AUTOMATED BLADDER CANCER SCREENING WITH DEEP LEARNING ALGORITHMS

DARTMOUTH

edit.

Tanay Janmanchi, Matthew Spektor, Avilash Angirekula & Dennis Li

Abstract

- Novel Approach: Developed a cytology pipeline that segments nucleus and cytoplasm with a UNet and computes nuclear-to-cytoplasmic (N/C) ratio per cell, then bombines cell-level features—optionally via a specimen graph
- Clinical Impact: Improves consistency and throughput of urine cytology by turning subjective assessments into reproducible, quantitative measurements that can aid triage and downstream decision support.
- Future Directions: Scale to whole-slide deployment with integrated cell detection, extend to multi-class cytology labels, and conduct external validation across labs and scanners.
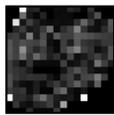
Methodology

Datasets
- Urothelial Cell Dataset (Model Training): ~300 urothelial cells, each with ground-truth segmentation masks for nucleus and cytoplasm.
  - Train and benchmark segmentation approaches and to compute the nucleus-to-cytoplasm (N/C) ratio.
- Specimen Cell Dataset (Clinical Validation): 25 cells per patient, grouped by diagnostic category.
  - Investigate whether the computed N/C ratio can serve as a predictive biomarker across patient diagnoses.

Preprocessing & Feature Engineering
- Grayscale Conversion: Converted urothelial dataset images to grayscale to simplify intensity-based analysis and establish a baseline for N/C ratio prediction.
- Feature Extraction: Generated multiple feature maps to capture complementary information:
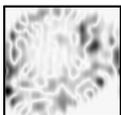
Gabor filters – highlighted local texture and frequency patterns.	GLCM (Gray-Level Co-occurrence Matrix) – captured statistical measures of local texture (contrast, homogeneity, correlation).	Sobel operator – detected edges and emphasized boundary intensity between nucleus and cytoplasm.	Multiscale features – captured structural patterns across different image resolutions and intensities.

Modeling
- K-Means Clustering: Applied to feature maps for unsupervised segmentation of nucleus vs. cytoplasm regions, enabling N/C ratio calculation.
- Random Forest (RF): Trained on extracted features to predict N/C ratios and classify cells, comparing computational predictions to ground-truth measurements.

Evaluation
- Compared predicted N/C ratios against ground-truth segmentation-derived ratios for accuracy.
- Assessed biomarker utility by analyzing N/C ratio distributions across patient diagnosis groups in the specimen dataset.

Introduction

Challenges in Urine Cytology
- Overlapping cells and low contrast break simple thresholding rules
- Manual review is slow and subjective, causing inconsistent N/C ratio estimates
- Specimen-to-specimen variability (prep, scanner) hurts reproducibility

Proposed Solution & Next Steps
- Train a UNet to segment nucleus & cytoplasm and compute per-cell N/C ratio
- Use simple baselines (Otsu/morphology, KMeans+texture) for comparison and interpretability
- Combine N/C across cells to the specimen for diagnosis support

Data & Evaluation Plan
- Urothelial cell set with expert masks → split into train/val/test with patient/specimen separation
- Specimen dataset (~25 cells/patient) with diagnosis labels for downstream association studies
- Preprocess (or tile WSIs): detect cells, crop, drop whitespace/debris/out-of-focus regions
- Report per-class Dice/IoU, per-cell N/C error, specimen-level correlations; include qualitative overlays for reviewer verification

Results

Segmentation and Feature Output
- Each segmentation method—UNet, KMeans, and Random Forest—produced arrays of pixel-level values between 0 and 1, representing the probability of a region being part of the nucleus or cytoplasm.
- These values represent the likelihood that a given pixel belongs to the nucleus or cytoplasm.
- Low values (close to 0) indicate background or empty regions, while values near 1 correspond to dense, filled-in cellular structures.


```
array([[0.25261159, 0.38812835, 0.1019802 , 0.31619544, 0.13163579, 0.1925663,
0.43393689, 0.14457287, 0.15793921, 0.60597242, 0.80309521, 0.0360488 , 0.0346577 ,
0.35569048, 0.04574554,0.41502096, 0.32888741, 0.19138345, 0.46018537,
0.36438186, 0.00806673, 0.2649845 , 0.37968875, 0.1815911 , 0.36064289, 0.19673168,
0.65604151, 0.12777501, 0.12166405, 0.56624547, 0.32622135, 0.58285507, 0.22946005,
0.61076531, 0.30216946, 0.31245248, 0.38026198, 0.70317443, 0.25243231, 0.51862919,
0.0275162 , 0.31710401, 0.15292656, 0.39803464, 0.10301465, 0.32308876,
0.09079486, 0.65975918, 0.39717922, 0.63598433, 0.37569001, 0.31195177, 0.1681291 ,
0.35780464, 0.20441367, 0.21020821, 0.0743784 , 0.14622558, 0.55169748, 0.46863978,
0.02785515, 0.30196104, 0.07609407, 0.31178087, 0.3108113 , 0.39929353, 0.27235574,
0.4509296 , 0.71095882, 0.30378774, 0.25824561, 0.27664266, 0.30284931,
0.42490341, 0.3593172, 0.18489693, 0.24464067, 0.18504853, 0.52637887, 0.20491283,
0.18725814, 0.58085815, 0.29902539, 0.11941516, 0.09073394, 0.43172001, 0.20206191,
0.17034108, 0.03573258, 0.27744543, 0.35527726, 0.16153373, 0.39780234, 0.14631854,
0.39538227, 0.07852506, 0.29991402, 0.58948897, 0.05124466, 0.20302904,
0.50842855, 0.55011781, 0.24017405, 0.53722311, 0.44692406, 0.31049182, 0.25759646,
0.26720031, 0.31107396, 0.62188718, 0.51558144, 0.31650931, 0.33558801, 0.23032805,
0.03809153, 0.36980708, 0.10427719, 0.48389411, 0.20473692, 0.17551761,
0.05304044]])
```
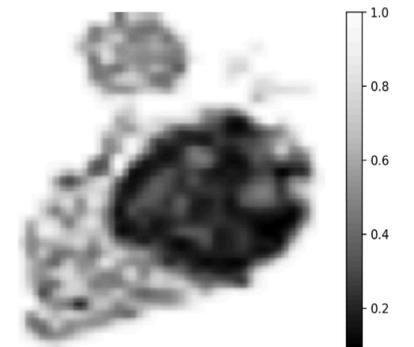
Comparative Performance & Model Insights
- The UNet model consistently outperformed baseline methods (Otsu thresholding, KMeans, and Random Forest) in segmentation accuracy, especially around complex or overlapping boundaries.
- Evaluation metrics—such as Dice coefficient and IoU—confirmed the UNet's superior performance in capturing both nucleus and cytoplasm regions.
- Feature-based approaches like Gabor filters, Sobel edges, and GLCM textures were effective for unsupervised clustering, but produced noisier masks and less consistent N/C ratio calculations compared to UNet.
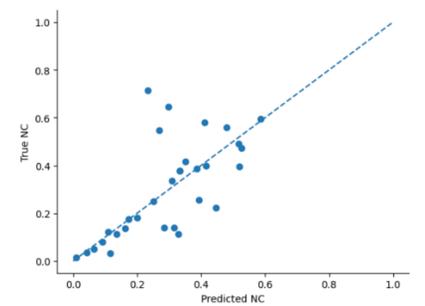
N/C Ratio Calculation & Diagnostic Trends
- Using the segmentation masks, nuclear and cytoplasmic areas were measured and used to compute the per-cell N/C ratio.
- These ratios were then aggregated at the patient level (25 cells per patient), and grouped by diagnosis category.
- Malignant samples showed elevated average N/C ratios, while benign samples showed lower values—confirming the biological relevance of the feature. The UNet-derived ratios aligned most closely with ground truth, showing reduced variance and better class separation in statistical analysis.
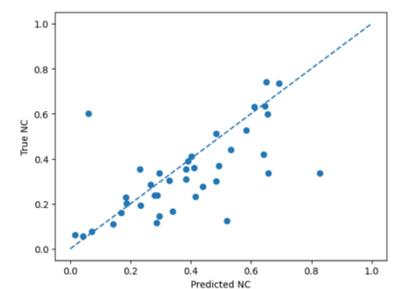
Clinical Relevance and Interpretation
- Clear differences in N/C ratio distributions emerged across diagnostic categories.
- Higher average N/C ratios correlated with malignant or atypical samples.
- UNet-derived ratios showed stronger alignment with expert-labeled clinical diagnoses, suggesting utility as a quantitative biomarker.

Conclusion

Potential for future impact
- Provides a quantitative and reproducible biomarker (N/C ratio) to support bladder cancer screening, reducing subjectivity and observer variability
  - Makes diagnostics more consistent across labs and healthcare systems.
- Can serve as a decision-support tool for clinicians, improving diagnostic confidence.
- Potential to extend beyond bladder cancer into other cytology-based screenings (e.g., cervical cancer, lung cytology).

Limitations
- Dataset size is small (~300 cells for training, 25 cells/patient for validation), which may limit generalizability and cause UNet and clustering approaches to overfit.
- Specimen variability – Differences in preparation, staining or imaging equipment can reduce reproducibility across labs.
  - Overlapping cells, debris or low contrast images may cause the model to struggle or be inaccurate

Future Directions
- Explore additional features beyond N/C Ratio (texture, shape descriptors, chromatin density) for improved prediction
- Find ways to integrate into clinical practice
- Extend models to multi-class cytology labels (negative, atypical, suspicious, positive) as opposed to using binary markers