Biolens: AI-Driven Diatom Analysis for Environmental Monitoring and Risk Detection

Mr. Pragadheesh Thirumal M[1], Akiladithya R[2], Dhivakar B S[3], Sumitha I[4], Mahalakshmi J[5]
Assistant Professor, Department of Computer Science and Engineering
Student, Department of Computer Science and Engineering

***Abstract*—Aquatic ecosystems face increasing threats from eutrophication, pollution, and oxygen depletion, which affect biodiversity, water security, and public health. Traditional monitoring methods are often reactive, detecting environmental degradation only after damage occurs, while manual microscopic analysis of bioindicators is resource-intensive and struggles with highly clustered samples. This paper proposes BioLens, an advanced AI-driven environmental monitoring framework that integrates digital microscopy, a two-stage deep learning pipeline, and automated ecological scoring for early risk detection. To overcome the challenge of overlapping microorganisms in real-world samples, the system utilizes StarDist, an instance segmentation model, to precisely extract isolated diatoms.**

These extractions are subsequently processed by a YOLO (You Only Look Once) classification engine trained to identify 81 distinct diatom genera with a validation accuracy of 98.12%. The AI predictions are then mathematically mapped to a 1-to-5 biological tolerance scale to compute a quantitative Water Quality Index (WQI). BioLens provides scalable, real-time ecosystem assessment, shifting aquatic management from reactive response to proactive environmental risk mitigation.

I. INTRODUCTION

Aquatic ecosystems worldwide are increasingly threatened by pollution, environmental degradation, and water-related disasters that impact both biodiversity and public health. In this context, proactive and timely environmental monitoring is essential to prevent large-scale ecological damage. Among various biological indicators, diatoms are highly effective due to their rapid response to changes in water chemistry, nutrient levels, and pollutant concentration, making them suitable for early detection of ecosystem disturbances before visible damage occurs. However, conventional diatom analysis relies heavily on manual microscopy and expert-driven workflows, which limit scalability and delay ecological assessment.

To address the need for intelligent and scalable monitoring, BioLens introduces an AI-driven framework for automated diatom extraction, classification, and ecological scoring, positioned within the domain of Computer Vision, Deep Learning, and Intelligent Environmental Monitoring Systems. The system employs a two-stage deep learning pipeline, where StarDist performs precise instance segmentation of overlapping diatoms, followed by a customized YOLO-based classifier for rapid recognition of 81 diatom species. The classified species are further mapped to a 1–5 biological tolerance scale to compute a Water Quality Index (WQI), enabling reliable ecosystem assessment and supporting early warning-based environmental risk mitigation.

II. RELATED WORK

Research in automated microscopic image analysis and aquatic ecosystem monitoring has expanded with the integration of machine learning and deep learning techniques. Baharin et al. [1] proposed an image enhancement technique called Integrated Automatic Background Removal (IABR) to improve the quality of low-resolution diatom microscopic images. Their approach focuses on preprocessing and background removal to enhance image clarity for more accurate analysis of diatom structures.

Israde-Alcántara [2] introduced a geometric approach for diatom identification using binary masks of concentric rings. This method analyzes structural patterns and shape characteristics in microscopic images to approximate diatom classification. However, the approach relies heavily on handcrafted features and may face limitations when dealing with highly diverse diatom species.

Machine learning has also been applied to broader aquatic ecosystem monitoring. Tang et al. [3] developed a spatio-temporal deep learning model for predicting algal blooms in Lake Okeechobee using multi-source environmental data, demonstrating the effectiveness of deep learning in ecological prediction. Similarly, Fernández-Fernández et al. [4] used Long Short-Term Memory (LSTM) networks to forecast cyanobacterial blooms. Their study showed that LSTM models can provide reliable predictions even with incomplete spatio-temporal data.

For biological image classification, Gunduz and Gunal [5] proposed a lightweight convolutional neural network that achieves accurate classification while maintaining computational efficiency, making it suitable for resource-constrained environments.

While prior studies have contributed substantially to image preprocessing, ecological forecasting, and deep learning-driven classification, the majority of existing approaches focus on individual aspects rather than providing an end-to-end solution. To bridge this gap, the proposed BioLens system integrates deep learning-based diatom classification with a scalable automated microscopic image analysis framework, enabling efficient and reliable environmental monitoring.

III. PROBLEM STATEMENT

Despite the importance of diatom-based ecosystem monitoring, existing assessment methods remain predominantly reactive and operationally inefficient. Current workflows depend on manual identification and counting of diatoms, making the process slow, labor-intensive, and prone to inconsistency. The challenge becomes more severe in microscopic samples containing densely clustered organisms, overlapping structures, and background debris, where accurate extraction and species recognition are difficult to achieve. These limitations prevent existing systems from supporting timely, large-scale, and real-time environmental surveillance, creating a significant gap between the need for early aquatic disaster prediction and the capability of present monitoring techniques.

IV. PROPOSED SYSTEM

The proposed BioLens system is an intelligent environmental monitoring platform designed for the early detection of ecological disturbances in aquatic ecosystems through microscopic diatom analysis and a specialized deep learning pipeline. Unlike conventional monitoring workflows that depend on periodic manual sampling and delayed laboratory interpretation, BioLens offers a scalable, automated, and real-time solution by integrating digital microscopy, mobile-based data acquisition, and artificial intelligence. The architecture establishes a unified framework connecting field researchers with centralized machine learning services, enabling efficient image acquisition, automated species-level classification, and quantitative ecosystem risk assessment.

A. Microscopic Image Acquisition and Data Upload

The initial stage of the framework focuses on the acquisition of microscopic diatom images from collected water samples. Field researchers or environmental monitoring personnel capture images using a digital microscope connected to a mobile device. The BioLens mobile interface provides a secure and user-friendly environment for uploading crowded microscopic images along with optional metadata such as sampling location, timestamp, and environmental conditions. This process facilitates the creation of a structured real-world dataset for downstream analysis. The uploaded images are securely transmitted to the centralized backend server, where they are queued for automated processing.

B. Instance Segmentation and Image Preprocessing

Microscopic water samples often contain densely clustered diatoms, overlapping organisms, and significant background debris, which complicate direct species recognition. To address this challenge, the uploaded images are processed through a StarDist-based instance segmentation module. This stage accurately detects organism boundaries, separates overlapping structures, and extracts clean diatom instances from noisy backgrounds. By generating isolated organism crops, the segmentation phase significantly enhances the reliability and precision of the subsequent classification stage.

C. High-Speed YOLO-Based Diatom Classification

The extracted diatom instances are then forwarded to a customized YOLO-based deep learning classifier, which forms the core of the BioLens identification pipeline. YOLO is adopted due to its fast inference capability and strong classification accuracy compared to conventional CNN-based models. The trained architecture performs species-level prediction across 81 distinct diatom taxonomic classes and outputs both the predicted label and an associated confidence score. This enables rapid and scalable microscopic species recognition suitable for real-time environmental monitoring.

D. Ecological Scoring and WQI Risk Assessment

To transform species predictions into meaningful ecological intelligence, the backend server maps each identified diatom class to a Master Tolerance Dictionary, where species are assigned a biological tolerance score ranging from 1 (clean-water sensitive) to 5 (pollution tolerant). Using the aggregated species counts and tolerance weights, the system computes a Water Quality Index (WQI) that reflects the ecological health of the monitored water body. Based on the computed WQI range, the system generates risk-level alerts such as Safe, Caution, and Unsafe, which are transmitted to the user interface to support timely preventive actions and proactive environmental risk mitigation.

V. SYSTEM MODULES

1. Image Acquisition Module This module is responsible for capturing microscopic images of water samples containing diatoms. A digital microscope connected to a mobile device is used to collect high-resolution images of dense aquatic samples. The BioLens mobile application allows field researchers to securely upload these images along with optional metadata such as sampling location, date, and environmental conditions. This module ensures standardized data collection and facilitates structured dataset creation for automated server-side analysis.

2. Instance Segmentation Module Real-world microscopic images often contain dense, overlapping clusters of diatoms and background debris that confuse standard image processing algorithms. To resolve this, this module replaces traditional preprocessing with an advanced instance segmentation model powered by StarDist. The model identifies the morphological boundaries of complex, touching organisms and automatically performs pixel-perfect crops. This isolates each diatom from the noisy background, ensuring the classification engine receives clean, individual inputs.

3. Feature Extraction Module In this module, the deep learning backbone of the YOLO architecture automatically extracts critical morphological features from the segmented diatom crops. The convolutional layers analyze complex structural elements such as diatom shape, bilateral or radial symmetry, texture patterns, and distinct frustule structures. This automated feature extraction is highly optimized for speed and completely eliminates the need for manual feature engineering.

4. Diatom Classification Module The classification module utilizes a highly customized YOLO (You Only Look Once) deep learning architecture, chosen over baseline models like EfficientNet for its superior inference speed and Top-1 accuracy. The model analyzes the extracted features and predicts the exact diatom genus across 81 distinct taxonomic classes. Alongside the predicted class, the model generates a mathematical confidence score for each prediction, enabling rapid, high-throughput identification of microscopic species.

5. Ecological Scoring Module
Replacing traditional manual counting, this module translates the AI classification results into standardized ecological metrics. It utilizes a comprehensive Master Tolerance Dictionary that maps all 81 detected diatom genera to a biological tolerance scale ranging from 1 (Very Sensitive) to 5 (Highly Tolerant). By tallying the specific species detected in the image and applying these tolerance weights, the module calculates an aggregate score for the entire sample.

6. Risk Detection Module This module evaluates the immediate ecological risk using the data generated by the scoring module. The system calculates a final, quantitative Water Quality Index (WQI) by dividing the weighted tolerance scores by the total number of diatoms analyzed. The system continuously evaluates this WQI against strict environmental thresholds. If the index crosses into moderate or severe pollution ranges, the system automatically generates early

warning alerts indicating immediate environmental toxicity risks.

VI. TECHNOLOGY STACK

A. Programming Languages

Python serves as the primary programming language for implementing the deep learning pipeline, backend services, and Water Quality Index (WQI) computation, owing to its rich scientific computing and machine learning ecosystem. JavaScript and TypeScript are employed for developing the cross-platform mobile application and interactive web frontend, ensuring scalable and responsive user interaction.

B. Deep Learning Frameworks

The Ultralytics framework is adopted for training, optimization, and deployment of the YOLO-based diatom classification model. For the segmentation stage, the StarDist biological imaging library is integrated on top of TensorFlow or PyTorch backends, enabling efficient tensor computation, model training, and GPU-accelerated inference.

C. Image Processing and Segmentation Libraries

Conventional image processing libraries such as OpenCV and PIL are utilized for image formatting, resizing, and preprocessing tasks. In addition, specialized biological instance segmentation libraries are employed to accurately detect and isolate star-convex overlapping diatom structures from noisy microscopic backgrounds, significantly improving downstream classification accuracy.

D. Model Architecture

The proposed system employs a specialized two-stage deep learning architecture rather than a conventional single-stage CNN baseline. The first stage utilizes StarDist (Star-Convex Object Detection) as the segmentation engine to isolate touching and overlapping microorganisms. The extracted organism instances are then processed by a custom-trained YOLO architecture, which serves as the high-speed species classification engine across 81 distinct taxonomic classes.

E. Mobile Application for Data Acquisition

The mobile interface for field researchers is developed using React Native, enabling seamless cross-platform deployment across iOS and Android devices. The application facilitates secure upload of raw crowded microscopic images along with associated environmental metadata, ensuring efficient field-level data acquisition.

F. Backend Framework

The backend infrastructure is implemented using FastAPI or Flask, exposing a RESTful API layer for communication between the mobile interface, web dashboard, and AI inference services. The backend manages image routing through the segmentation and classification pipelines, performs WQI computation, and returns the resulting ecological risk metrics to the frontend systems.

G. Database

A NoSQL database architecture, such as MongoDB or Firebase, is used to store raw microscopic images, AI-generated classification outputs, and associated metadata including sampling location and timestamp. This supports structured long-term temporal ecological analysis and scalable environmental data management.

VII. SYSTEM ARCHITECTURE

A. Data Acquisition Layer This layer is responsible for collecting the primary environmental data required for the system. The process begins with the physical gathering of water samples, which are subsequently examined using a digital microscope. The visual data captured during this examination results in high-resolution microscopic images. Finally, a dedicated mobile application is utilized by the user to capture these images along with relevant contextual information, initiating the upload process to the main system.

B. Upload via Mobile App This specific process acts as the vital communication bridge connecting the physical data acquisition layer to the analytical processing layer. Through the mobile application, the core microscopic image data is securely transmitted to the backend server. Alongside the image itself, the app automatically packages and uploads essential metadata, specifically the geographical location coordinates and the exact time stamp of when the data collection occurred.

C. Backend Processing Layer (Backend Server) Operating as the core server-side component, this layer handles all complex data analysis and processing through a highly organized, sequential workflow. Initially, the uploaded microscopic images bypass standard preprocessing and undergo advanced instance segmentation using StarDist to isolate and cleanly extract individual diatoms from crowded, overlapping backgrounds. These extracted crops are then fed directly into a customized YOLO (You Only Look Once) deep learning model. This high-speed engine is responsible for generating predictions across 81 distinct diatom genera alongside confidence scores that calculate the certainty of those predictions. As a final step in this layer, the system maps these predictions to a Master Tolerance Dictionary to calculate a quantitative Water Quality Index (WQI).

D. Risk Detection & Alert Generation Following the analytical processing sequence, this critical stage utilizes the conclusive WQI outputs to conduct a thorough potential risk assessment. If the system detects significant environmental threats, ecological imbalances, or harmful anomalies based on the calculated pollution tolerance thresholds, it automatically generates targeted alerts (e.g., Safe, Caution, or Unsafe) to warn users or relevant stakeholders of potential danger.

E. Application Layer Serving as the primary user-facing interface, this final layer allows individuals to seamlessly visualize and interact with the complex results generated by the backend server. It provides comprehensive, interactive web-based dashboards that deliver a clear, broad overview of all processed system data. Through these interfaces, users can visually explore species distribution and water health via advanced interactive charts and dynamic visual analytics. Additionally, any critical risk alerts generated during the previous processing stage are prominently displayed within this layer to ensure immediate user awareness.

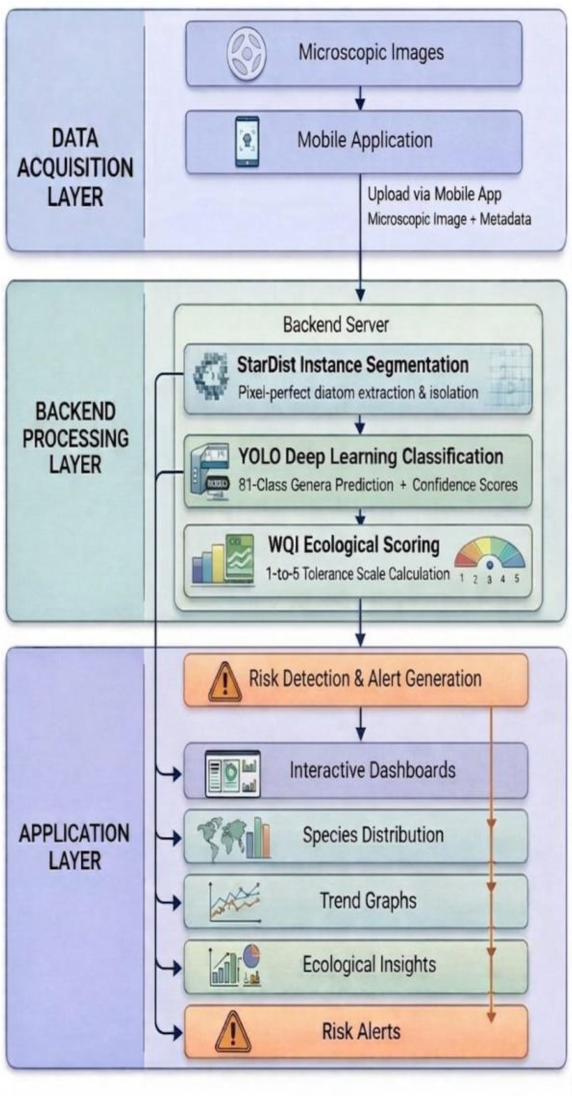

(System Architecture Diagram)

VIII. ADVANTAGES OF PROPOSED SYSTEM

A. Precision in Complex Microscopic Samples
Real-world aquatic samples often contain overlapping, touching diatoms and significant background noise, which severely affects conventional analysis pipelines. By integrating the StarDist instance segmentation model, the proposed system achieves highly precise isolation of individual diatom instances from complex microscopic backgrounds. This accurate extraction minimizes background interference and ensures that the downstream classification stage receives clean, isolated organism data, significantly improving reliability.

B. High-Speed Multi-Class Classification

The customized YOLO-based classification architecture provides robust feature extraction combined with rapid inference performance. Unlike traditional baseline classifiers, the proposed framework efficiently distinguishes among 81 distinct diatom genera, making it highly suitable for large-scale and real-time environmental monitoring applications.

C. Quantitative Ecological Scoring

Beyond species-level recognition, BioLens transforms AI predictions into actionable ecological intelligence. By mapping the identified 81 diatom species to a 1–5 biological tolerance scale, the framework computes a Water Quality Index (WQI) that quantitatively evaluates the severity of aquatic ecosystem pollution. This integration extends the system beyond conventional computer vision tasks toward environmental decision support.

D. Transfer Learning Efficiency

The adoption of transfer learning enables the YOLO classifier to leverage knowledge from pre-trained feature representations, thereby reducing training time and improving classification accuracy. This is particularly beneficial in biological image analysis, where collecting and annotating large-scale labeled datasets is often difficult and resource-intensive.

E. Improved Generalization and High Accuracy

To improve model robustness, the system employs data augmentation techniques that introduce variations in image orientation, scale, and structural appearance. Combined with adaptive learning rate scheduling, this strategy enables the model to achieve a validation accuracy of 98.12%. The minimal difference between training and validation performance further indicates strong generalization capability with reduced overfitting.

F. Advanced Visual Analytics and Confidence-Aware Monitoring

In addition to species predictions, the system outputs confidence scores for each classification result, enabling reliability-aware analysis. These outputs are integrated into interactive visual dashboards, including Sankey diagrams, sunburst charts, and WQI gauge indicators, allowing researchers to intuitively analyze ecological flows, species distributions, and environmental safety trends in real time.

IX. FUTURE WORK

Geospatial Ecological Mapping: The system can be extended to support geospatial mapping by integrating Geographic Information Systems (GIS). By linking the calculated Water Quality Index (WQI) and specific diatom classifications with geographic sampling coordinates, researchers can visualize ecological health across entire water bodies. Such mapping would assist in identifying precise pollution hotspots and supporting data-driven environmental management.

High-Throughput Batch Processing: To improve efficiency in large-scale studies, the platform can be upgraded to support parallel batch image classification. This feature would allow users to upload hundreds of complex microscopic images simultaneously. By leveraging cloud GPUs, the StarDist and YOLO pipeline could process these batches concurrently, significantly reducing analysis time for large ecological monitoring programs.

Explainable Artificial Intelligence (XAI): Integrating explainable AI techniques can improve the transparency and reliability of the YOLO classification results. Visualization methods such as Grad-CAM or saliency maps can be implemented to highlight the specific morphological regions (e.g., frustule patterns or symmetry) that influenced the model’s prediction. This capability helps researchers understand the AI’s reasoning and verify its taxonomic accuracy.

Continuous Active Learning (Human-in-the-Loop): Future iterations could introduce an active learning feedback loop. If the model is uncertain about a highly mutated or rare diatom species, the image can be flagged for review by a human taxonomist. The expert’s correction would then be fed back into the YOLO training dataset, allowing the BioLens model to continuously learn and improve its 81-class accuracy over time.

Automated Report Generation: The platform can include an automated reporting module that enables users to export classification results, dashboard visualizations, and WQI insights in formats such as PDF or CSV. These reports can summarize species predictions and dataset statistics, making it easier to document findings and share results with environmental agencies.

Research Collaboration and Administrative Tools: Administrative tools can be added to provide system-level analytics, including classification frequency, dataset usage, and geographic distribution. Collaboration features would allow multiple researchers across different laboratories to share datasets, annotations, and WQI outcomes, promoting cooperative research.

Scalable Cloud-Based Deployment and API Security: Deploying the backend architecture on cloud platforms such as AWS, Microsoft Azure, or Google Cloud can dramatically improve scalability. To maintain secure access to the BioLens engine, the system will implement robust authentication mechanisms and API rate limiting. These measures protect sensitive research data and ensure stable inference performance during high-traffic periods.

IoT Water Sensor Integration: Future enhancements will include integration with IoT-based water quality hardware sensors that monitor real-time physicochemical parameters such as pH, turbidity, dissolved oxygen, and temperature. Fusing this live sensor data with the biological WQI scores generated by the diatom classifications will provide a complete, holistic understanding of aquatic ecosystem health.

X. RESULT AND DISCUSSION

The proposed BioLens system was evaluated on a curated and highly complex microscopic aquatic image dataset comprising 81 distinct diatom genera. Experimental findings demonstrate that the proposed two-stage deep learning integration of StarDist-based instance segmentation and customized YOLO-based classification effectively addresses the traditional challenges posed by overlapping organisms and background debris. By leveraging transfer learning, the YOLO classifier utilized pre-trained feature representations, significantly reducing training time while maximizing species-level classification performance. This strategy proved especially beneficial in biological image analysis, where acquiring large-scale, perfectly isolated labeled datasets is inherently difficult.

To improve robustness against the variability present in real-world aquatic samples, the training pipeline incorporated spatial and pixel-level data augmentation techniques. These transformations exposed the model to variations in orientation, scale, and illumination, thereby improving its ability to generalize across unseen and densely populated microscopic slides. Furthermore, the use of adaptive learning rate scheduling ensured stable convergence and efficient optimization throughout training. The minimal performance gap between training and validation phases indicates strong feature generalization and reduced risk of overfitting, confirming that the model learned meaningful morphological representations of diatom structures.

The final classification framework achieved a validation accuracy of 98.12%, demonstrating strong predictive performance on unseen test samples. More importantly, the significance of the BioLens framework extends beyond image recognition into actionable ecological intelligence. By mapping the highly accurate YOLO predictions and associated confidence scores to a 1–5 biological tolerance scale, the system successfully generated a reliable Water Quality Index (WQI) for ecosystem risk assessment. The integration of these outputs into interactive hierarchical dashboards further validates the practicality of the system for proactive environmental monitoring, ecological research, and biological risk detection.

XI. CONCLUSION:

BioLens demonstrates the transformative potential of integrating accessible data acquisition with state-of-the-art deep learning to modernize environmental monitoring and automate complex biological classification. By overcoming the limitations of traditional manual microscopy through a robust two-stage AI pipeline—utilizing StarDist for precise instance segmentation and a custom YOLO architecture for high-speed, 81-class taxonomic identification—the system establishes a new standard for processing densely populated, real-world aquatic samples. Furthermore, by mathematically translating these automated AI predictions into a standardized Water Quality Index (WQI) and presenting the results through professional-grade, interactive hierarchical dashboards, BioLens transcends basic computer vision to become a fully realized environmental decision-support tool. Overall, the project highlights the critical role of advanced computational engineering in environmental science, seamlessly connecting deep learning, ecological mathematics, and interactive data analytics to proactively support sustainable water resource management and global biodiversity conservation.

REFERENCES

[1] Mohd Aiman Syahmi Kamarul Baharin, Ahmad Shahrizan Abdul Ghani, Syafiq Qhushairy Syamsul Amri, Normawaty Mohammad-Noor, Hasnun Nita Ismail, "Enhancement of Low-Quality Diatom Images using Integrated Automatic Background Removal (IABR) Method from Digital Microscopic Image," 2023.
[2] Isabel Israde-Alcántara, "Binary Masks of Concentric Rings as a Method to Approximate Identification of Diatoms Using Images," 2020.
[3] Tang, Y., Feng, Y., Fung, S., Ruiz Xomchuk, V., Jiang, M., Moore, T., Beckler, J., "Spatio-Temporal Deep Learning-based Algal Bloom Prediction for Lake Okeechobee Using Multi-Source Data Fusion," 2022.
[4] Raúl Fernández-Fernández, Eva Besada-Portas, Antonio Quesada, José A. López-Orozco, Samuel Cirés, "LSTM Networks Provide Efficient Cyanobacterial Blooms Forecasting Even with Incomplete Spatio-Temporal Data," 2024.
[5] Huseyin Gunduz, Serkan Gunal, "A Lightweight Convolutional Neural Network (CNN) Model for Diatom Classification: DiatomNet," 2024.