

Smart Karawala: A Batch-Centric AI–IoT Framework for Predictive and Automated Dry-Fish Processing

Ridmi Vivipem
Faculty of Computing
Sri Lanka Institute of Information
Technology
Malabe, Sri Lanka
it22639844@my.sliit.lk

Chankama Gunasekara
Faculty of Computing
Sri Lanka Institute of Information
Technology
Malabe, Sri Lanka
it22127396@my.sliit.lk

Jayani Kalansooriya
Faculty of Computing
Sri Lanka Institute of Information
Technology
Malabe, Sri Lanka
it22554536@my.sliit.lk

Milan Sanjaya
Faculty of Computing
Sri Lanka Institute of Information
Technology
Malabe, Sri Lanka
it22189394@my.sliit.lk

Jenny Krishara
Faculty of Computing
Sri Lanka Institute of Information
Technology
Malabe, Sri Lanka
jenny.k@sliit.lk

Dinuka Wijendra
Faculty of Computing
Sri Lanka Institute of Information
Technology
Malabe, Sri Lanka
dinuka.w@sliit.lk

Abstract—Small-scale dry-fish processing in Sri Lanka depends heavily on operator experience and open-air drying, producing inconsistent quality, spoilage risk, undocumented waste, and fragmented batch histories. Existing work addresses environmental monitoring, machine-learning drying prediction, process control, and food traceability as separate functions, and rarely returns a prediction to the drying equipment as an operating set-point. This paper presents Smart Karawala, an integrated Artificial Intelligence (AI) and Internet of Things (IoT) framework that unifies these functions within a single batch-centric processing lifecycle. A unique batch identifier links preparation data, continuous temperature, humidity, gas and weight sensing, machine-learning prediction of the recommended drying temperature and drying time, rule-based spoilage and over-drying assessment, automated feedback-based drying control, and sustainability records. A two-phase start protocol converts the predicted parameters into the chamber's control profile and records a batch as drying only after the chamber independently confirms a valid sensor state. An independent safety layer clamps every recommended set-point, faults on missing sensor data, and can stop the heater regardless of the model output. On a physics-informed modelled drying dataset the remaining-time regressor reached a cross-validated coefficient of determination of 0.96 (mean absolute error 1.1 minutes) under run-level grouped validation. We report the implementation and integration status of each subsystem and define the field experiments required for full validation.

Index Terms—artificial intelligence, dry-fish processing, Internet of Things, machine learning, predictive analytics, process control, smart drying, traceability

I. INTRODUCTION

Drying is one of the oldest food-preservation methods: reducing available moisture suppresses microbial growth and enzymatic activity, extending shelf life. In Sri Lanka roughly a quarter of the marine fish harvest is dried or smoked, and

dried fish (*karawala*) is a staple processed seafood product supporting coastal livelihoods. Reported post-harvest quality losses in the sector are high, and traditional open-sun drying is associated with spoilage, dust and insect contamination, and weather-dependent inconsistency [1].

Conventional processing remains highly manual. Operators judge cleaning losses, salt quantity, salting duration, drying conditions, and final dryness from experience, while open drying yards expose the batch to changing temperature and humidity. The consequences are interlinked: environmental fluctuation changes the rate of moisture removal; inappropriate conditions cause under-drying, over-drying, or spoilage; and waste volumes and batch histories are seldom recorded.

The Internet of Things (IoT), machine learning (ML), and closed-loop control offer routes to replace experience-dependent judgement with measurable, data-driven workflows. Implemented in isolation, however, their combined value is limited. A drying-time estimate is more useful when continuously updated from live observations; a spoilage indicator is more informative alongside the physical progression of the same batch; and a control decision is more defensible when the evidence behind it is retained against the batch record. A recent review of intelligent food drying identifies closed-loop integration of sensing, prediction, and control as an open direction [2].

This paper proposes *Smart Karawala*, a unified framework that treats preparation, sensing, prediction, control, verification, and traceability as stages of one processing lifecycle. The framework combines continuous chamber temperature, relative humidity, batch weight, gas-sensor response, and batch-specific preparation information. These observations support

recommended-temperature prediction, initial and dynamically updated drying-time estimation, spoilage-risk assessment, automated control actions, and drying-completion decisions.

Contributions. (i) A batch-centric AI-IoT architecture in which one identifier links preparation, environmental and weight sensing, ML predictions, control actions, safety events, and sustainability records across loosely coupled services and a mobile operator application. (ii) A two-phase drying-start protocol in which ML-predicted parameters become the chamber’s control profile and the analytics service records a batch as drying only after the chamber independently confirms a valid sensor state and running heater. (iii) An independent safety layer – set-point clamping, sensor-completeness faulting, no-resume-after-restart, and a debounced over-drying auto-stop – that operates independently of, and can override, the ML path. (iv) A reported reconciliation failure mode: a service holding a *logical* record of drying must not be treated as authoritative about whether drying is *physically* occurring. (v) An evidence-based implementation and evaluation status report with a defined protocol for the field experiments still required.

The research question is whether integrating heterogeneous sensing, predictive analytics, automated control, and event-level logging provides a coherent and safe framework for digitising small-scale dry-fish processing.

II. RELATED WORK

A. IoT and Intelligent Food Drying

Mishra *et al.* [3] developed an IoT drying system for food quality monitoring and control, showing that real-time measurement and actuation stabilise drying conditions. Alvinika *et al.* [4] regulated an automatic fish-drying chamber from temperature and humidity using fuzzy logic on a low-cost microcontroller, confirming that closed-loop control of a fish dryer is achievable with inexpensive hardware. Such systems mark a shift toward instrumented processing, but their measurements describe the drying *environment* rather than the physical state of the product, and a review of the field [2] observes that many deployed systems remain reactive.

B. Machine Learning for Drying Prediction

Rather than assuming a fixed duration, a predictive system updates its remaining-time estimate as observations accumulate. Rakholia *et al.* [5] predicted meat drying time from historical process data using an ensemble model for production planning, and Ahmad *et al.* [6] estimated grain moisture in a drying silo from environmental parameters. Classical thin-layer models remain the physical baseline for such estimators; Ghanem *et al.* [7] fitted them to solar-dried salted tilapia fillets with close agreement to measured moisture ratios. Across this literature, predictions are usually generated offline or as advisory outputs, and few systems return a predicted parameter to the drying equipment as a control set-point.

C. Product-State Sensing and Traceability

Setyagraha *et al.* [8] built a portable electronic nose from low-cost metal-oxide gas sensors and machine-learning classifiers to discriminate seafood freshness, supporting an inexpensive gas sensor as an odour proxy. For traceability, Cromwell *et al.* [9] reviewed digital traceability in global fish supply chains and highlighted integration as a prerequisite for higher traceability levels, while Tsang *et al.* [10] demonstrated a blockchain-and-IoT architecture for perishable-food traceability. That research concentrates on catch, cold chain, and distribution, treating processing itself as a single opaque event.

The literature therefore establishes the individual feasibility of sensing, prediction, control, and traceability, but leaves a system-level gap: predictions are seldom returned as equipment set-points, drying temperature is typically operator-selected, and processing decisions are rarely preserved in the traceability record.

III. PROPOSED INTEGRATED FRAMEWORK

A. Architecture and Batch-Centric Data Model

Smart Karawala is organised around the lifecycle of an individual batch rather than around separate applications (Fig. ??). It is realised as four loosely coupled services communicating over HTTP and persisting to a document database, together with a cross-platform mobile application: an authentication service; a batch and sustainability service that owns batch data, cleaning-loss and salt estimation, salting progression, waste records, and traceability reporting; an environmental-monitoring service that runs the drying-chamber state machine and actuator control at the edge; and a predictive-analytics service that produces the ML predictions, rule-based risk assessments, and orchestrates drying.

Each batch receives a unique identifier linking preparation information (fish variety, raw and cleaned weight, salt quantity and duration), environmental observations, weight progression, analytical predictions and risk assessments, sustainability information, and processing and traceability events. The analytics service reads batch data from the batch service and live readings from the monitoring service over their public interfaces, so no service touches another’s database.

B. End-to-End Processing Workflow

The framework follows a *sense–predict–decide–control–verify–trace* cycle. A batch is created with its fish variety and raw weight; cleaning loss and salt requirement are estimated and the batch is salted, with drying gated until salting is recorded complete. Before drying, the analytics service predicts the recommended temperature and total time. Those values are registered with the chamber as a control profile; the chamber validates its own sensors and starts; only then does the analytics service record the batch as the single active drying batch. During drying, a periodic loop pulls the latest reading, re-estimates the remaining time, evaluates spoilage and over-drying risk, and – at high over-drying risk – stops the heater. Completion is detected from weight and elapsed

time, the chamber runs a cooling phase, and every significant event is written against the batch identifier.




Fig. 1: Smart Karawala architecture linking the mobile application, predictive analytics, batch/sustainability services, and the edge drying controller.

IV. INTELLIGENT MONITORING AND PREDICTIVE PROCESSING

A. Environmental and Physical Sensing

An edge microcontroller streams a periodic sensor block over a serial link from a digital temperature and humidity sensor, a separate temperature probe, a metal-oxide gas sensor, and a load cell with a signal-conditioning amplifier. Malformed lines are rejected rather than parsed into data, and a reading is used for control only when every required field is present and physically plausible. The environmental state is $E_t = [T_t, H_t]$, and a safe operating condition is

$$C_E(t) = \begin{cases} 1, & T_{\min} \leq T_t \leq T_{\max} \wedge H_t \leq H_{\max} \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

A software-simulation mode reproduces the device message format for development and demonstration.

B. Weight Progression

The load cell continuously measures batch weight. From the captured initial weight W_0 , the normalised ratio, percentage reduction, and rate of change are

$$\begin{aligned} R_W(t) &= W_t/W_0, & L_W(t) &= 100(W_0 - W_t)/W_0, \\ r_W(t) &= (W_{t-\Delta t} - W_t)/\Delta t. \end{aligned} \quad (2)$$

Weight progression describes physical evolution that cannot be obtained from chamber temperature alone. Because the two time predictors share this signal, the analytics service applies a reliability gate: the weight-loss rate is used only once enough time has passed and enough mass has been lost for the ratio to clear sensor noise; before that, the pre-drying estimate is displayed.

C. Gas and Odour Sensing

The metal-oxide gas sensor provides an odour-related signal associated with volatile compounds released during protein degradation. The raw reading is mapped to discrete smell levels using configurable prototype thresholds; consistent with [8], this is treated as an uncalibrated odour proxy and not as a direct spoilage measurement.

D. Pre-Drying Parameter Prediction

Before drying begins, the system predicts the recommended chamber temperature and expected total drying time from conditions known in advance,

$$[\hat{\theta}, \hat{R}_0] = f_{\phi}(\text{variety}, W_0, H_0, G_0). \quad (3)$$

This differs from approaches that treat drying temperature as a fixed operator input: here temperature is itself predicted and in turn determines the expected duration. The predictor is a multi-output regression ensemble; a gradient-boosted and a random-forest candidate are compared by cross-validation and the stronger retained. Where a live reading is unavailable at prediction time, a documented default is substituted and the result is marked accordingly, so the operator can distinguish a sensor-grounded estimate from a fallback.

E. Dynamic Remaining-Time Prediction

During drying, accumulated observations are used to re-estimate progression. The predictive feature vector at time t is

$$Z_t = [\text{variety}, W_0, W_t, T_t, H_t, \tau_t, r_W(t)], \quad (4)$$

where τ_t is elapsed drying time, and the estimate is $\hat{R}_t = f_{\theta}(Z_t)$. The prediction is refreshed periodically. To keep the displayed countdown steady under sensor noise, the client applies a short rolling average and permits an upward correction only when it exceeds a tolerance.

F. Spoilage and Recommendation Assessment

Spoilage risk is assessed by a transparent rule over the odour level, humidity, temperature, and stalled weight loss:

$$s = \sum_k w_k \mathbb{K}[c_k], \quad P_s = \begin{cases} \text{High}, & s \geq s_{\text{hi}} \\ \text{Medium}, & s_{\text{lo}} \leq s < s_{\text{hi}} \\ \text{Low}, & s < s_{\text{lo}}, \end{cases} \quad (5)$$

where each c_k is a threshold condition and w_k its weight. The interface reports each contributing reading with its value, expected range, and direction of influence, so model limitations remain visible. A learned spoilage classifier is planned once labelled data are available. A companion rule set turns the current readings and risk level into a short textual drying recommendation.

G. Sustainability Intelligence

At batch creation the platform estimates the expected cleaning waste and the required salt quantity and salting duration from per-variety coefficients, and stores them against the batch rather than as isolated calculations. Salting progress is estimated from elapsed time against the recommended duration. When waste is recorded, a recycler-collection message is composed and stored against the batch together with the waste type, quantity, timestamp, handling method, and batch identifier.

V. DECISION, CONTROL, AND TRACEABILITY

A. Automated Drying Control

The chamber is governed by a state machine with states $\text{READY} \rightarrow \text{DRYING} \rightarrow \text{COOLING} \rightarrow \text{COMPLETED}$, plus STOPPED and FAULT . In automatic mode the heater is switched with a fixed hysteresis about the target temperature, and the exhaust fan responds to the humidity target. A control action is generated as $U_t = \pi(S_t, E_t)$, where S_t is the combined processing state and U_t is heater activation, fan activation, continuation, protective intervention, or termination.

B. Independent Safety Layer

Two safety properties are enforced by construction. Every recommended temperature is clamped to a configured operating envelope before it can reach the controller,

$$\theta_{\text{safe}} = \min(\max(\hat{\theta}, T_{\min}), T_{\max}), \quad (6)$$

so an erroneous prediction cannot command an unsafe set-point; this clamp is applied in the analytics service, re-applied when drying starts, and re-applied again on the client. Independently, the chamber validates sensor completeness on every reading and transitions to a FAULT state that de-energises the heater after a bounded number of consecutive missing or implausible readings, on a weight increase beyond a noise allowance, and on any service restart – a heating session is never silently resumed. A separate debounced check raises over-drying risk when the batch is at or past its target dryness while still being heated, or when the chamber runs sustainedly above target; at high risk the analytics service issues a stop command to the chamber. Safety therefore remains independent of model output.

C. Completion Verification

The prototype evaluates two independent completion criteria,

$$D_t = \begin{cases} 1, & W_t \leq W_0/k \text{ or } \tau_t \geq \hat{R}_0 \\ 0, & \text{otherwise,} \end{cases} \quad (7)$$

where k is a configurable completion ratio. The weight criterion provides physical evidence of progression, while the elapsed-time criterion acts as a prototype fallback based on the pre-drying estimate rather than as a laboratory-validated dryness measurement. On completion the heater is turned off and the exhaust fan runs for a configured cooling period before the batch is marked complete. The weight-loss

completion criterion is a configurable prototype parameter, not an established universal standard, and requires per-variety experimental validation.

D. Logical and Physical State Reconciliation

Prototype operation exposed a failure mode worth reporting. The analytics service and the chamber maintain separate notions of whether drying is occurring. When the chamber halted – through a safety fault or a lost sensor link – the analytics service continued to report an active batch and the interface continued a countdown derived from the pre-drying estimate: the displayed state was internally consistent yet physically false. The resolution is to make the chamber authoritative: the client reads the chamber’s live session status, suspends the countdown, and explains the discrepancy whenever the two disagree. A service holding a logical record of an activity must not be permitted to assert that the activity is physically occurring.

E. Event-Level Traceability

Each significant processing event is stored as $T_k = (BID, t_k, type_k, value_k, source_k)$. Recorded events span batch creation, cleaning-loss and salt estimation, salting start and completion, control-profile receipt, drying start, actuator commands, missed-sensor and fault events, weight-completion and drying-completion, predictions, risk assessments, and the waste record. Traceability completeness can subsequently be evaluated by comparing the number of recorded processing events with the expected event set. The objective is not merely to identify where a product originated, but to retain a digital history of how the batch was processed.

VI. IMPLEMENTATION AND EXPERIMENTAL METHODOLOGY

A. Implementation

Four asynchronous Python services communicate over HTTP and persist to a document database, with an in-memory fallback for control state. The edge service communicates with the chamber over serial commands for heater, fan, light, and tare. A cross-platform mobile application presents live readings, predictions, risks, actuator status, alerts, and batch history. Drying starts only after the predicted control profile is registered and the chamber independently validates its sensor state.

The implemented Smart Karawala prototype and a dried-fish product obtained during prototype operation are shown in Fig. 2. The physical setup integrates the sensing and actuation hardware required for environmental monitoring, weight tracking, and automated drying control.




Fig. 2: Physical implementation of Smart Karawala: (a) IoT-enabled drying chamber and (b) dried-fish product obtained during prototype operation.

B. Dataset and Model Training

The two drying-time models were trained on a drying dataset of 3,000 pre-drying cases and 2,767 in-process observations across 420 drying profiles, spanning twelve fish varieties. Feature values and drying-time targets were derived from thin-layer hot-air drying theory [6] and anchored to measured oven runs. Features and targets follow (3) and (4). Because observations from one drying profile are strongly autocorrelated, the dynamic model was evaluated with five-fold grouped cross-validation, with every profile assigned wholly to one fold; the pre-drying model, whose rows are independent batches, used standard five-fold cross-validation. Three candidate regressors were compared per task – linear regression, random forest, and gradient boosting – and the strongest retained.

C. Evaluation Metrics

Regression models are evaluated with the coefficient of determination R^2 , mean absolute error (MAE), and root-mean-square error (RMSE),

$$\text{MAE} = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i|, \quad \text{RMSE} = \sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2}. \quad (8)$$

Future field evaluation will measure temperature deviation, humidity-band exposure, actuator latency, sensor availability, communication success, and traceability completeness.

VII. RESULTS AND DISCUSSION

A. Predictive Model Performance

Table I reports the cross-validated performance of the candidate regressors. For the dynamic remaining-time task, evaluated with grouped cross-validation by drying profile, the random forest achieved $R^2 = 0.956 \pm 0.006$ with a mean absolute error of 1.10 minutes and an RMSE of 1.50 minutes, ahead of gradient boosting and linear regression. For the joint pre-drying prediction of temperature and total time, the random forest reached $R^2 = 0.745 \pm 0.007$; the gradient-boosted model was within one standard deviation. Fig. 3 shows out-of-fold predicted versus actual remaining time under grouped validation, with points distributed close to the identity line across the operating range.




Fig. 3: Out-of-fold predicted versus actual remaining drying time for the random-forest regressor under grouped five-fold cross-validation. Each drying profile is held wholly in one fold; the dashed line denotes ideal agreement.

TABLE I: Predictive Model Performance (Five-Fold Cross-Validation)

Task	Model	R^2	MAE	RMSE
Remaining time ^a (minutes)	Random Forest	0.956	1.10	1.50
	Gradient Boosting	0.950	1.23	1.61
	Linear Regression	0.913	1.59	2.12
Pre-drying ^b (temp. + time)	Random Forest	0.745	–	–
	Gradient Boosting	0.745	–	–
	Linear Regression	0.656	–	–

^aGrouped cross-validation by drying profile; MAE/RMSE in minutes.

^bJoint two-target regression; R^2 averaged over temperature (°C) and time (h) targets, so a single MAE/RMSE is not reported.

The pre-drying task is harder than the dynamic task, as expected: before drying starts only the variety, weight, humidity, and gas reading are known, whereas the dynamic model additionally sees the realised weight-loss trajectory, which is the dominant predictor of how much time remains. This is also why the framework shows the pre-drying estimate only until the weight signal becomes reliable and then hands over to the dynamic model.

B. System Functional Validation

Table II distinguishes cross-validated ML functions from software-level functional and unit validation. The two-phase start, risk rules, safety clamp, completion/cooling sequence, and state reconciliation are integrated end-to-end; controller and sensor-parser tests cover malformed readings, completion events, and heater cut-off.

TABLE II: System Functional Validation

Function	Evidence	Status
Pre-drying prediction	5-fold model report	Integrated; cross-validated
Remaining-time prediction	Grouped 5-fold report	Integrated; cross-validated
Two-phase start	Sensor validation before active state	Integrated; functional test
Risk + auto-stop	Rule output and stop command	Integrated; functional test
Safety + sensor faults	Clamp, bounded-failure fault	Integrated; unit-tested
Completion + cooling	Weight/time rule; cooling state	Integrated; unit-tested
State reconciliation	Chamber status gates countdown	Integrated; functional test
Waste/salt + traceability	Batch-linked coefficients/events	Integrated; domain validation pending

C. Discussion

Two prototype findings are important. First, prediction and safety must remain separate: the configured envelope can override a model output before actuation. Second, the chamber must be authoritative about physical drying state; this removed a failure mode in which a logical active-batch record continued a countdown after the chamber had stopped. Intermittent load-cell dropouts also triggered protective faults, showing that sensor robustness is a prerequisite for reliable intelligent control. These results establish integration behaviour, not superiority over conventional drying.

VIII. LIMITATIONS AND FUTURE WORK

The models are cross-validated on a physics-informed modelled dataset, not an independent corpus of logged field runs. Weight loss is not equivalent to laboratory moisture content, and the completion ratio, gas mapping, alert thresholds, and temperature envelope remain configurable engineering parameters. Spoilage assessment is rule-based because labelled ground truth is unavailable. The serial protocol confirms command transmission but does not read back relay state, so actuator execution is not independently acknowledged. Traceability is batch-keyed across services but the unified timeline is only partially consolidated, and no controlled baseline comparison has yet been performed. Future work will collect multi-variety field runs with reference moisture measurements, add actuator acknowledgement, calibrate sensors and thresholds, train a labelled spoilage classifier, consolidate traceability, and compare Smart Karawala with conventional processing using control stability, completion error, manual interventions, and record completeness.

IX. CONCLUSION

Smart Karawala integrates batch preparation, environmental and weight sensing, ML drying prediction, rule-based risk assessment, feedback-based control, safety intervention, and event-level traceability within one batch-centric workflow. The two-phase start returns predicted parameters to the chamber only after sensor validation, while an independent safety path

constrains set-points and can stop heating. On the physics-informed modelled dataset, grouped five-fold validation of the remaining-time random forest produced $R^2 = 0.956$, MAE = 1.10 min, and RMSE = 1.50 min. The integrated prototype also exposed two system-level requirements: physical device state must override conflicting logical state, and sensing reliability is fundamental to safe automation. Field trials are still required to quantify benefit over conventional processing.

ACKNOWLEDGMENT

The researchers thank the academic supervisors, domain experts, and dry-fish processors who supported this work.

REFERENCES

- [1] Food and Agriculture Organization of the United Nations, "Postharvest losses in the multiday fisheries sector of sri lanka," FAO, Rome, Italy, 2021, [Online]. Available: <https://www.fao.org/flw-in-fish-value-chains/>.
- [2] J. Guo, X. Zhang, Y. Liu, J. Wu, H. Xu, H. Xiao, Z. Ai, and P. Gong, "Intelligent monitoring, predicting, and control technology of food drying: recent advances, challenges, and future prospects," *Food Control*, vol. 180, p. 111666, 2026.
- [3] N. Mishra, S. K. Jain, N. Agrawal, N. K. Jain, N. Wadhawan, and N. L. Panwar, "Development of drying system by using internet of things for food quality monitoring and controlling," *Energy Nexus*, vol. 11, p. 100219, Sep. 2023.
- [4] Y. Alvinika, D. B. Setyohadi, and M. Sulistyoningih, "IoT-based monitoring and design of automatic fish drying equipment using fuzzy logic," *IOP Conference Series: Earth and Environmental Science*, vol. 704, no. 1, p. 012042, 2021.
- [5] R. Rakholia, A. L. Suárez-Cetrulo, M. Singh, and R. S. Carbajo, "AI-driven meat food drying time prediction for resource optimization and production planning in smart manufacturing," *IEEE Access*, vol. 13, pp. 22 420–22 428, 2025.
- [6] F. Ahmad, M. S. Younis, R. U. Zahid, and L. A. Shahid, "Machine learning based grain moisture estimation for real-time monitoring of high-temperature paddy drying silo," in *Proceedings of the IEEE 23rd International Multitopic Conference (INMIC)*, 2020, pp. 1–6.
- [7] T. H. M. Ghanem *et al.*, "Thin-layer modeling, drying parameters, and techno-enviro-economic analysis of a solar dried salted tilapia fish fillets," *Scientific Reports*, vol. 15, p. 5073, 2025.
- [8] M. R. M. Setyagraha, H. A. Nurqamaradillah, L. M. Hermawan, N. R. Pratama, L. Novamizanti, and D. R. Wijaya, "A portable real-time electronic nose for evaluating seafood freshness using machine learning," *IEEE Access*, vol. 13, pp. 82 000–82 019, 2025.
- [9] J. Cromwell, C. Turkson, M. Dora, and F. A. Yamoah, "Digital technologies for traceability and transparency in the global fish supply chains: a systematic review and future directions," *Marine Policy*, vol. 178, p. 106700, Aug. 2025.
- [10] Y. P. Tsang, K. L. Choy, C. H. Wu, G. T. S. Ho, and H. Y. Lam, "Blockchain-driven IoT for food traceability with an integrated consensus mechanism," *IEEE Access*, vol. 7, pp. 129 000–129 017, 2019.