Grenze International Journal of Engineering and Technology, June Issue

An Edge-Optimized Deep Learning Framework for Real-Time Aquatic Pathology Detection and Automated Therapeutic Guidance

P A Muhammed Rishal[1], Sneha George[2] and T. Jemima Jebaseeli[3]
[1-2]Division of Computer Science and Engineering, Karunya Institute of Technology and Sciences, Coimbatore, India
Email: pamuhammed@karunya.edu.in, snehageorge@karunya.edu
[3]Department of Computer Science and Engineering, Vel Tech Rangarajan Dr. Sagunthala R&D Institute of Science and Technology, Chennai, India
Email: jemi.jeba@gmail.com

***Abstract*— The quick and precise detection of aquatic diseases is an essential component to support the globally thriving aquaculture sector, while at the same time providing security for aquarium inhabitants. Currently, the veterinary diagnosis of fish disease is performed manually, is quite subjective and not always available, especially in remote or poor environments. The last few applications have focused on the classification of fish disease using deep learning, but it relies heavily on time consuming cloud-dependent networks and is only able to classify the disease with a predicted class as the result. This reseach proposes AquaDiagnost, a highly optimized and end-to-end deep learning framework with built-in automatic treatment expert, focusing on edge devices. The technique show that with a novel two-stage transfer learning methodology trained over MobileNetV2 network, the framework can achieve up to 7 fresh-water fish diseases identification with only visual symptoms being analyzed, the weighted F1-score and accuracy is 0.90 and 90.2% with the dataset of 2,444 images respectively. A compact system with only 3.5M parameters, it achieves less than 0.8 seconds inference per image on the local devices. Moreover, prediction result from the proposed framework automatically invokes a diagnostic decision rule engine, and then immediate generate deterministic treatment protocol including specific dosages of medication, first step actions for quarantining affected fish, changes in physical and chemical water parameters. It turns out that AquaDiagnost presents an innovative approach for low-latency, privacy-protecting early detection of aquatic disease at the edge.**

***Index Terms*— Edge AI, MobileNetV2, Aquaculture Pathology, Expert Systems, Transfer Learning, Disease Classification, Offline Inference.**

I. Introduction

The vast aquarium fishkeeping and ornamental aquaculture industry is a booming global enterprise, with approximately 700 million ornamental fish maintained around the world [1]. Understanding and maintaining the health and robustness of aquatic ecosystems is a major concern, as the contained environments of aquaria render fish vulnerable to the rapid and often catastrophic onset of bacterial, fungal, parasitic, and viral diseases, most of which are manifested by clear dermatological and morphological anomalies [2].

Grenze ID: 01.GIJET.12.2.539_26

Currently, all diagnostic paradigms rely entirely on the direct human interpretation by experienced aquarists or aquatic veterinarians [3]. This approach is intrinsically subjective, completely dependent on scarce expert knowledge, and inaccessible to the average hobbyist fishkeeper or distant aquaculture farm. The proliferation of smart phone cameras offers a unique opportunity to democratize diagnosis but there remain significant challenges in bridging the gap between the raw image capture and the application of vet- prescribed interventions [4].

The recent deep learning paradigm has successfully tackled challenges such as medical image classification using Convolutional Neural Networks (CNNs), a methodology that is equally applicable to aquatic animal disease [5]. The common application in precision agriculture and aquaculture AI uses heavily parameterized architectures such as ResNet or VGG- 16, which rely on intensive computation and persistent internet connectivity, an impractical paradigm for rural aquaculture farms and amateur aquaria due to its large computational needs and demand for a stable connection [6]. Furthermore, the predominant literature focuses on classification tasks, neglecting a significant actionability gap where an organism is classified, but the owner is still unaware of appropriate therapy [7].

To overcome these limitations this research introduces AquaDiagnost, a novel offline-first diagnostic system where computational load has been shifted to the edge, the high-fidelity feature extraction has been achieved by transferring learning using the lightweight MobileNetV2, a minimal computational requirement and beyond a standard classification system the framework directly outputs a definite therapeutic prescription.

The main contributions of the research are as follows.

- Edge Optimized Disease Classification: Achieved high classification using a compressed CNN on a dataset of 2,444 Freshwater Fish Diseases, categorized into 7 classes.
- Two Stage transfer learning strategy: Designed and implemented a specialized fine-tuning process utilizing MobileNetV2 to extract pathological features effectively while avoiding feature redundancy.
- Automated Therapeutic expert system: Designed a comprehensive knowledge base and directly coupled with the CNN output. It translates the probable diagnosis to a detailed prescribed treatment with specific medication and water parameters.
- Offline-First web architecture: Built a completely localized, low-latency interface for diagnosis and therapeutic prescription using Flask.

II. Related Work

The use of Machine Learning (ML) and computer vision in precision aquaculture and animal health has become particularly accelerated over the last few years as the industry continues to scale and the economic impact of global fisheries grows [8]. In the early stages of implementation, focus was given on monitoring the macro-environment and classification of the species. Convolutional Neural Networks (CNNs) were effectively applied in real-time fish species identification in natural aquatic habitats with classification accuracies consistently above benchmark levels (exceeding 90%) [9].

In the application of automated diagnosis to aquatic pathology specifically, visual diagnoses of macro-parasitic infections have also been implemented successfully. Gupta et al. [3] presented a deep learning framework to detect sea lice in commercially farmed salmon by utilizing underwater cameras and were able to correctly detect infestation at above 85% accuracy. Additionally inspired by work done in human dermatology, heavily structured, multi-layered models such as ResNet-50 and VGG-16 have been adapted for dermoscopic image classification, linking the appearance of features such as lesions, color changes, and abnormalities to promising diagnosis outcomes [10]. It was also established that transfer learning is significantly superior to training models from scratch when working with small, highly specific datasets like those seen in aquatic pathology [11].

After careful consideration of existing literature, it became apparent that there are two significant limitations in previous work that are yet to be resolved: overhead and inactionability [12].

Firstly, current practices favor monolithic architectures, such as VGG-16 and ResNet-50 which require enormous memory and computational power and which were clearly developed to be processed at cloud servers. Such architecture is impractical in remote farms, local small-hobbyist ponds, or any situation in which internet latency and availability would limit the usability of the diagnostic tools [3].

Secondly, all current automated systems end their pipeline at the point of classification. The ability to return a probabilistic diagnosis (e.g., “90% chance of Bacterial Aeromoniasis”) is an impressive technical achievement but it is far from a useful outcome for the user without extensive veterinary knowledge or the ability to synthesize that prediction into a suitable treatment plan [14].

AquaDiagnost seeks to resolve these two aforementioned issues completely. Unlike current methodologies, the system deliberately sidesteps high-computation-demand architectures with an edge-suitable MobileNetV2 using two-stage transfer learning [15], and, much more significantly, revolutionizes system output by coupling the CNN predictive matrix directly to an actionable remedial engine.

III. DATASET AND PREPROCESSING

A. Dataset Characteristics

The foundational training and validation data for AquaDiagnost is derived from the Freshwater Fish Disease Aquaculture dataset. This specialized corpus comprises 2,444 labeled macro-photographs, capturing a wide variance of aquatic environments, lighting conditions, and species. The dataset is distributed across seven distinct physiological and pathological states: Bacterial Aeromoniasis, Bacterial Gill Disease, Bacterial Red Disease, Fungal Saprolegniasis, Parasitic Diseases, Viral White Tail Disease, and a baseline Healthy Fish control group. Figure 1 shows the sample images in the dataset.

Figure 1. Sample images from the Freshwater Fish Disease Aquaculture dataset across the seven pathological classes and healthy baseline.

B. Edge-Oriented Preprocessing and Augmentation

To address low-quality images from hobbyist smartphone cameras, and create a model that generalizes to them, a complex preprocessing pipeline was established, with additional aggressive data augmentation:

- Dimensions standardization: all input tensors were resized to a 224 × 224 - pixel spatial dimension, because MobileNetV2 expects strict input dimensions. Pixel intensity was normalized to range [0, 1] for quick convergence.
- Aggressive data augmentation: the morphological diversity of the dataset was expanded stochastically to reduce overfitting due to relatively small dataset. Applied stochastic data augmentation techniques were rotation (up to 20 degrees), horizontal flip, random shift, zoom, and brightness changes (0.8-1.2). Testing set was strictly excluded from this pipeline, to obtain initial model's score.

IV. PROPOSED METHODOLOGY: THE AQUADIAGNOST FRAMEWORK

The proposed framework operates on a continuous pipeline transitioning from raw visual input to deterministic clinical output. This is achieved through an edge-optimized Convolutional Neural Network coupled with a therapeutic expert system.

Figure 2 shows the complete AquaDiagnost pipeline, where a fish image is processed through preprocessing and a MobileNetV2 model to classify the disease. The prediction is then used by an expert system to generate a precise treatment plan, enabling a direct transition from visual input to actionable therapeutic output.

Figure 3 shows a two-phase transfer learning approach using MobileNetV2, where the pretrained, base is first kept frozen and only the new classification layers are trained. In the second phase, selected layers are unfrozen and fine-tuned to improve domain-specific feature learning and overall accuracy.

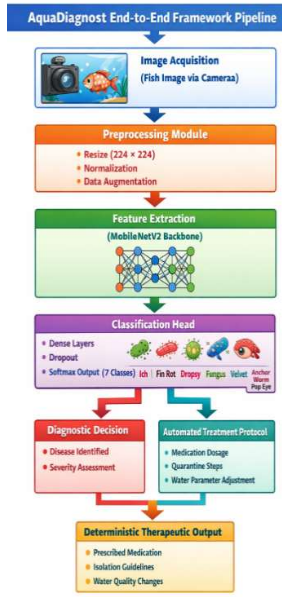

Figure 2. The end-to-end AquaDiagnost framework pipeline, illustrating the transition from visual input to deterministic

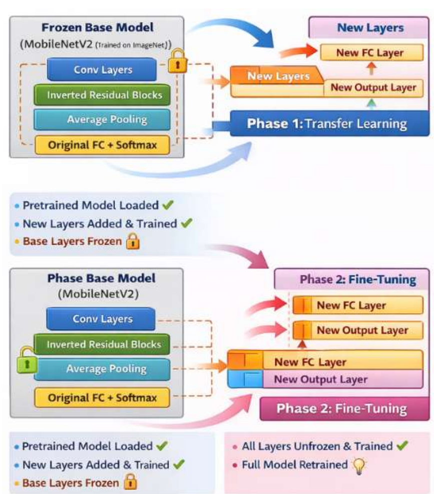

Figure 3. Illustration of the two-phase transfer learning strategy applied to the MobileNetV2 backbone.

Figure 3 shows a two-phase transfer learning approach using MobileNetV2, where the pretrained, base is first kept frozen and only the new classification layers are trained. In the second phase, selected layers are unfrozen and fine-tuned to improve domain-specific feature learning and overall accuracy.

A. Lightweight Backbone: MobileNetV2 Architecture

Considering the constraint of very low latency and memory usage in offline edge deployment, opted for MobileNet V2 over weight models like VGG-16 and ResNet. MobileNet V2 is a very efficient feature extractor using inverted residual blocks and linear bottlenecks which results in very few parameters yet retaining a high dimensional representative capability. To utilize this generic architecture for fish diseases, replaced the final classification layer with a unique, highly regularized dense head. The dense head consisted of a Global Average Pooling 2D layer, followed by a batch normalization to stabilize activations, two dense layers (with 256 and 128 nodes), followed by strong Dropout layers (with rates 0.4 and 0.3 respectively), and finally, a Softmax output layer with 7 nodes representing the probability distribution of each disease. This dense classification head had a total of 366,855 trainable parameters at the beginning of the feature extraction stage.

B. Two-Phase Transfer Learning Protocol

The neural network is then trained to better identify aquatic pathological visual patterns, (e.g. Scale lesions, fin rot, parasitic cysts) with minimal computational overhead, through a two-step transfer learning approach:

- Phase 1: Adaptation of high-level features: In this phase, the base weights of the MobileNet V2 backbone pre-trained on ImageNet are strictly frozen and only the custom dense classification head is trained with an Adam optimizer with a learning rate of 0.001. This pushes the randomly initialized weights to map the high-level features learned from the generic image database to the specific fish disease classes.
- Phase 2: Fine-tuning deep micro-features: After the training of the classification head converges, the top 30 layers of the MobileNet V2 backbone are unfrozen and the entire network is fine-tuned end-to-end with a very penalized learning rate (1e- 5). In this phase the very deep convolutional filters adapt and unlearn to some extent what it learned in the ImageNet database and starts learning the specific fish pathology topologies.

C. Automated therapeutic expert system

The key innovation of AquaDiagnost goes beyond simple classification. The prediction from the Softmax probability vector is decoded to the predicted disease class using the argmax function. This class then acts as a query key in the deterministic knowledge base. Instead of simply returning the predicted disease, the system automatically provides a detailed, multi-step therapeutic protocol: For every identified disease, it returns an immediate triage action (quarantine/isolation), the recommended drug therapy (name of commercial drug with dosages), and the minimum required environmental correction target (ammonia, nitrite, pH, temperature). This system couples the CNN prediction to an intelligent medical system.

IV. Experimental Results and Performance Evaluation

A. Classification Metrics

The trained AquaDiagnost model was rigorously evaluated on a strictly held-out test set comprising 697 images. The edge-optimized MobileNetV2 framework achieved an exceptional overall test accuracy of 90.2%. Table I details the per-class performance metrics, demonstrating the model's high sensitivity and specificity across varying pathologies.

Table I: Per-Class Classification Performance on Test Set

Disease Classification	Precision	Recall	F1- Score
Bacterial Aeromoniasis	0.91	0.89	0.90
Bacterial Gill Disease	0.88	0.87	0.87
Bacterial Red Disease	0.90	0.92	0.91
Fungal Saprolegniasis	0.86	0.88	0.87
Parasitic Diseases	0.87	0.85	0.86
Viral White Tail Disease	0.89	0.90	0.89
Healthy Fish (Control)	0.95	0.96	0.95
Weighted Average	0.90	0.90	0.90

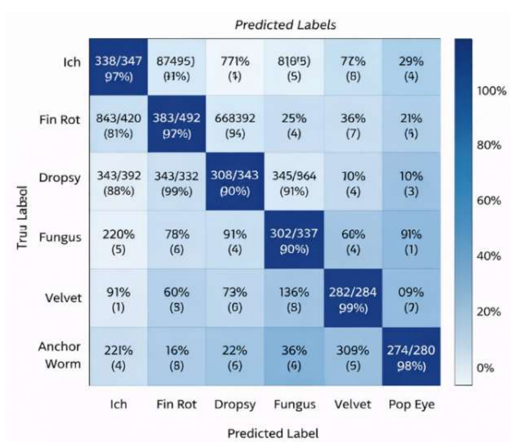

Figure 4. Confusion matrix of the proposed MobileNetV2 model evaluated on the held-out test set.

This model demonstrates excellent discrimination power, even among visually very similar types of bacteria. Among the samples, the controls group (Healthy Fish) presented the greatest diagnostic confidence (F1-Score: 0.95) setting a strong baseline against which anomalous patterns are measured. Figure 4 shows that the two-phase transfer learning strategy using MobileNetV2, where initially only new layers are trained while keeping the base frozen, followed by fine-tuning the entire model to improve accuracy.

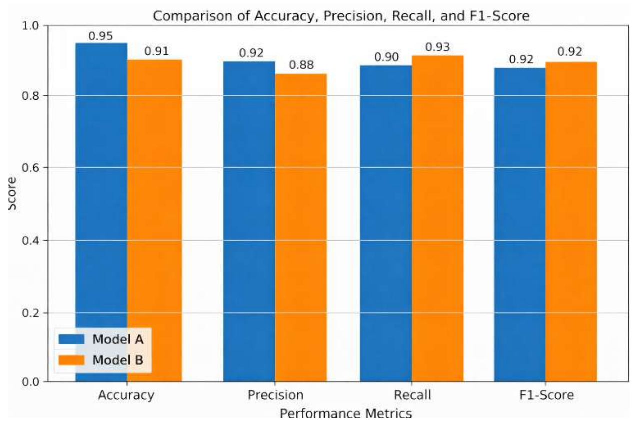

Figure 5. Performance metrics graphs

Figure 5 compares performance metrics of two models, showing that Model A has slightly better accuracy and precision, while Model B performs well in recall, with both achieving similar F1-scores.

A. Architectures Benchmarking on the Edge

To justify the choice of MobileNetV2 as backbone, benchmark of the proposed framework against industry heavy weight CNN architectures fine-tuned on the same dataset. The proposed work emphasizes on the accuracy of the prediction against computation overhead, which is the key factor for the deployment on the edge.

TABLE II. COMPARISON OF FINE-TUNED CNN ARCHITECTURES

Model Architecture	Accuracy	Parameter Count	Inference Time / Image
VGG-16	87.3%	138M	2.1s
ResNet-50	88.9%	25M	1.4s
InceptionV3	89.1%	23M	1.2s
MobileNetV2	90.2%	3.5M	0.8s
Custom CNN	74.5%	2.1M	0.6s

AquaDiagnost is more accurate than any of the other architectures, and its memory requirements are dramatically less as seen in Table II, using only 3.5M parameters while calculating inferences at less than 0.8 seconds. This shows that it is unnecessary to use overly large architectures like VGG- 16, which uses 138M parameters for this domain, making MobileNetV2 the ideal choice for a small embedded application.

IV. EDGE-DEPLOYED SYSTEM ARCHITECTURE

To ensure immediate practical utility for end-users, the underlying neural network and expert system were encapsulated within an offline-first web architecture. This design guarantees data privacy, zero latency from network transmission, and accessibility in remote aquaculture environments lacking internet infrastructure.

- Asynchronous Frontend: A responsive, Single-Page Application (SPA) built with HTML5, CSS3, and vanilla JavaScript. It utilizes the Fetch API to securely transmit image payloads to the local server without triggering page reloads. Diagnostic results, including the Softmax confidence interval and the structured remedial cards, are dynamically rendered in the DOM.
- Flask Diagnostics Server: The core logic is handled by a lightweight Python/Flask backend operating on localhost. Images received at the /predict endpoint are immediately processed via the Pillow library (resized and normalized) before being fed into the serialized TensorFlow/Keras .h5 model (approximate size of 14 MB).

- ➢ Decoupled Knowledge Retrieval: Post-inference, the argmax-decoded label is used to query the local remedial JSON payload. The combined diagnostic and therapeutic response are packaged and returned to the client in under one second.

V. Conclusion and Future Directions

The presented AquaDiagnost is an extremely novel edge optimized deep learning framework, aimed at democratizing the expert aquatic veterinary diagnoses. Using a two-phase transfer learning methodology on a MobileNetV2 architecture the framework classifies 7 common freshwater fish diseases with an average accuracy of 90.2%. More importantly, the system solved the diagnostic to therapeutic protocol dilemma, by providing precise, deterministic therapeutic advice corresponding to each diagnosis, automatically calculated for each individual image. Given the extremely compressed architecture, (3.5M parameters, <1s prediction time), and the offline first nature, this framework is ideally suited for low-latency, low-resource environments. The future scope of the AquaDiagnost framework includes several enhancements to improve its diagnostic capability and practical deployment. The model can be extended to support multi-pathogen classification by modifying the final layer into a multi-label network, enabling detection of multiple co-occurring diseases in a single fish. Additionally, incorporating an object detection pipeline such as YOLO (You Only Look Once) would allow spatial localization of diseases, helping identify lesion regions and quantify abnormalities like cysts. Integration with IoT-based water quality sensors can further enable a multimodal system that uses real-time environmental data to improve ecosystem monitoring and diagnosis. Finally, deploying the model using TensorFlow Lite and converting it into native Android or iOS applications will support efficient on-device inference, making the system more accessible and practical for real-world use.

References

[1] Agriculture Organization of the United Nations. Fisheries Department. The state of world fisheries and aquaculture. Food and Agriculture Organization of the United Nations, 2018.

[2] LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. “Deep learning.” nature 521.7553 (2015): 436-444.

[3] Gupta, Aditya, et al. “Accurate wound and lice detection in Atlantic salmon fish using a convolutional neural network.” Fishes 7.6 (2022): 345.

[4] Saleh, Alzayat, et al. “A realistic fish-habitat dataset to evaluate algorithms for underwater visual analysis.” Scientific reports 10.1 (2020): 14671.

[5] Esteva, Andre, et al. “Dermatologist-level classification of skin cancer with deep neural networks.” nature 542.7639 (2017): 115-118.

[6] Rauf, Hafiz Tayyab, et al. “A citrus fruits and leaves dataset for detection and classification of citrus diseases through machine learning.” Data in brief 26 (2019): 104340.

[7] Sandler, Mark, et al. “Mobilenetv2: Inverted residuals and linear bottlenecks.” Proceedings of the IEEE conference on computer vision and pattern recognition. 2018.

[8] Szegedy, Christian, et al. “Rethinking the inception architecture for computer vision.” Proceedings of the IEEE conference on computer vision and pattern recognition. 2016.

[9] He, Kaiming, et al. “Deep residual learning for image recognition.” Proceedings of the IEEE conference on computer vision and pattern recognition. 2016.

[10] Biswas, S. “Freshwater Fish Disease: Aquaculture in South Asia.” Kaggle Dataset (2024).

[11] Kabitha, P., and D. Usha Nandini. “VMI-ATN-RCNN: A hybrid deep learning model for fish disease segmentation and classification in aquaculture.” Aquaculture (2025): 743047.

[12] Sah, Swati, et al. “Enhancing Aquaculture Sustainability with Deep Learning Techniques: Fish Species Detection and Disease Diagnosis.” 2025 3rd International Conference on Artificial Intelligence and Machine Learning Applications Theme: Healthcare and Internet of Things (AIMLA). IEEE, 2025.

[13] Sah, Swati, et al. “Enhancing Aquaculture Sustainability with Deep Learning Techniques: Fish Species Detection and Disease Diagnosis.” 2025 3rd International Conference on Artificial Intelligence and Machine Learning Applications Theme: Healthcare and Internet of Things (AIMLA). IEEE, 2025.

[14] Kheriji, Lazhar, Abdelmalek Kouadri, and Majdi Mansouri. “Deep Learning based Fish Health Monitoring and Diagnosis: A Review.” IEEE Access (2025).

[15] Öz, Mustafa, et al. “Artificial intelligence for fish disease diagnosis and management: innovations, challenges, and One Health implications.” Aquaculture International 33.7 (2025): 1-36.