

A Smart IOT-Driven Digital Platform for Sustainable Dry Fish Processing in Sri Lanka

RP-026-079

Project Proposal Report

Gunasekara D T C D P

B.Sc. (Hons) Degree in Information Technology Specialized in
Information Technology


Department of Information Technology

Sri Lanka Institute of Information Technology Sri Lanka


March 2026

DECLARATION

I declare that this is my own work, and this proposal does not knowingly incorporate any material previously submitted for a degree or diploma at any other university or higher education institution, nor does it contain any material that has been previously published or written by another person apart from where proper acknowledgment is given in the text.

Name	Student ID	Signature
GUNASEKARA D T C D P	IT22127396	

The above candidate is conducting research for their undergraduate Dissertation under my supervision.

Name	Date	Signature
Ms. Jenny Krishara	2026/03/14	

ABSTRACT

Sri Lanka's dry fish processing industry remains one of the most important yet underserved sectors of the country's coastal economy. Quality control throughout the processing cycle relies entirely on subjective manual inspection by individual processors, resulting in inconsistent grading outcomes, undetected contamination events, and product rejections that cause significant financial losses. No automated, objective, or repeatable quality grading system exists for Sri Lankan dry fish.

This research addresses that gap by developing an AI-powered computer vision module specifically designed for automated dry fish quality grading and contamination detection within a controlled drying prototype environment. The module uses a Convolutional Neural Network trained on a locally collected image dataset of Sri Lankan dry fish batches, captured under standardised artificial lighting conditions within the prototype's enclosed chamber. The camera system supports two operational modes: a live view stream accessible to producers through the Flutter mobile application, and a periodic automatic capture mode that feeds images into the quality grading pipeline at regular intervals throughout the drying process.

A central novelty of this research is the self-learning verification station, implemented as a Next.js professional dashboard. Images for which the model returns low-confidence predictions are automatically routed to this station for review by domain professionals with expertise in dry fish quality assessment. Verified classifications from professionals are added to the model's retraining dataset, enabling continuous improvement over successive retraining cycles. This mechanism eliminates the dependency on a large pre-existing labelled dataset and provides a practical pathway to a highly accurate, locally validated model without reliance on any external third-party analysis service.

The module integrates with the broader IoT platform, passing grading results to the predictive analytics component for enriched completion time and spoilage risk predictions, and storing all grading evidence within each batch's traceability record. Technologies used include Python, TensorFlow, Keras, OpenCV, FastAPI, Node.js, Flutter, and Next.js.

Keywords: Computer Vision, CNN, Dry Fish Quality Grading, Contamination Detection, Self-Learning Model, Verification Station, Sri Lanka, Controlled Drying, IoT Integration, Transfer Learning

ACKNOWLEDGEMENT

I would like to express my sincere gratitude to my supervisor, Ms. Jenny Krishara, for her continuous guidance, constructive feedback, and unwavering support throughout the development of this research. Her expertise in software systems and research methodology has been instrumental in shaping the quality and direction of this work.

I also extend my heartfelt thanks to my co-supervisor, Dr. Dinuka Wijendra, for her valuable academic input and technical guidance, which helped strengthen the research methodology and analytical rigour of this study.

Special thanks go to our external supervisor, Mr. Dinuth Jayasekara, whose first-hand experience as a dry fish business owner provided an invaluable practical perspective. His knowledge of quality expectations, contamination challenges, and processing realities in the Sri Lankan context was essential in ensuring that the proposed computer vision module is practically relevant and grounded in real-world conditions.

I also acknowledge my project group members for their collaboration and teamwork throughout the development of the integrated platform. Finally, I thank my family for their patience and support throughout this research journey.

TABLE OF CONTENTS

DECLARATION	2
ABSTRACT.....	3
ACKNOWLEDGEMENT	4
TABLE OF CONTENTS	5
LIST OF TABLES	7
LIST OF FIGURES.....	8
LIST OF ABBREVIATIONS	8
LIST OF APPENDICES	9
1 INTRODUCTION.....	10
1.1 Background and Context.....	10
1.1.1 Analysis of Existing Solutions and Their Limitations	11
1.1.2 SDG Alignment.....	12
1.2 Literature Review.....	12
1.2.1 CNN-Based Image Classification	12
1.2.2 Transfer Learning for Small Datasets	12
1.2.3 Computer Vision for Food Quality Grading	13
1.2.4 Defect Detection in Food Products	13
1.2.5 Mobile-Based Food Quality Inspection	13
1.2.6 Fungal Contamination in Dried Fish.....	13
1.2.7 Lightweight CNN Architectures	14
1.2.8 IoT-Based Monitoring in Food Processing	14
1.2.9 Digital Traceability in Fish Supply Chains	14
1.2.10 Data Augmentation for Robust Classification.....	14
1.2.11 Food Loss and Waste in Low-Income Settings	15
1.3 Research Gap.....	16
2 OBJECTIVE	17

2.1	Main Objective	17
2.2	Specific Objectives.....	17
3	METHODOLOGY	18
3.1	Controlled Drying Prototype and Camera Setup.....	19
3.2	Dual Camera Operating Modes	20
3.3	Image Dataset Collection	20
3.4	Self-Learning Model and Verification Station.....	21
3.4.1	Model Development.....	21
3.4.2	Self-Learning Verification Station Workflow	21
3.5	Technology Stack.....	22
3.6	Ethical Considerations.....	23
3.7	Validation Strategy and Performance Metrics	24
4	HIGH-LEVEL SYSTEM ARCHITECTURE	24
4.1	Overview of the Computer Vision Component.....	24
4.2	Data Flow — Computer Vision Module.....	25
4.3	Scalability and Security	27
5	USER REQUIREMENTS	27
5.1	Stakeholder Identification	27
5.2	Functional Requirements	28
5.3	Non-Functional Requirements	29
6	COMMERCIALIZATION PLAN	31
6.1	Target Market.....	31
6.2	Value Proposition.....	31
6.3	Revenue Model and Pricing	32
6.4	Competitive Advantage and Intellectual Property.....	32
7	BUDGET AND JUSTIFICATION	33

7.1	Budget Breakdown	33
8	<i>WORK BREAKDOWN STRUCTURE (WBS)</i>	34
9	<i>GANTT CHART</i>	35
10	<i>REFERENCES LIST</i>	36
11	<i>APPENDICES</i>	38
	Appendix A : Informed Consent Form.....	38
	Appendix B : Image Grading Criteria Reference	39
	Appendix C : Risk Analysis	40
	Appendix D : Verification Station Dashboard - UI Description	41
	Appendix E : High level complete System Architecture	42
	Appendix F : Image processing architecture	43

LIST OF TABLES

Table 1 - Comparison of Existing Research Methods and Proposed System (Computer Vision Component).....	16
Table 2 – Technologies that use for the system	23
Table 3 - Data Flow summary table	26
Table 4 - Functional Requirements	29
Table 5 - Non-Functional Requirements	30
Table 6 - Revenue Model and Pricing.....	32
Table 7 - Budget Breakdown	34

LIST OF FIGURES

Figure 1 - System Overview Diagram	19
Figure 2 - Data Flow chart.....	26
Figure 3 - Work Breakdown Structure.....	34
Figure 4 - Project Timeline Gantt Chart	35

LIST OF ABBREVIATIONS

Abbreviation	Full Form
AI	Artificial Intelligence
API	Application Programming Interface
CNN	Convolutional Neural Network
CV	Computer Vision
DHT	Digital Humidity and Temperature (Sensor)
ESP32	Espressif Systems Microcontroller
IoT	Internet of Things
ML	Machine Learning
REST	Representational State Transfer
SDG	Sustainable Development Goal
UI	User Interface
WBS	Work Breakdown Structure

LIST OF APPENDICES

Appendix A : Informed Consent Form.....	38
Appendix B : Image Grading Criteria Reference	39
Appendix C : Risk Analysis	40
Appendix D : Verification Station Dashboard - UI Description	41
Appendix E : High level complete System Architecture	42
Appendix F : Image processing architecture	43
Appendix G : Single Component Diagram - Image processing part	44

1 INTRODUCTION

1.1 Background and Context

Sri Lanka's dry fish industry, known locally as karawala, is a culturally significant and economically vital component of the country's coastal food system. It provides an affordable, high-protein food source to millions of Sri Lankans across all income levels while supporting the livelihoods of tens of thousands of fishing and processing families in coastal communities including Negombo, Chilaw, Mannar, Jaffna, Trincomalee, and Matara.

Despite its importance, quality control in the dry fish industry remains entirely manual and subjective. Processors visually assess dried fish batches using personal experience and judgment, with no standardised grading criteria, no digital records of processing conditions, and no systematic method for detecting contamination or quality defects before product enters the market. This approach produces inconsistent quality assessments, misses contamination events, and leaves no traceable quality evidence within the supply chain.

Current open-air drying practices in Sri Lanka expose fish batches to direct contact with rocks, nets, and sandy ground, making contamination from insects, stray animals, birds, and environmental remains a near-certainty rather than an exception. Field clarifications confirm that fireflies, dogs, and other animals regularly interact with drying product. Processors acknowledge that product processed under current conditions has a maximum safe usable period of approximately two weeks before quality begins to deteriorate noticeably, severely limiting market reach and export viability.

Fungal contamination is among the most serious biological threats to dry fish quality. Research confirms that the dominant fungal species affecting dried fish products, including *Aspergillus flavus*, *Aspergillus parasiticus*, *Penicillium citrinum*, and *Cladosporium* species, require sustained relative humidity above 70 to 85 percent to initiate and sustain growth [9]. Sri Lanka's coastal climate routinely exposes open-air drying batches to humidity well above these thresholds, creating repeated opportunities for fungal colonisation that processors have no practical means to prevent or detect until visible mould appears on the product surface.

This research proposes and develops a computer vision module for automated quality grading and contamination detection, situated within a broader smart IoT-enabled platform that introduces a controlled drying prototype to eliminate the contamination risks endemic to traditional processing. The controlled chamber maintains humidity below the fungal growth threshold, provides standardised artificial lighting for reliable image capture, and integrates with IoT sensors and predictive analytics to create a fully digitised processing environment.

1.1.1 Analysis of Existing Solutions and Their Limitations

Existing approaches to fish quality assessment in Sri Lanka are exclusively manual. Processors rely on visual cues including colour, surface appearance, texture, and smell, with outcomes that vary significantly between individuals, across lighting conditions, and with changes in processor fatigue or experience. There are no standardised grading criteria applied consistently across the industry.

At the international level, computer vision and deep learning approaches to food quality grading are well established in industrialised markets. Automated inspection systems exist for fruit grading, grain sorting, and seafood processing in high-income settings. However, these systems are designed for large-scale operations with controlled laboratory environments, expensive high-resolution industrial cameras, and training datasets that bear no resemblance to Sri Lankan fish species, processing methods, or environmental conditions. Their cost and design assumptions make them entirely inaccessible and inapplicable for small-scale Sri Lankan coastal processors.

No existing system provides an automated, locally contextualised quality grading solution for Sri Lankan dry fish that operates within a real-world processing environment, accounts for the practical constraints of small producers, and improves continuously over time through professional verification.

This research directly addresses this gap by developing a CNN-based computer vision module with a self-learning verification station, trained on locally collected data from Sri Lankan processing contexts and integrated within an enclosed, hygienically controlled drying prototype.

1.1.2 SDG Alignment

- SDG 2 - Zero Hunger: Reducing post-harvest losses and improving food safety in a sector serving millions of Sri Lankan consumers.
- SDG 9 - Industry, Innovation and Infrastructure: Introducing AI-driven quality innovation to a traditional informal industry.
- SDG 12 - Responsible Consumption and Production: Minimising food waste through early defect detection and controlled processing.
- SDG 14 - Life Below Water: Supporting sustainable and responsible fisheries through traceability and quality assurance.

1.2 Literature Review

The following review examines twelve scholarly sources and existing systems relevant to the computer vision quality grading component, identifying key findings and limitations that justify the novelty of this research.

1.2.1 CNN-Based Image Classification

The study demonstrated that deep CNN architectures significantly outperform traditional feature engineering approaches for image recognition tasks, establishing the foundational principle that a CNN trained on domain-specific data can learn discriminative visual features automatically [1]. This principle directly guides the model development approach in this research.

1.2.2 Transfer Learning for Small Datasets

The study established the theoretical and empirical basis for transfer learning, demonstrating that models pre-trained on large general-purpose datasets such as ImageNet can be fine-tuned on smaller domain-specific datasets with strong performance outcomes [2]. Studies in food quality inspection have demonstrated successful transfer learning with as few as 300 to 500 training images, which is directly relevant given the dataset constraints of this research.

1.2.3 Computer Vision for Food Quality Grading

The research applied CNNs to fish freshness classification, achieving classification accuracy above 90 percent under controlled imaging conditions [3]. Bhargava and Bansal reviewed deep learning approaches for food quality inspection, confirming that CNN-based models generalise effectively across grading tasks when trained on domain-specific datasets, and that transfer learning substantially reduces minimum dataset size requirements [4]. Both studies note that model generalisation to real-world conditions, particularly varying lighting, remains a key challenge.

1.2.4 Defect Detection in Food Products

The research investigated surface defect detection in food and industrial products using neural network approaches, reporting that mould growth and discolouration are reliably detectable through colour space analysis and texture feature extraction [5]. These findings directly inform the defect labelling strategy of this research, which targets discolouration, uneven drying, mould-like appearance, and insect contamination signs as primary defect categories.

1.2.5 Mobile-Based Food Quality Inspection

The study validated that food quality classification models trained on smartphone-captured images can achieve performance comparable to high-resolution industrial cameras, provided that image capture guidelines are followed and training data covers varied conditions [6]. This supports the decision to use mobile phone image capture as the primary user-facing input method, ensuring accessibility without specialised hardware requirements.

1.2.6 Fungal Contamination in Dried Fish

The study investigated fungal spoilage in dried and salted fish products, identifying *Aspergillus flavus*, *Aspergillus parasiticus*, *Penicillium citrinum*, and *Cladosporium* species as the dominant fungal contaminants in dried fish processed under humid open-air conditions [7]. Their findings confirmed that sustained relative humidity above 70 to 85 percent is required for fungal growth initiation. Maintaining processing environment humidity below this threshold effectively prevents

fungal colonisation, directly validating the humidity control strategy of the proposed drying prototype.

1.2.7 Lightweight CNN Architectures

The research introduced the MobileNet family of lightweight CNN architectures specifically designed for resource-constrained deployment environments [8]. MobileNetV2 and EfficientNetB0 offer strong classification performance with significantly lower computational requirements than full-scale architectures, making them appropriate choices for the inference service that must respond within 10 seconds per image under normal network conditions.

1.2.8 IoT-Based Monitoring in Food Processing

The research demonstrated that IoT-based environmental monitoring using ESP32-class microcontrollers and standard DHT sensors can improve drying process consistency and provide early warning of unsafe conditions [9]. Their work establishes the hardware basis for the sensor integration layer of the broader platform within which the computer vision module operates.

1.2.9 Digital Traceability in Fish Supply Chains

The research conducted a systematic review of digital technologies for traceability in global fish supply chains, identifying that integrated platforms combining quality grading data with supply chain records represent a critical emerging priority [10]. Their review identified that small-scale processing contexts in low-income settings remain significantly underserved, directly motivating this research.

1.2.10 Data Augmentation for Robust Classification

The research demonstrated that data augmentation strategies including random brightness adjustment, contrast variation, horizontal flipping, and rotation substantially improve model robustness across lighting conditions without requiring specialised imaging hardware [11]. These

strategies are adopted in the training pipeline of this research to address the variability in natural lighting that may be encountered during site visit data collection.

1.2.11 Food Loss and Waste in Low-Income Settings

The research reported that post-harvest losses in fish processing in low-income settings account for a disproportionate share of economic loss and environmental impact in food supply chains [12]. This finding reinforces the importance of early quality detection, as the computer vision module's ability to identify defects at processing stage directly reduces the likelihood of rejections and wastage at market stage.

Technologies / Methods	Paper [1]	Paper [2]	Paper [3]	Paper [4]	Paper [5]	Proposed System
Automated Quality Grading	✓	✓	✓	✓	✓	✓
Defect Detection (mould/insects/colour)	✗	✗	✓	✗	✗	✓
Locally Trained Dataset (Sri Lanka)	✗	✗	✗	✗	✗	✓
Self-Learning / Continuous Improvement	✗	✗	✗	✗	✗	✓
Professional Verification Station	✗	✗	✗	✗	✗	✓
Mobile Image Capture Support	✗	✗	✗	✓	✗	✓

Technologies / Methods	Paper [1]	Paper [2]	Paper [3]	Paper [4]	Paper [5]	Proposed System
Standardised Lighting for Grading	X	X	X	X	X	✓
Live Camera / Real-Time View	X	X	X	X	X	✓
Integration with IoT Sensor Data	X	X	X	X	X	✓
Supply Chain Traceability Link	X	X	X	X	X	✓
Low-Cost / Small-Scale Applicable	X	X	X	✓	X	✓
Works in Controlled Environment	X	X	X	X	X	✓

Table 1 - Comparison of Existing Research Methods and Proposed System (Computer Vision Component)

1.3 Research Gap

The literature review reveals four interconnected research gaps that collectively justify this work.

First, no documented system provides automated visual quality grading specifically for Sri Lankan dry fish. The absence of a locally collected and labelled image dataset for Sri Lankan fish species, drying conditions, and defect types means that no model trained on existing international data can be reliably applied in this context. This research creates the first such dataset and the first locally validated grading model.

Second, all existing food quality grading systems are static; they are trained once and deployed without any mechanism for continuous improvement based on real-world experience. The proposed self-learning verification station introduces a continuous improvement pathway grounded in domain professional expertise, enabling the model to become progressively more accurate as it encounters new cases and receives verified corrections.

Third, existing systems do not integrate visual quality grading outputs with IoT sensor data from the same processing batch in a unified prediction pipeline. In this research, camera-derived quality signals (grade, detected defects, confidence) are passed alongside sensor readings to the predictive analytics module, creating a richer and more accurate prediction input than sensor data alone.

Fourth, the challenge of fungal contamination, which is the dominant quality threat in current Sri Lankan processing conditions, has not been addressed through a software-driven controlled environment approach. By maintaining humidity below fungal growth thresholds, the prototype eliminates the root cause of this contamination, and the camera system provides real-time visual confirmation of processing outcomes that manual inspection cannot reliably provide.

2 OBJECTIVE

2.1 Main Objective

To develop and evaluate an AI-powered computer vision module that automatically grades dry fish quality and detects visible contamination and defects using a Convolutional Neural Network with a self-learning verification station, thereby replacing subjective manual grading with a repeatable, objective, and continuously improving quality assurance system for Sri Lanka's dry fish processing industry.

2.2 Specific Objectives

1. To collect and label a representative image dataset of Sri Lankan dried fish batches captured under real-world processing conditions and standardised artificial lighting, covering quality grades A, B, and C, and defect categories including discolouration, uneven drying, mould-like growth, and insect contamination signs.

2. To design and document a standardised, repeatable grading criteria framework that maps visual characteristics to quality grades and defect labels, developed in consultation with domain experts in Sri Lankan dry fish quality assessment.
3. To train and optimise a CNN model using transfer learning on the collected dataset, achieving a minimum macro-averaged F1-score of 0.75 on the held-out test set.
4. To implement a dual-mode camera system supporting live streaming for producer remote monitoring and periodic automatic capture for quality grading, both operating under standardised artificial lighting within the controlled drying prototype.
5. To develop and deploy a Next.js Verification Station dashboard enabling domain professionals to review uncertain model predictions, provide verified corrections, and contribute to a continuously expanding retraining dataset that improves model accuracy over successive cycles.
6. To integrate the computer vision module with the broader IoT platform, passing grading results to the predictive analytics component and storing all grading evidence within each batch's traceability record.

3 METHODOLOGY

This research adopts a Design Science Research (DSR) methodology, appropriate for projects that create and evaluate information technology artefacts to solve real-world problems. The methodology is structured across six phases: prototype hardware setup, dataset collection, grading criteria definition, model development and training, verification station development, and integration and evaluation.




Figure 1 - System Overview Diagram

3.1 Controlled Drying Prototype and Camera Setup

The computer vision module operates within and depends on a custom-engineered controlled drying prototype designed by the research team. The prototype is an enclosed drying chamber that eliminates contamination risks endemic to traditional open-air processing. The chamber incorporates temperature control through infrared heating bulbs and heated air circulation, active humidity management that maintains the processing environment below 65 percent relative humidity (below the 70 to 85 percent fungal growth threshold established in the literature), weight monitoring via a load cell sensor, and environmental sensors at multiple internal positions.

A camera module is mounted within the prototype chamber with standardised artificial lighting panels positioned at a defined colour temperature and angle to provide consistent, uniform illumination across the drying surface. This controlled lighting environment is a deliberate design decision that directly addresses one of the primary challenges identified in the food quality grading literature: lighting variability causing model performance degradation. By controlling the lighting within the chamber, images captured for quality grading are consistent in colour balance, exposure, and shadow pattern regardless of external lighting conditions or time of day.

3.2 Dual Camera Operating Modes

The camera system operates in two distinct modes, both managed through the platform:

- **Live View Mode:** When activated by the producer through the Flutter mobile application, the camera streams live video of the drying batch in real time. This allows producers to visually monitor their fish remotely from anywhere using their smartphone without needing to open the chamber. Live view is activated on demand and deactivated when not needed, conserving bandwidth and system resources.
- **Periodic Capture Mode:** The camera automatically captures still images of the drying batch at configurable intervals throughout the processing cycle. These images are the primary input to the quality grading and contamination detection pipeline. Interval frequency is adjustable based on batch type and drying stage. Images are always captured under the standardised artificial lighting conditions within the chamber.

These two modes are independent and do not interfere with each other. A producer may activate live view at any time without disrupting the scheduled periodic capture cycle.

3.3 Image Dataset Collection

Prior to building the self-learning model, an initial dataset is collected through field visits to dry fish processing sites in coastal Sri Lankan communities. Photographs of dried fish batches are captured under varied natural conditions, including different fish species, drying stages, and visible defect types. This field-collected data supplements the images captured within the controlled prototype chamber and broadens the model's exposure to real-world processing variation.

All images will be labelled by two independent reviewers following the standardised grading criteria framework. A minimum target dataset of 300 to 500 labelled images is established, consistent with transfer learning benchmarks in the literature. Disagreements between reviewers are resolved through discussion or referral to the external supervisor Mr. Dinuth Jayasekara, whose industry expertise provides authoritative quality assessment for disputed cases.

3.4 Self-Learning Model and Verification Station

3.4.1 Model Development

The quality grading model is a CNN developed using Python with TensorFlow and Keras. A pre-trained MobileNetV2 or EfficientNetB0 architecture is selected as the base model, with the final classification layer replaced and configured for the target quality grade and defect classification tasks. Transfer learning is applied by initially freezing base layer weights before progressively unfreezing layers during fine-tuning.

Data augmentation techniques applied during training include random horizontal flipping, rotation, brightness adjustment, contrast variation, and zoom variation. These augmentations improve model robustness across the real-world image conditions encountered both during field data collection and within the prototype chamber.

The model produces two output types: a quality grade classification (A, B, or C) with associated confidence score, and defect detection flags identifying which specific defect categories are present in the image as a multi-label classification output.

3.4.2 Self-Learning Verification Station Workflow

The verification station is the core novel contribution of this component. It enables the model to improve continuously over time through structured domain professional input, without requiring the research team to manually label all future data. The workflow operates as follows:

1. The camera captures a periodic image. The image is submitted to the CNN model via the FastAPI inference service.
2. If the model returns a prediction meeting the defined confidence threshold, the result is automatically accepted and recorded in the batch traceability record.
3. If the model returns a low-confidence prediction, the image is automatically flagged and routed to the Verification Station dashboard, accessible to domain professionals through the Next.js application.

4. The domain professional reviews the flagged image, examines the model's tentative prediction, and provides a verified correct classification including the accurate quality grade and applicable defect labels.
5. The verified classification is stored as a ground truth label for that image. The corrected data point is added to the model's retraining dataset.
6. At defined intervals, the model is retrained on the expanded dataset including all newly verified examples. Each retraining cycle produces an improved model with broader coverage of real-world cases.
7. Over time, the proportion of images requiring verification station referral decreases as the model's confidence covers a growing range of cases, demonstrating measurable improvement.

This workflow ensures that all quality assessments within the platform are either made by a sufficiently confident model or verified by a qualified domain professional. It provides a documented pathway to a highly accurate, locally validated grading model grounded entirely in Sri Lankan dry fish quality expertise.

3.5 Technology Stack

Component	Technology	Justification
CV Model Development	Python, TensorFlow, Keras	Industry-standard deep learning; extensive transfer learning support
Image Processing	OpenCV	Efficient preprocessing, augmentation, colour space analysis

Component	Technology	Justification
Inference API	FastAPI (Python)	High-performance async REST; sub-10-second response for image grading
Verification Dashboard	Next.js	Server-side rendering; fast data-intensive professional review interface
Mobile Application	Flutter	Cross-platform iOS and Android from single codebase; live camera stream support
Backend API	Node.js (Express)	Lightweight, event-driven; well-suited for real-time IoT data and REST services
Database	MongoDB	Flexible document schema for varied batch records, sensor logs, image metadata
IoT + Camera Hardware	ESP32, DHT22, Load Cell, ESP32-CAM	Low-cost, widely documented; dual-mode camera support within prototype
Artificial Lighting	LED panels, controlled colour temperature	Standardised lighting for consistent colour-accurate image capture

Table 2 – Technologies that use for the system

3.6 Ethical Considerations

All images collected during field site visits are obtained with the informed written consent of site owners and processors. Consent forms are provided in both Sinhala and English. Participants are clearly informed of the research purpose, image usage, and their right to withdraw. No personal identification information is linked to images. Data is stored securely with role-based access

controls. Images submitted to the inference service are not stored permanently within the inference layer; only structured grading results are retained within batch traceability records.

3.7 Validation Strategy and Performance Metrics

The dataset is split into training (70%), validation (15%), and test (15%) sets. Model performance is evaluated on the held-out test set using accuracy, macro-averaged F1-score, precision, recall, and a confusion matrix identifying class-specific misclassification patterns.

A robustness evaluation assesses model performance separately on images captured under different artificial lighting intensities within the prototype chamber and on natural-light field images, quantifying any performance variation between controlled and uncontrolled capture conditions.

The self-learning improvement trajectory is measured by tracking the proportion of images requiring verification station referral over successive retraining cycles. A declining referral rate indicates measurable model improvement. User feedback is collected from producers and domain professionals following a small operational trial to assess usability and grading output alignment with expert judgment.

4 HIGH-LEVEL SYSTEM ARCHITECTURE

4.1 Overview of the Computer Vision Component

The computer vision module is one of four tightly integrated components within the broader dry fish processing platform. It operates as an independent microservice that receives batch images from the camera system within the controlled drying prototype, processes them through the trained CNN inference engine, and returns structured grading results to the central Node.js backend. The module communicates with the predictive analytics component by providing quality grade and defect data as enriched inputs, and with the traceability component by contributing grading evidence to each batch record.

4.2 Data Flow — Computer Vision Module

Step	Actor	Action	Output
1	Camera (Periodic Mode)	Captures still image at defined interval within prototype chamber	Raw image file
2	Node.js Backend	Receives image, validates format and size, forwards to CV inference service	Validated image
3	FastAPI Inference Service	Preprocesses image; runs CNN model; generates grade + defect output	Grade, confidence, defect flags
4a	High Confidence Path	Result accepted automatically; stored in batch traceability record	Batch record updated
4b	Low Confidence Path	Image flagged; routed to Verification Station dashboard	Verification task created
5	Domain Professional	Reviews flagged image; provides verified grade and defect labels	Verified ground truth label
6	Retraining Pipeline	Adds verified data to retraining dataset; model retrained at defined interval	Improved model deployed
7	Predictive Analytics Module	Receives grading result as enriched input alongside sensor data	Updated predictions

Step	Actor	Action	Output
8	Flutter Mobile App	Displays grading result, defect indicators, and confidence on batch dashboard	Producer informed

Table 3 - Data Flow summary table




Figure 2 - Data Flow chart

4.3 Scalability and Security

The FastAPI inference service is stateless and horizontally scalable. Model weights are loaded once at service startup and cached in memory, minimising per-request inference latency. All API communications use HTTPS. The Verification Station dashboard is accessible only to users with a verified domain professional role, enforced through the platform's authentication layer. Images are not stored permanently in the inference service. Grading results in batch records are governed by role-based access controls consistent with the traceability module's security design.

5 USER REQUIREMENTS

5.1 Stakeholder Identification

- Dry Fish Producers and Processors: Activate live camera view, receive automated grading results, and monitor batch quality through the Flutter mobile application.
- Domain Professionals and Quality Inspectors: Access the Next.js Verification Station to review low-confidence model predictions and provide verified quality classifications.
- Buyers and Retailers: View grading histories within batch traceability records before purchase decisions.
- Regulators and Certification Bodies: Access batch grading records for compliance monitoring and certification support.
- System Administrators: Manage user roles, schedule model retraining cycles, and monitor inference service health.

5.2 Functional Requirements

ID	Requirement	Priority
FR01	The system shall capture still images of the drying batch automatically at configurable intervals and submit them to the quality grading model.	High
FR02	The system shall provide a live video stream of the drying batch accessible through the Flutter mobile application on producer demand.	High
FR03	The system shall classify each batch image into quality grade A, B, or C and return a confidence score.	High
FR04	The system shall detect and report visible defects in each image including discolouration, uneven drying, mould-like growth, and insect contamination signs.	High
FR05	The system shall automatically route low-confidence predictions to the Next.js Verification Station dashboard for domain professional review.	High
FR06	The system shall allow domain professionals to review flagged images, correct quality grades and defect labels, and submit verified classifications.	High
FR07	The system shall store verified classifications in the model retraining dataset and support scheduled retraining cycles.	High

ID	Requirement	Priority
FR08	The system shall display grading results including grade, defect indicators, and confidence score on the batch dashboard within 10 seconds of image submission.	High
FR09	The system shall store all grading results, verified classifications, timestamps, and image references within the batch traceability record.	High
FR10	The system shall pass grading results to the predictive analytics module as enriched inputs for completion and spoilage prediction.	High
FR11	The system shall reject invalid image uploads and return a clear error message to the user.	Medium
FR12	The system shall maintain an audit log of all Verification Station activities including reviewer identity, review time, and classification changes.	Medium

Table 4 - Functional Requirements

5.3 Non-Functional Requirements

ID	Requirement	Category
NFR01	The inference service shall return a grading result within 10 seconds of receiving a valid image under normal network conditions.	Performance
NFR02	The live camera stream shall deliver video with latency of no more than 3 seconds from capture to display on the mobile app.	Performance

ID	Requirement	Category
NFR03	The CNN model shall achieve a minimum macro-averaged F1-score of 0.75 on the held-out test dataset after initial training.	Accuracy
NFR04	All API communications shall be conducted over encrypted HTTPS connections.	Security
NFR05	The Verification Station shall be accessible only to users with a verified domain professional role.	Security
NFR06	Image data submitted for grading shall not be retained permanently within the inference service layer.	Privacy
NFR07	The mobile application grading interface shall be usable by producers with limited digital literacy following a brief orientation.	Usability
NFR08	The system shall maintain at least 99% inference service availability during active drying hours.	Reliability
NFR09	The self-learning model's verification referral rate shall show measurable decline across successive retraining cycles.	Improvement

Table 5 - Non-Functional Requirements

6 COMMERCIALIZATION PLAN

6.1 Target Market

The primary target market is small to medium-scale dry fish producers and processing cooperatives across Sri Lanka's coastal communities. Estimated at 80,000 to 100,000 active fishing households, a significant proportion engage in fish drying and processing as a primary or supplementary income source. Secondary markets include fish traders seeking certified quality documentation, exporters requiring compliance evidence, and government food safety agencies.

Medium-term expansion opportunities include other dried seafood and agricultural products across Sri Lanka, and regional markets in South Asia including Bangladesh, southern India, and the Maldives, where similar small-scale dried seafood processing industries face comparable quality control challenges.

6.2 Value Proposition

- For producers: Automated grading replaces inconsistent manual inspection, reduces market rejections, extends the product's safe usable period beyond the current two-week maximum, and provides documented quality evidence supporting better price negotiation.
- For buyers and retailers: Access to quality-verified, traceable product with documented batch grading histories reduces procurement risk.
- For regulators: Digital quality grading records simplify compliance monitoring without requiring physical inspection of every batch.

6.3 Revenue Model and Pricing

Tier	Price	Features
Basic (Free)	LKR 0 / month	Up to 20 batch gradings/month, basic dashboard, limited traceability records
Producer	LKR 1,500 / month	Unlimited gradings, full dashboard, complete traceability, predictive analytics
Business	LKR 5,000 / month	All Producer features + multi-user, Verification Station access, advanced analytics, export compliance docs
Enterprise	Custom	White-label, API access, government/cooperative bulk licensing, implementation consulting

Table 6 - Revenue Model and Pricing

6.4 Competitive Advantage and Intellectual Property

The primary competitive advantage is local specificity. The model trained on Sri Lankan dry fish data, the verification station workflow, and the integrated controlled-environment platform constitute proprietary assets that cannot be quickly replicated. The self-learning verification station mechanism, which routes uncertain predictions to domain professionals and uses verified outcomes for continuous retraining, is a novel system design that existing commercial and academic solutions do not incorporate.

The trained model weights, labelled image dataset, grading criteria framework, verification station codebase, and integrated platform are protectable through copyright registration. The specific methodology of combining IoT sensor data with visual quality grading in a traceability-linked self-learning platform presents a basis for a patent application.

7 BUDGET AND JUSTIFICATION

7.1 Budget Breakdown

Item	Description	Cost (LKR)
Prototype Chamber (shared)	Enclosure, heating, insulation — shared across team	25,000 (shared)
Heating System (shared)	Infrared bulbs, heated air fan — shared across team	15,000 (shared)
IoT Sensors (shared)	DHT22 x4, load cell + HX711 — shared across team	8,000 (shared)
ESP32 Microcontrollers	ESP32 x2 for sensor integration	3,000
Camera Module	ESP32-CAM for dual-mode operation within prototype	2,500
Artificial Lighting Panels	LED panels with standard colour temperature for CV capture	6,000
Image Dataset Collection	Transport, printing, participant refreshments for site visits	15,000
Cloud GPU (CV Training)	Google Colab Pro or AWS for CNN model training	8,000
Cloud Hosting (CV Service)	FastAPI inference service + Next.js dashboard hosting	5,000

Item	Description	Cost (LKR)
Documentation and Printing	Reports, consent forms, grading criteria reference cards	3,000
Contingency (10%)	Unexpected costs in prototype or data collection	7,050
TOTAL (this component)		~77,550

Table 7 - Budget Breakdown

8 WORK BREAKDOWN STRUCTURE (WBS)




Figure 3 - Work Breakdown Structure

9 GANTT CHART




Figure 4 - Project Timeline Gantt Chart

10 REFERENCES LIST

- [1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “ImageNet classification with deep convolutional neural networks,” *Communications of the ACM*, vol. 60, no. 6, pp. 84–90, Jun. 2017, doi: 10.1145/3065386.
- [2] S. J. Pan and Q. Yang, “A survey on transfer learning,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 22, no. 10, pp. 1345–1359, Oct. 2010, doi: 10.1109/TKDE.2009.191.
- [3] X. Zhang, Y. Liu, and C. Li, “Fish freshness classification using convolutional neural networks,” *Food Control*, vol. 110, Art. no. 107052, 2020, doi: 10.1016/j.foodcont.2019.107052.
- [4] A. Bhargava and A. Bansal, “Fruits and vegetables quality evaluation using computer vision: A review,” *Journal of King Saud University – Computer and Information Sciences*, vol. 33, no. 3, pp. 243–257, Mar. 2021, doi: 10.1016/j.jksuci.2018.06.002.
- [5] P. Jiang, Y. Chen, B. Liu, D. He, and C. Liang, “Real-time defect detection using Bayesian forecasting and deep neural networks,” *IEEE Transactions on Industrial Informatics*, vol. 18, no. 7, pp. 4744–4752, Jul. 2022, doi: 10.1109/TII.2021.3134237.
- [6] N. Kang, W. Zhang, and Y. Liu, “Mobile-based food quality inspection using transfer learning and lightweight CNNs,” *Food and Bioprocess Technology*, vol. 15, no. 4, pp. 823–837, 2022, doi: 10.1007/s11947-022-02777-1.
- [7] A. Riba, R. Nout, and A. Savelkoul, “Fungal spoilage in dried and salted fish products,” *International Journal of Food Microbiology*, vol. 119, nos. 1–2, pp. 68–75, 2007, doi: 10.1016/j.ijfoodmicro.2007.07.052.
- [8] A. G. Howard *et al.*, “MobileNets: Efficient convolutional neural networks for mobile vision applications,” *arXiv preprint*, arXiv:1704.04861, 2017. [Online]. Available: <https://arxiv.org/abs/1704.04861>
- [9] N. Mishra *et al.*, “Development of drying system by using internet of things for food quality monitoring and controlling,” *Energy Nexus*, vol. 11, Art. no. 100219, Sep. 2023, doi: 10.1016/j.nexus.2023.100219.

- [10] J. Cromwell, C. Turkson, M. Dora, and F. A. Yamoah, "Digital technologies for traceability and transparency in global fish supply chains," *Marine Policy*, vol. 178, Art. no. 106700, Aug. 2025, doi: 10.1016/j.marpol.2025.106700.
- [11] L. Perez and J. Wang, "The effectiveness of data augmentation in image classification using deep learning," *arXiv preprint*, arXiv:1712.04621, 2017. [Online]. Available: <https://arxiv.org/abs/1712.04621>
- [12] J. Gustavsson, C. Cederberg, U. Sonesson, R. van Otterdijk, and A. Meybeck, *Global Food Losses and Food Waste: Extent, Causes and Prevention*. Rome, Italy: FAO, 2011.

11 APPENDICES

Appendix A : Informed Consent Form

RESEARCH CONSENT FORM

IT4010 Research Project | SLIIT University | RP-026-079

Study Title: AI-Based Computer Vision for Dry Fish Quality Grading in Sri Lanka

You are being invited to participate in a research study by allowing the research team to capture photographs of your dry fish batches. Images will be used solely to train a quality grading model as part of an undergraduate research project at SLIIT University.

☐ I agree to allow photographs of my dry fish batches to be captured for this research.

☐ I understand the images will be used for academic research purposes only.

☐ I understand my identity will not be linked to any published findings.

Name: _____

Signature: _____

Date: _____

Appendix B : Image Grading Criteria Reference

Grade / Defect	Visual Description	Decision
Grade A	Uniform colour, fully and evenly dried, clean surface, no contamination	Acceptable Premium –
Grade B	Minor colour variation, mostly dried, minor surface irregularities	Acceptable Standard –
Grade C	Significant defects present; not suitable for standard market	Reject
Discolouration	Dark patches, yellowing, or uneven colour across surface	Flag Defect
Uneven Drying	Areas visibly wetter or lighter, inconsistent drying surface	Flag Defect
Mould-like Growth	White, grey, or green fuzzy patches or surface film	Flag Defect
Insect Contamination	Visible insects, larvae, eggs, or surface damage from infestation	Flag Defect

Appendix C : Risk Analysis

Risk	Likelihood	Impact	Mitigation
Insufficient labelled images	Medium	High	Multiple site visits; aggressive augmentation; start collection early
Model overfitting (small dataset)	Medium	High	Transfer learning; dropout regularisation; cross-validation
Poor lighting variation robustness	Low	Medium	Standardised artificial lighting within chamber minimises this risk
Low professional engagement with Verification Station	Medium	High	External supervisor and community contacts as initial reviewers
Camera module hardware failure	Low	Medium	Backup camera module procured; modular mount allows replacement
Dataset labelling inconsistency	Medium	High	Dual-reviewer protocol with inter-rater reliability check before training
Backend integration delays	Low	Medium	Define API contract early; develop inference service with mock data in parallel


Appendix D : Verification Station Dashboard - UI Description

The Next.js Verification Station dashboard displays the following elements for each image flagged for professional review:


1. Batch ID, timestamp, and drying stage at the time of image capture.
2. The flagged batch image displayed at full resolution.
3. The model's tentative quality grade (A/B/C) and confidence score displayed prominently below the image.
4. The model's detected defect flags listed as candidate labels.
5. A correction panel allowing the professional to select the correct quality grade and tick all applicable defect labels that accurately describe the image.
6. A submit button that records the professional's verified classification and adds it to the retraining dataset.
7. A batch history panel showing all previous grading events for the same batch, enabling contextual review.
8. An overall statistics panel showing total reviews completed, model confidence distribution, and retraining dataset size.

The layout is designed to minimise the time required per review while maximising the quality and consistency of verified label data collected for model improvement.

Appendix E : High level complete System Architecture



Appendix F : Image processing architecture



Appendix G : Single Component Diagram - Image processing part

