Automated Bladder Cancer Screening with Deep Learning Algorithms

Ameya R, Joycelyn C, David A, Aqib S, Abhi K, Gagan G,
Alos Diallo, Joshua J. Levy PhD
Emerging Diagnostic and Investigative Technologies, Department of Pathology, Dartmouth Hitchcock Medical Center

Dartmouth Health

ABSTRACT

- Background: Bladder cancer has a high recurrence rate, requiring lifelong surveillance. Urine cytology is noninvasive but labor-intensive and subject to variability in assessing nuclear-to-cytoplasmic (N/C) ratio.
- Objective: Develop an automated segmentation pipeline to objectively measure N/C ratio in urine cytology images.
- Methods: Trained a U-Net with a ResNet-50 encoder on 300 urothelial cell images with ground-truth nucleus and cytoplasm masks and compared performance with Otsu thresholding.
- Results: U-Net achieved a Spearman correlation coefficient of 0.844 with the ground-truth N/C ratios, compared with 0.765 for Otsu thresholding. Otsu failed to detect nuclei in 9/200 cells.
- Conclusion: Mean N/C ratio increased across diagnostic categories (0.185 → 0.250 → 0.397 → 0.488), supporting N/C ratio as a quantitative biomarker and demonstrating the potential of deep learning for more reproducible bladder cancer screening.

INTRODUCTION

Challenges in Urine Cytology:
- Manual screening of urine cytology specimens is time consuming and laborious, can also lead to subjective N/C ratio estimates, hurting reproducibility
- Overlapping cells, debris, and low contrast break simple thresholding rules, can lead to misclassification

Intended Solutions & Next Steps:
- Train a U-Net model to segment the nucleus, cytoplasm, and background of cell image while also computing per cell N/C ratio
- Use a simple baseline such as Otsu thresholding for comparison and interpretability of per cell N/C ratio across the four diagnostic categories (atypical suspicious, positive, negative)
- Combine N/C ratios across cells for the specimen for diagnosis support

Data & Evaluation Plan:
- Urothelial cell dataset: 200 with ground truth/cytoplasm masks
- Specimen dataset: 25 cells/patients across diagnostic categories
- Grayscale preprocessing
- Compared U-NET and Otsu Method
- Evaluated using Dice score and Spearman coefficient correlation
- Analyzed N/C ratio across diagnostic categories

METHODS

Datasets

Urothelial Cell Dataset (model training):
- ❖ 300 urothelial cells, each with ground truth segmentation masks for nucleus and cytoplasm
- ❖ Used to train and benchmark segmentation approaches, and to calculate the nucleus-to-cytoplasm (N/C) ratio

Specimen Cell Dataset:
- ❖ Contains 25 cells per patient, grouped by diagnostic category
- ❖ Used to investigate whether the calculated N/C ration can be used as a predictive biomarker across patient diagnosis

Preprocessing, Models, Evaluation Metrics

Preprocessing:
- ❖ Grayscale conversion: Converted urothelial dataset images to grayscale for better intensity based analysis and to establish a baseline for N/C ratio prediction

Models:
- ❖ Otsu's Method: Segmentation Approach
- ❖ UNET: Deep Learning Approach

Evaluation Metrics:
- ❖ Dice Score Coefficient
- ❖ Spearman correlation
- ❖ Analyzed N/C ratio variation by diagnosis

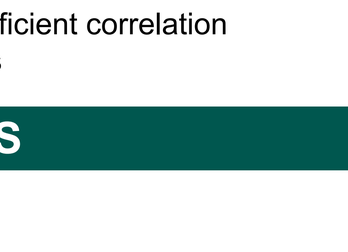

UNET

Otsu's Method:
- ❖ Classes: 3 (multi-Otsu thresholding: background, cytoplasm, nucleus)
- ❖ Preprocessing: grayscale conversion, intensity rescaling to 0-1, Gaussian smoothing (σ = 1.5, tuned via grid search), inverted so nuclei have higher intensity
- ❖ Post-processing: fill small holes in segmented regions (area threshold = 200 pixels) to clean up noisy boundaries
- ❖ N/C ratio = nucleus pixels / (nucleus + cytoplasm pixels)

UNET:
- ❖ Epochs: 50
- ❖ Batch size: 8
- ❖ Loss function: 50% Cross_Entropy_Loss + 50% Dice Loss
- ❖ Optimizer: Adam, learning rate 1e-4
- ❖ N/C ratio = nucleus pixels / (nucleus + cytoplasm pixels)
- ❖ Resize to 256×256, normalize 0–1, channels-first conversion
- ❖ 3 classes for segmentation: background, cytoplasm, nucleus

RESULTS

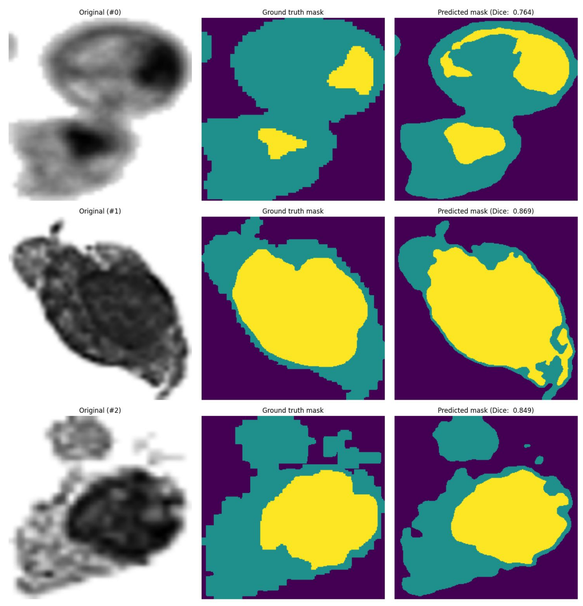

Otsu's Method mask
Original image, ground truth, and Otsu-predicted masks

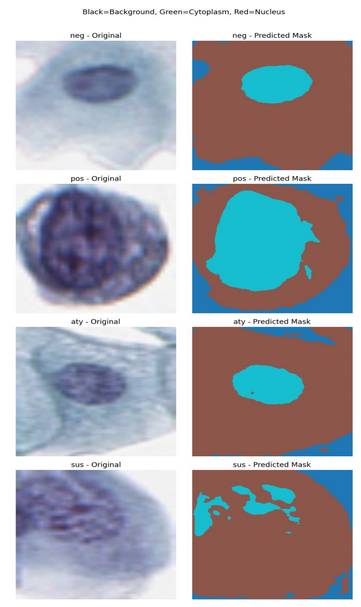

UNET mask
Original image and UNET-predicted masks

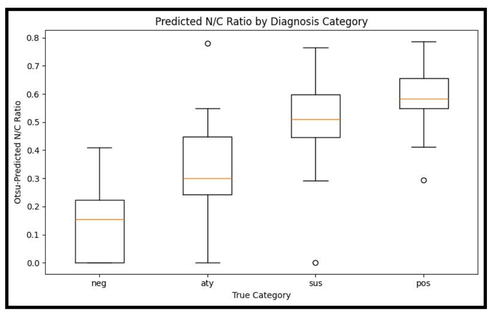

Otsu's Method: SCC = 0.765

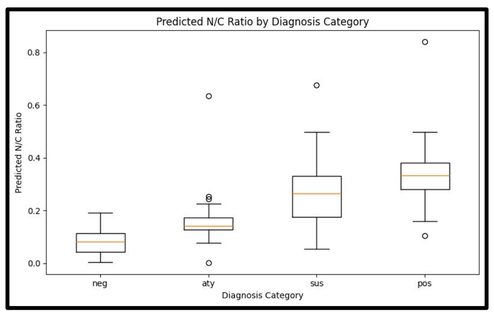

UNET: SCC = 0.844

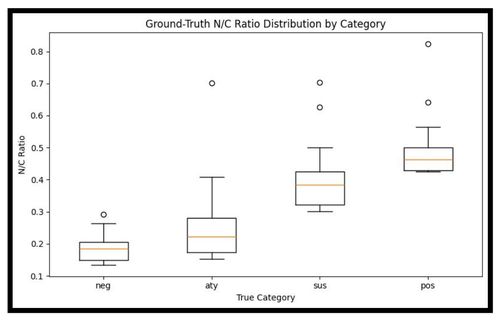

Ground truth

CONCLUSION

Future Impact:
- **Scale & Improve**: Expand to larger, multi-institution datasets and combine N/C ratio with additional cellular features for more robust predictions.
- **Target Hard Cases**: Focus on atypical/suspicious cells and improve performance where pathologists have the most difficulty reaching agreement.
- **Whole-Specimen Analysis**: Automate cell detection across entire specimens and generate patient-level screening scores.

Limitations:
- limited dataset size (~300 cells for training, 25 cells/patient validation), which limits generalizability, and causes there to be a risk of overfitting
- N/C ratio may be insufficient, bladder cancer diagnosis depends on other morphological features
- Image quality challenges could lead to segmentation errors, which could then lead to potential misclassification

Future Directions:
- **Expand the Dataset:** Include more patients, cells, and diverse clinical data from multiple institutions.
- **Improve Detection:** Refine segmentation and cell detection for crowded or atypical cells.
- **Enhance Biomarkers:** Combine N/C ratio with nuclear texture, chromatin density, and shape features.
- **Clinical Validation:** Test the model on independent datasets and evaluate its potential for real-world screening.

edit.