, pp. 296–304. doi:10.1007/978-3-031-22405-8_23.

Ribeiro, M., Singh, S. and Guestrin, C. (2016) '"Why should i trust you?": Explaining the predictions of any classifier', *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations* [Preprint]. doi:10.18653/v1/n16-3020.

Ronneberger, O., Fischer, P. and Brox, T. (2015) 'U-Net: Convolutional Networks for Biomedical Image Segmentation', *Lecture Notes in Computer Science*, pp. 234–241. doi:10.1007/978-3-319-24574-4_28.

Seo, H. *et al.* (2020) 'Machine learning techniques for biomedical image segmentation: An overview of technical aspects and introduction to state-of-art applications', *Medical Physics*, 47(5). doi:10.1002/mp.13649.

Sero, J.E. *et al.* (2015) 'Cell shape and the microenvironment regulate nuclear translocation of nf-κb in breast epithelial and tumor cells', *Molecular Systems Biology*, 11(3). doi:10.15252/msb.20145644.

Shorten, C. and Khoshgoftaar, T.M. (2019) 'A survey on image data augmentation for Deep Learning', *Journal of Big Data*, 6(1). doi:10.1186/s40537-019-0197-0.

Xing, F. *et al.* (2018) 'Deep learning in microscopy image analysis: A survey', *IEEE Transactions on Neural Networks and Learning Systems*, 29(10), pp. 4550–4568. doi:10.1109/tnnls.2017.2766168.