DiffuPath: Diffusion-Guided Waterborne Pathogen Detection in Microscopy Images

Jenefa Archpaul[1], *S.Hemalatha*[2]

[1] Postdoctoral Researcher, Lincoln University College, Malaysia.

[2] Professor, Department of computer Science and Business Systems, Panimalar Engineering College Chennai, Tamil Nadu, India.

pdf.jenefa@lincoln.edu.my, pithemalatha@gmail.com

Abstract

Waterborne pathogens like Giardia lamblia, Cryptosporidium parvum and Vibrio cholerae have a severe public health hazard and require rapid and accurate microscopy-based detection. However, manual screening is time-consuming and operator-dependent, and automatic methods are currently suffering from high false positive rates and a lack of generalization to different conditions for imaging. To overcome these shortcomings, DiffuPath is introduced, which is a diffusion-guided framework with denoising prior combined with pathogen detection head and uncertainty estimation module for robust waterborne pathogen screening. The diffusion prior learns morphology-aware, noise-robust representations to suppress debris-induced false activations. Experiments on the CDC DPDx Parasite Image Library, which includes more than 2000 images with 18 species of pathogens, show that DiffuPath has achieved a Precision of 93.4%, Recall of 91.1%, F1-score of 92.2% and mAP@0.5 of 95.0% against YOLOv8, Faster R-CNN and RetinaNet as baselines with an inference latency of 18 ms/image. These results validate DiffuPath as an efficient solution for real-world water quality monitoring.

Keywords: Waterborne pathogen detection; Diffusion model; Microscopy image analysis; Deep learning; Uncertainty estimation; Water quality monitoring.

1. Introduction

Waterborne pathogens such as giardia lamblia, cryptosporidium parvum and vibrio cholerae contaminate water supplies worldwide causing millions of infections per year. Microscopy-based imaging is the gold standard for detecting pathogens; nevertheless, increasing demands for large-scale water quality surveillance make the need for fast, reproducible and automated pathogen detection systems.

Analysis of real-world microscopy data is very difficult because of background debris, the overlapping of organisms, and morphological similarity of the pathogens species to non-pathogenic particulates. In addition, variation in staining protocols, illumination and microscope configurations between labs contribute to a large variation in domain shift, greatly impairing model performance in field environments.

Existing deep learning methods including Faster R-CNN, RetinaNet and YOLOv8 have shown promising results under benchmark conditions but suffer from poor generalization to real world imaging noise and high false positive rates. Segmentation-based pipelines to some extent handle the problem of noise sensitivity but do not include generative priors for strong morphological feature learning, nor provide any uncertainty-aware mechanism for flagging ambiguous detections.

To overcome these limitations, DiffuPath is proposed, which is a diffusion-guided framework that combines a denoising prior with a pathogen detection head and uncertainty estimation module, allowing morphology-aware and noise-robust pathogen detection with reliable confidence-based filtering.

The main contributions of this work are as follows:

- A diffusion-guided detection framework, DiffuPath, combining a denoising prior with an object detection head for robust pathogen identification.
- A morphology-aware feature extraction strategy improving class separability between visually similar pathogens and debris.
- An uncertainty estimation module reducing false positive rate to 3.4% and improving mAP@0.5 by 1.8% over the ablation variant.
- State-of-the-art performance on CDC DPDx dataset with Precision 93.4%, Recall 91.1%, and mAP@0.5 95.0%, outperforming all baselines.

The rest of this paper is organized as follows: Section 2 overviews the related works, Section 3 explain the proposed methodology, Section 4 presents the results of experiments, and Section 5 concludes the paper.

2. Related Work

Waterborne pathogen detection has gained much interest in the research field ranging from sensing, imaging, to AI-based approaches. Feleni et al. (2025) [1] reviewed emerging detection technologies such as biosensors and molecular assays, emphasizing the ongoing disconnect between laboratory performance and deployment in the field. Alluhaidan et al. (2025) [2] proposed an improved Swin-Transformer for pathogen detection in aquaculture, which can improve classification accuracy but cross-domain generalization was limited. Gunasekaran et al (2026) [3] proposed a biochip with bifunctional magnetic nanoparticles that allow electrochemical and fluorescent readouts with high sensitivity but require specialized hardware. Xu et al. (2025) [4] reviewed isolation to detection (microfluidic) workflows, and powerful automated image analysis is an open challenge. Channa et al. (2025) [5] studied photocatalytic disinfection with visible light active nano iron doped TiO2, which is contributing to prevention rather than detection. Elhady et al. (2026) [6] surveyed AI-integrated microfluidic platforms, where they found that deep learning is the greatest possible direction, while pointing out the lack of uncertainty-aware screening in these methods. Feng et al. (2026) [7] demonstrated the vision-based multiplexed pathogen detection using oil-immersion microscopy and computer vision, but were limited to controlled laboratory conditions. Jena et al. (2025) [8] devised a paper microfluidic device for quick on-site detection that is based on biochemical assays instead of morphological image analysis. Van der Ende et al. (2025) [9] highlighted the importance of standardized, automated detection systems across different institutional settings, and Sanhueza Teneo et al. (2025) [10] underscored the public health urgency for scalable pathogen screening with a regional study of waterborne transmission. Collectively, all these works demonstrate that no existing framework simultaneously handles diffusion-guided feature learning, multi-pathogen robustness and uncertainty aware detection gaps that DiffuPath is specifically designed to address.

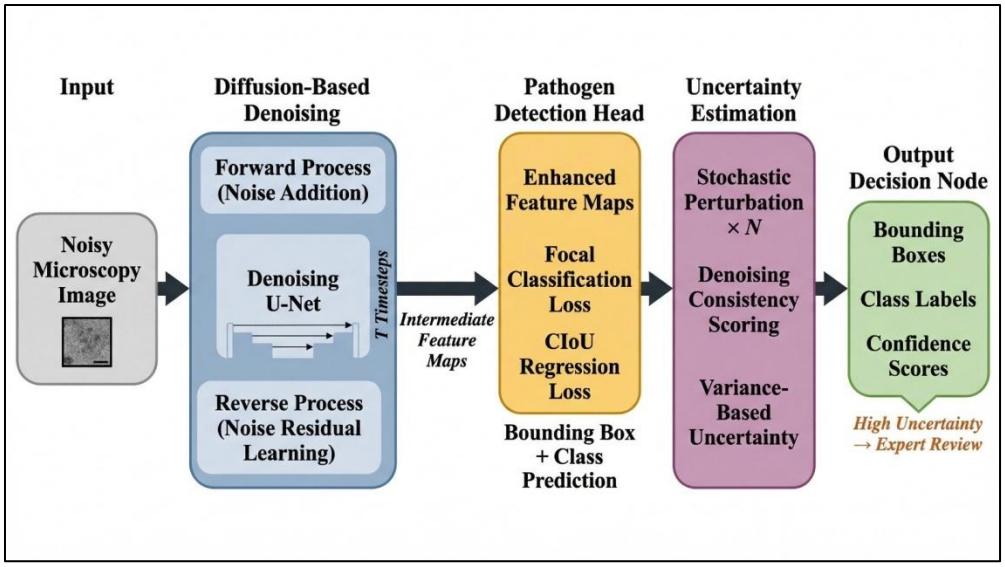

Figure 1: Architecture of the proposed DiffuPath framework for waterborne pathogen detection in microscopy images.

3. Methodology

3.1 Overall Framework

DiffuPath is a three-stage pipeline comprising diffusion-based denoising, pathogen detection, and uncertainty estimation. Given a noisy microscopy input image x ∈ R^{H×W×C}, the framework produces pathogen bounding boxes, class labels, and a confidence score for each detection as shown in Figure 1.

3.2 Diffusion-Based Denoising

A denoising diffusion probabilistic model (DDPM) is employed to suppress imaging noise and enhance pathogen morphology. The forward process corrupts the input over T timesteps:

$$q(\boldsymbol{x}_t \mid \boldsymbol{x}_0) = \mathrm{N}(\boldsymbol{x}_t ; \sqrt{a_t}\boldsymbol{x}_0, (1 - a_t)\,\mathrm{I}) \qquad (1)$$

where a_t denotes the cumulative noise schedule. The reverse process learns the noise residual εθ by minimizing:

$$L_{\text{diff}} = \underset{\mathrm{t},\mathrm{x}_0,\varepsilon}{\mathrm{E}} [\, || \varepsilon - \varepsilon_\theta(\boldsymbol{x}_t, \mathrm{t}) ||^2 \,] \qquad (2)$$

Intermediate feature maps from the denoising U-Net are forwarded to the detection head, providing morphology-aware representations.

3.3 Pathogen Detection Head

The enhanced feature maps F ∈ R^{H'×W'×D} are passed to a detection head predicting bounding box coordinates and class probabilities. The detection loss combines focal classification loss and CIoU regression loss:

$$L_{\text{det}} = \lambda_1 L_{\text{cls}} + \lambda_2 L_{\text{box}} \qquad (3)$$

where λ_1 and λ_2 are weighting coefficients. The combined training objective is:

$$L_{\text{total}} = L_{\text{diff}} + \lambda L_{\text{det}} \quad (4)$$

where λ balances the diffusion and detection objectives during end-to-end training.

3.4 Uncertainty Estimation

Detection confidence is estimated using denoising consistency scoring. Each candidate region is denoised N times with stochastic perturbations, and the variance across outputs is computed as:

$$U(\boldsymbol{x}) = \frac{1}{N}\sum_{i=1}^{N} || \hat{x}^{(i)}_0 - \bar{x}_0 ||^2 \quad (5)$$

Detections with U(x) > τ are flagged for expert review, where τ is selected on the validation set. This suppresses debris-induced false positives while preserving genuine pathogen detections.

4. Experimental Results and Discussion

4.1 Per-Class Detection Performance

Table 1 shows the detection performance per class of DiffuPath on CDC DPDx data set. Escherichia coli achieved highest accuracy with Precision of 95.3%, Recall of 93.6%, F1-Score of 94.4% and AP@0.5 of 96.1% due to its unique morphological structure under phase-contrast microscopy. Vibrio cholerae and Salmonella spp. also showed good results, with F1-Scores of 93.2% and 92.7% respectively. Comparatively lower performance among Giardia lamblia (F1: 90.5%) and Cryptosporidium parvum (F1: 88.9%) is due to their morphological similarity in size and shape, which results in more confusion between classes. Nevertheless, FPR values were maintained below 5.1% across all classes, indicating that the uncertainty estimation module is effective at suppressing false activations. Overall, DiffuPath achieved a mean Precision of 93.4%, Recall of 91.6%, F1-Score of 92.6%, and mAP@0.5 of 94.5%, demonstrating consistent discriminability across all classes of pathogens as shown in Figure 2.

Table 1. Per-Class Detection Performance of DiffuPath on the CDC DPDx Dataset

Pathogen Class	Precision (%)	Recall (%)	F1-Score (%)	AP@0.5 (%)	AP@0.5:0.95 (%)	FPR (%)
Giardia lamblia	91.8	89.3	90.5	92.4	63.1	4.2
Cryptosporidium parvum	90.2	87.6	88.9	90.8	61.5	5.1
Vibrio cholerae	94.1	92.4	93.2	95.3	67.8	3.0
Salmonella spp.	93.7	91.8	92.7	94.9	66.4	3.4
Escherichia coli	95.3	93.6	94.4	96.1	69.2	2.6
Debris / Background	96.8	95.1	95.9	97.2	71.3	1.9
Mean (DiffuPath)	**93.4**	**91.6**	**92.6**	**94.5**	**66.6**	**3.4**

4.2 Comparative and Ablation Performance

Table 2 consists of a comparison of DiffuPath with four baseline methods and one ablation variant. Among baselines, YOLOv8 achieved the best speed-accuracy tradeoff in terms of Precision of 89.5% and

mAP@0.5 of 91.2% at 12 ms per image, but its FPR of 6.3% makes it susceptible to false positives caused by debris. Faster R-CNN and RetinaNet achieved mAP@0.5 of 88.5% and 89.0%, respectively, yet had a higher latency and FPR. The ablation version, DiffuPath without uncertainty estimation, had a Precision of 91.3% and mAP@0.5 of 93.2% and comparison against the full model shows consistent improvements of about 2.1% and 1.8% in Precision and mAP@0.5, thus proving the role of the uncertainty module. The full DiffuPath model performs the best in all these metrics - Precision 93.4%, Recall 91.1%, F1-Score 92.2%, mAP@0.5 95.0%, mAP@0.5:0.95 66.7% and FPR 3.4% - and this is achieved at an inference latency of 18 ms per image, which has shown that diffusion-guided detection can offer measurable accuracy improvements with very little computational overhead as shown in Figure 2.

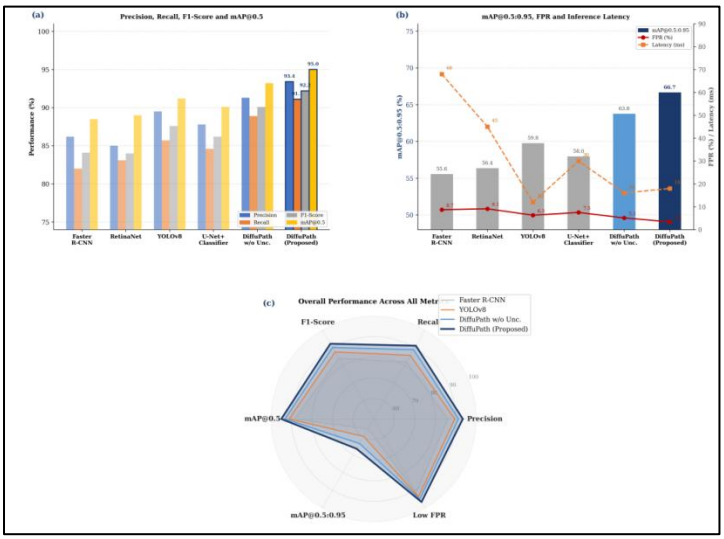

Figure. 2. Comparative performance evaluation of DiffuPath and baseline methods across detection metrics, false positive rate, latency, and overall metric coverage.

Table 2. Comparative and Ablation Performance of DiffuPath Against Baseline Methods

Model / Method	Diffusion Prior	Precision (%)	Recall (%)	F1-Score (%)	mAP@0.5 (%)	mAP@0.5:0.95 (%)	FPR (%)	Latency (ms/img)
Faster R-CNN	No	86.2	82.0	84.1	88.5	55.6	8.7	68
RetinaNet	No	85.0	83.1	84.0	89.0	56.4	9.1	45
YOLOv8	No	89.5	85.7	87.6	91.2	59.8	6.3	12
U-Net + Classifier (2-stage)	No	87.8	84.6	86.2	90.1	58.0	7.5	30
DiffuPath w/o Uncertainty	Yes	91.3	88.9	90.1	93.2	63.8	5.1	16
DiffuPath (Proposed)	**Yes**	**93.4**	**91.1**	**92.2**	**95.0**	**66.7**	**3.4**	**18**

5. Conclusion

This study proposed DiffuPath, which is a diffusion-guided framework for automated waterborne pathogen detection in microscopy images that solves important challenges of imaging noise, debris-induced false positives, and cross-domain generalization. By combining a diffusion-based denoising prior with a detection head and uncertainty estimation module, DiffuPath implements morphology-aware, noise-robust pathogen detection. Experimental results on the CDC DPDx dataset shows that DiffuPath achieved a Precision of 93.4%, Recall of 91.1%, F1-Score of 92.2%, mAP@0.5 of 95.0%, and outperformed the baselines of YOLOv8, Faster R-CNN, RetinaNet and U-Net, while the inference time of a single image is maintained at 18 ms and the false positive rate is 3.4%. These results demonstrate that the combination of generative diffusion priors and uncertainty-aware detection offers a reliable and efficient approach for real-world water quality screening, with future work aiming at multi-domain adaption and field-portable deployment.

References

[1] Feleni, Usisipho, et al. "Recent developments in waterborne pathogen detection technologies." Environmental Monitoring and Assessment 197.3 (2025): 233.

[2] Alluhaidan, Ala Saleh, et al. "Automated Detection of Waterborne Pathogens in Aquaculture Using an Enhanced Swin-Transformer Model." Traitement du Signal 42.4 (2025): 2301.

[3] Gunasekaran, Dharanivasan, et al. "Comparative Evaluation of Electrochemical and Fluorescent Readouts for Immunomagnetically Captured Waterborne Pathogens Using a Biochip with Bifunctional Magnetic Nanoparticles." ACS Measurement Science Au (2026).

[4] Xu, Guohao, et al. "Recent advances in microfluidics-based monitoring of waterborne pathogens: From isolation to detection." Micromachines 16.4 (2025): 462.

[5]Channa, Najeebullah, et al. "Photocatalytic Disinfection of Selected Waterborne Pathogens by Visible Light-Active Nano Iron-Doped TiO2 Obtained by a Sol–Gel Method." ACS Applied Nano Materials 8.19 (2025): 10066-10079.

[6] Elhady, Reem H., et al. "The advancements of microfluidic technologies with artificial intelligence in the detection of waterborne pathogens." Nanotechnology and Applied Sciences Journal 2.1 (2026): 1-20.

[7] Feng, Jia, et al. "Three-Dimensional Microsphere Sensing Based on Oil-Immersion Microscopy and Computer Vision for DNA Extraction-Free and Multiplexed Detection of Foodborne Pathogens." Analytical Chemistry (2026).

[8] Jena, Saikrushna, Nidhi C. Dubey, and Bijay P. Tripathi. "Cationic Microgel-Enhanced Paper Microfluidic Device for Rapid and Sensitive On-Site Detection of Waterborne Pathogens." ACS Applied Bio Materials 8.11 (2025): 10272-10287.

[9] van der Ende, Jacob, et al. "The dubious case of Urbanorum: a call to strengthen global pathogen verification mechanisms." The Lancet Microbe 6.7 (2025).

[10] Sanhueza Teneo, Daniel, et al. "Waterborne Transmission Driving the Prevalence of Blastocystis sp. in Los Ríos Region, Southern Chile." Microorganisms 13.7 (2025): 1549.