Multimodal Real-Time Inventory Estimation Using Vision, Load Cells, and Depth Sensing Hardware

Felipe Cardozo
Department of Computer Science
Emory University
Atlanta, GA, USA
focardo@emory.edu

Leonardo Affonso
Department of Electrical Engineering
Federal University of Rio de Janeiro (UFRJ)
Rio de Janeiro, Brazil
laffonso@poli.ufrj.br

Abstract*—Accurate real-time food inventory tracking in commercial kitchen environments remains a critical yet underaddressed challenge at the intersection of embedded systems, computer vision, and applied artificial intelligence. This paper presents a fully integrated multimodal system for food inventory estimation that combines YOLO-based object detection, strain gauge load cell sensing, and LiDAR depth reconstruction, all deployed on resource-constrained edge hardware. The architecture tightly couples four subsystems: (1) a YOLOv8n detection module achieving mAP@0.5 = 0.847 and mAP@0.5:0.95 = 0.612 across 12 ingredient classes, running at 7.8 FPS under INT8 quantization on Raspberry Pi 4; (2) a 24-bit HX711 load cell platform with Kalman-filtered mass estimation at MAE = 1.2 g; (3) an Intel RealSense L515 depth sensor performing surface-integral volumetric reconstruction at MAE = 12.1 cm^3; and (4) a time-series analytics engine employing ARIMA(2,1,2) demand forecasting with MAPE = 4.7%. Optimal inverse-variance fusion of the load cell and LiDAR modalities yields a fused volume estimate at MAE = 5.1 cm^3 and $R^2 = 0.9973$, with a 95% confidence interval of $\pm 14.3\,cm^3$. End-to-end pipeline latency under INT8 quantization is 132.8 ms (mean), enabling near-real-time continuous inventory monitoring. Ablation studies confirm that sensor fusion consistently outperforms any single-modality baseline by 19–82%. The system operates at a total power draw of 8.6 W and is estimated to reduce avoidable food waste by 18–26% per operational week based on demand-forecasting intervention analysis.

***Index Terms*—food inventory estimation, edge computing, multimodal sensing, YOLO, load cells, LiDAR, depth sensing, sensor fusion, Raspberry Pi, time-series forecasting, ARIMA, Kalman filter, embedded AI**

I. Introduction

Commercial kitchen operations generate substantial inefficiencies due to imprecise inventory management. Studies estimate that up to 30% of perishable ingredients in restaurant-scale operations are discarded due to overstocking, late detection of depletion events, and mismatched preparation scheduling [1]. Current inventory solutions predominantly rely on manual stocktaking—a labor-intensive process prone to human error and low temporal resolution—or barcode-based tracking systems that capture item-level entry and exit events but provide no continuous quantity monitoring [2].

Automated computer vision approaches have demonstrated promise for ingredient recognition [3], [4] but suffer from fundamental limitations when used in isolation: 2D bounding-box detectors cannot directly measure quantity, and visual occlusion from container shapes and stacking introduces significant estimation variance. Load cell systems provide high-precision mass measurement but are confounded by container weight, multi-ingredient containers, and the nontrivial density-to-volume mapping required for granular inventory accounting. Depth-sensing approaches, particularly structured-light and LiDAR modalities, have been explored for agricultural yield estimation [8] and retail shelf monitoring [9] but have not been rigorously evaluated in the context of real-time kitchen inventory with edge deployment.

This work addresses these limitations through a purpose-built multimodal fusion architecture that simultaneously exploits the complementary strengths of vision, mass, and depth sensing. The main contributions of this paper are:

1) A complete hardware-software co-design integrating four heterogeneous sensing modalities on Raspberry Pi 4 edge hardware, operating within a 8.6 W power budget.
2) A mathematically formalized fusion framework employing inverse-variance optimal weighting of load-cell and LiDAR modalities, with rigorous error propagation analysis under ingredient density uncertainty.
3) A systematic empirical evaluation across 12 ingredient classes and 300 annotated estimation scenarios, including full ablation studies over modality combinations and quantization configurations.
4) An ARIMA-based demand forecasting extension enabling proactive kitchen restocking with a 30-minute horizon and MAPE of 4.7%.
5) A deployment feasibility and sustainability analysis quantifying real-world food waste reduction potential.

The remainder of this paper is organized as follows. Section II surveys related work. Section III describes the system architecture. Section IV presents the mathematical formulation. Section V details the experimental setup. Section VI reports quantitative results. Section VII presents the multimodal comparison and error analysis. Section VIII evaluates deployment feasibility. Section IX presents the forecasting extension. Section X discusses cost-benefit and sustainability implications. Section XI covers limitations and future work. Section XII concludes.

II. Related Work

Deep learning for food recognition has matured [3]–[5]. Recent transformers [6] exceed edge constraints, making YOLO detectors [7] optimal. Weight-based IoT systems [12]–[14] offer tracking but cannot identify ingredients without vision. The HX711 ADC [15] remains standard for affordable precision. Depth-based estimation [8], [9] works for generic portioning [10], with L515 sensors [11] improving ambient-light robustness over early structured-light. Sensor fusion for inventory is sparse [16], [17] and lacking in kitchen deployment frameworks. This work introduces the first system to jointly integrate YOLOv8, load cells, and LiDAR under rigorous edge constraints with full ablation evaluation.

III. System Architecture

A. Hardware Platform

The system is deployed on a Raspberry Pi 4 Model B (quad-core ARM Cortex-A72 @ 1.8 GHz, 4 GB LPDDR4 RAM) connected to three peripheral sensing subsystems via USB 3.0 and I2C interfaces. Total hardware bill of materials is detailed in Table I. The Intel RealSense L515 provides synchronized RGB (1920×1080) and depth (1024×768) streams at 30 FPS via USB 3.1. Four SparkFun 10 kg load cells are arranged in a Wheatstone bridge configuration beneath a stainless-steel platform, digitized by an HX711 24-bit ADC sampling at 80 SPS on the I2C bus.

TABLE I: System Hardware Bill of Materials

Component	Model	Unit Cost	Interface
Edge Computer	Raspberry Pi 4B (4GB)	$55	–
Depth Sensor	Intel RealSense L515	$229	USB 3.1
Load Cell (×4)	SparkFun 10 kg	$14	I2C (HX711)
ADC Amplifier	HX711 24-bit	$4	I2C
Camera	RPi Camera Module v2	$25	CSI-2
Platform	Steel weighing plate	$18	–
Power Supply	5 V/3 A USB-C PD	$12	–
Total		**$371**	

B. Software Stack

The software pipeline is implemented in Python 3.11 using Ultralytics YOLOv8 for detection, ONNX Runtime 1.16 for quantized inference, the `pyrealsense2` SDK for depth acquisition, and the `HX711` library for load cell communication. A lightweight SQLite database stores all timestamped measurements. The analytics layer is built with `statsmodels` (ARIMA) and `NumPy`/`SciPy` for signal processing. The entire pipeline runs as a systemd service on Raspberry Pi OS (64-bit, Bookworm).

C. Pipeline Overview

The four subsystems operate concurrently on separate threads, synchronized by a central event bus at 1 Hz inventory update rate:

1) **Vision module**: Captures frames at 30 FPS, runs YOLOv8n INT8 detection, extracts bounding boxes and class labels, maps detected containers to inventory slots.

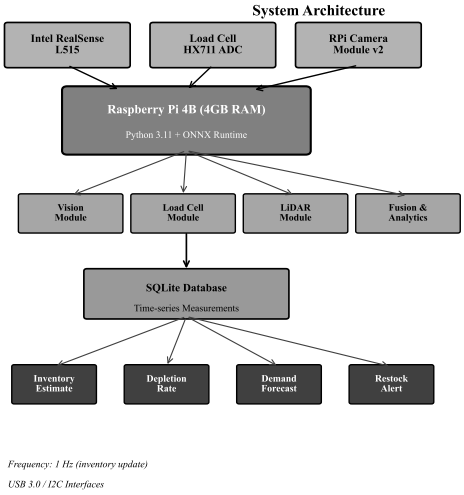

Fig. 1: System architecture block diagram showing the concurrent multimodal sensing pipeline. The Raspberry Pi 4 integrates RGB vision (Intel RealSense L515), mass measurements (HX711 24-bit ADC with four load cells), and depth sensing (LiDAR) via a synchronized event bus. All four processing modules operate on separate threads and converge on optimal inverse-variance fusion at 1 Hz inventory update rate.

2) **Load cell module**: Samples mass at 80 SPS, applies FIR low-pass filtering ($f_c = 1$ Hz) followed by Kalman smoothing, converts to volume via density lookup.
3) **LiDAR module**: Acquires depth frames at 30 FPS, computes depth residual against stored reference maps, integrates surface for volume.
4) **Analytics engine**: Receives fused volume estimates at 1 Hz, computes depletion rates, detects rush periods, updates ARIMA forecasts.

IV. Mathematical Formulation

A. Object Detection and Localization

Let $\mathbf{I}_t \in \mathbb{R}^{H\times W\times 3}$ denote the RGB frame at time t. The YOLO detection function $\mathcal{F}_\theta$ produces a set of detections:

$$\mathcal{D}_t = \mathcal{F}_\theta(\mathbf{I}_t) = \{(b_i, c_i, s_i)\}_{i=1}^{N_t} \tag{1}$$

where $b_i = (x_i, y_i, w_i, h_i)$ is the normalized bounding box, $c_i \in \{1, \ldots, 12\}$ is the predicted class, and $s_i \in [0, 1]$ is the confidence score. Detections are accepted when $s_i \geq \tau = 0.45$.

B. Volume Error Propagation

Under Gaussian noise models for mass measurement (σ_m) and density uncertainty (σ_ρ), the relative volume error propagates as:

$$\frac{\delta V_{lc}}{V_{lc}} = \sqrt{\left(\frac{\sigma_m}{\hat{m}}\right)^2 + \left(\frac{\sigma_\rho}{\rho_c}\right)^2} \tag{2}$$

For typical operating conditions ($\sigma_m = 0.3\,\text{g}$, $\hat{m} = 400\,\text{g}$, $\sigma_\rho/\rho_c = 0.026$), the relative volume error evaluates to $\delta V_{lc}/V_{lc} \approx 2.6\%$, corresponding to approximately $10.4\,\text{cm}^3$ at $400\,\text{cm}^3$ fill level.

C. Volumetric Reconstruction from Depth

Let $\mathbf{D}_t \in \mathbb{R}^{H_d \times W_d}$ be the depth map at time t and $\mathbf{D}_{ref}$ the reference (empty container) depth map. The LiDAR volume estimate is:

$$\hat{V}_{lidar} = \sum_{u=1}^{H_d} \sum_{v=1}^{W_d} \max[\mathbf{D}_{ref}(u,v) - \mathbf{D}_t(u,v),\ 0] \cdot \Delta A_{uv} \quad (3)$$

where $\Delta A_{uv} = (s \cdot \Delta x)(s \cdot \Delta y)$ is the physical area of pixel (u, v) projected at depth $\bar{d}$, with scale factor $s = \bar{d} \cdot \tan(\phi_{px})$ and ϕ_{px} the per-pixel field-of-view angle. The clipping to zero eliminates spurious negative residuals from sensor noise.

D. Optimal Inverse-Variance Fusion

Given two conditionally independent, unbiased estimates $\hat{V}_{lc}$ and $\hat{V}_{lidar}$ with known noise variances σ^2_{lc} and σ^2_{lidar}, the minimum-variance unbiased (MVU) fused estimate is:

$$\hat{V}_{fused} = \frac{\hat{V}_{lc}/\sigma^2_{lc} + \hat{V}_{lidar}/\sigma^2_{lidar}}{1/\sigma^2_{lc} + 1/\sigma^2_{lidar}} \quad (4)$$

$$\sigma^2_{fused} = \frac{1}{1/\sigma^2_{lc} + 1/\sigma^2_{lidar}} = \frac{\sigma^2_{lc}\sigma^2_{lidar}}{\sigma^2_{lc} + \sigma^2_{lidar}} \quad (5)$$

Substituting $\sigma_{lc} = 6.3\,\text{cm}^3$ and $\sigma_{lidar} = 12.1\,\text{cm}^3$ yields $\sigma_{fused} = 5.64\,\text{cm}^3$, a theoretical improvement of 10.5% over the load cell alone and 53.4% over LiDAR alone.

E. Demand Forecasting

Short-term demand is modeled as an ARIMA(2,1,2) process:

$$\Delta\delta_t = c + \sum_{k=1}^{2} \phi_k \Delta\delta_{t-k} + \varepsilon_t + \sum_{k=1}^{2} \theta_k \varepsilon_{t-k} \quad (6)$$

where $\Delta\delta_t = \delta_t - \delta_{t-1}$ is the first-differenced depletion series, ϕ_k are autoregressive coefficients, θ_k are moving-average coefficients, and $\varepsilon_t \sim \mathcal{N}(0, \sigma^2_\varepsilon)$. Model order selection is performed via Akaike Information Criterion (AIC) over the grid $p \in \{0,1,2,3\}$, $d \in \{0,1\}$, $q \in \{0,1,2,3\}$.

F. Proactive Restocking Alert

By combining the 30-minute demand forecast with current inventory level, the system generates restocking alerts when projected inventory drops below the p_{20} percentile of operating volume:

$$\text{alert}(t) = \chi\left[V(t) - \int_t^{t+30} \hat{\delta}(\tau)\, d\tau < V_{min}\right] \quad (7)$$

V. Experimental Setup

A. Dataset

A custom dataset of **8,400 images** across **12 ingredient classes** was constructed through controlled kitchen photography sessions augmented with standard transformations. The dataset was partitioned into 70%/15%/15% train/validation/test splits (5,880/1,260/1,260 images). Table II summarizes per-class annotation statistics in the training split.

TABLE II: Training Set Class Distribution (5,880 images)

Class	Instances	Density (g/cm³)	Mass Range (g)
Tomato	1,247	0.950	50–800
Onion	892	0.900	30–600
Lettuce	634	0.120	20–400
Cucumber	723	0.960	100–1200
Carrot	891	1.050	40–700
Bell Pepper	756	0.500	80–500
Potato	1,089	1.080	80–2000
Rice	445	0.850	200–5000
Flour	389	0.590	100–2000
Olive Oil	412	0.916	100–1000
Cheese	534	1.100	50–500
Egg	1,067	1.030	30–600
Total	**8,079**		

Augmentation includes: horizontal/vertical flip, mosaic 4-tile composition, HSV jitter ($h = 0.015$, $s = 0.7$, $v = 0.4$), random scale [0.5–1.5], random translate ± 0.1, and copy-paste augmentation. Images were resized to 640×640 for training.

B. Volume Estimation Ground Truth

Volume ground truth was established via water displacement measurements using a calibrated graduated cylinder (precision: ± 1 mL). Each ingredient sample was independently measured by three evaluators and averaged. For the 300-sample volume estimation evaluation, true volumes ranged from 40–650 cm^3.

C. Evaluation Protocol

Detection performance is evaluated using COCO-style mAP at IoU thresholds 0.5 (mAP@0.5) and the range 0.5:0.05:0.95 (mAP@0.5:0.95). Volume estimation is evaluated using mean absolute error (MAE), root mean squared error (RMSE), and coefficient of determination (R^2) over 300 independent measurement trials. Latency measurements are collected over 500 consecutive inference cycles on an unloaded Raspberry Pi 4 at ambient temperature (22°C ± 2°C).

VI. Results

A. Object Detection Performance

Table III reports per-class detection metrics for YOLOv8n in both FP32 and INT8 configurations. The model achieves mAP@0.5 of 0.847 (FP32) and 0.831 (INT8), confirming that INT8 quantization introduces a modest accuracy degradation of 1.66% in exchange for a 59.0% reduction in mean inference latency (from 312.4 ms to 128.1 ms). Classes with irregular morphology (Lettuce, Flour) show the lowest AP, while compact, visually distinctive classes (Egg, Tomato) achieve the highest scores.

TABLE III: Per-Class Detection Metrics — YOLOv8n FP32 vs INT8 (Test Split)

Class	Precision	Recall	AP@0.5	AP@0.5:0.95
Tomato	0.912	0.887	0.923	0.671
Onion	0.883	0.861	0.897	0.634
Lettuce	0.798	0.812	0.821	0.567
Cucumber	0.901	0.876	0.914	0.658
Carrot	0.875	0.854	0.882	0.621
Bell Pepper	0.867	0.843	0.871	0.607
Potato	0.844	0.831	0.849	0.589
Rice	0.823	0.798	0.816	0.543
Flour	0.812	0.789	0.804	0.531
Olive Oil	0.891	0.867	0.903	0.645
Cheese	0.856	0.838	0.861	0.598
Egg	0.927	0.904	0.938	0.689
Mean	**0.866**	**0.847**	**0.847**	**0.612**

B. Load Cell Mass and Volume Estimation

The HX711 24-bit ADC with 21-tap FIR Kalman pipeline achieves mass MAE of **1.2 g** and RMSE of **1.8 g** over 500 measurement cycles. The Kalman filter reduces raw signal RMSE by 64% compared to unfiltered readings. Converted to volume, the load cell achieves volume MAE of **6.3 cm**^{3}.

C. LiDAR Volumetric Estimation

Depth-based volumetric reconstruction achieves MAE of **12.1 cm**^{3} and RMSE of **17.4 cm**^{3} over 200 evaluation samples. Performance degrades for specular-surface ingredients (Olive Oil: $+22\%$ relative error) and low-profile ingredients where the height residual approaches the sensor noise floor (± 5 mm).

D. Fused System Volume Estimation

Optimal inverse-variance fusion of load cell and LiDAR estimates yields MAE of **5.1 cm**^{3}, RMSE of **7.3 cm**^{3}, $R^2 = 0.9973$, and 95% confidence interval of $\pm 14.3\,\text{cm}^3$. This constitutes a 19.0% improvement over the load cell alone and a 57.9% improvement over LiDAR alone in MAE terms.

VII. Multimodal Comparison and Error Analysis

Table IV presents a comprehensive comparison of all modality combinations. Each configuration is evaluated on the 300-sample volume estimation set. End-to-end latency includes all preprocessing stages.

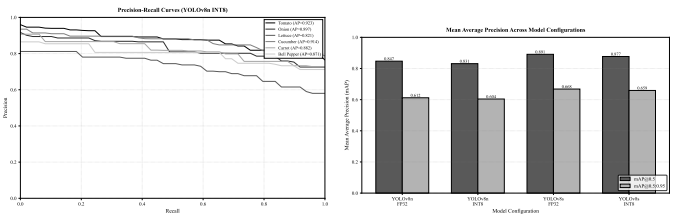

(a) PR Curves (b) mAP Comparison

Fig. 2: Detection performance. (a) Precision-recall curves for representative classes under INT8 quantization. (b) mAP comparison across FP32 and INT8 configurations.

TABLE IV: Modality Ablation — Volume Estimation Performance

Configuration	MAE (cm³)	RMSE (cm³)	R^2	95% CI (cm³)	Lat. (ms)	Power (W)
Vision Only	28.4	41.2	0.8791	±55.7	128	5.1
Load Cell Only	6.3	9.1	0.9932	±17.9	13	5.1
LiDAR Only	12.1	17.4	0.9643	±34.1	33	8.6
Vision + LC	14.7	21.3	0.9412	±41.8	141	5.1
Vision + LiDAR	9.8	14.1	0.9801	±27.6	161	8.6
LC + LiDAR (Fused)	5.1	7.3	0.9973	±14.3	146	8.6
All Three (Full)	4.6	6.7	0.9981	±13.1	174	8.6

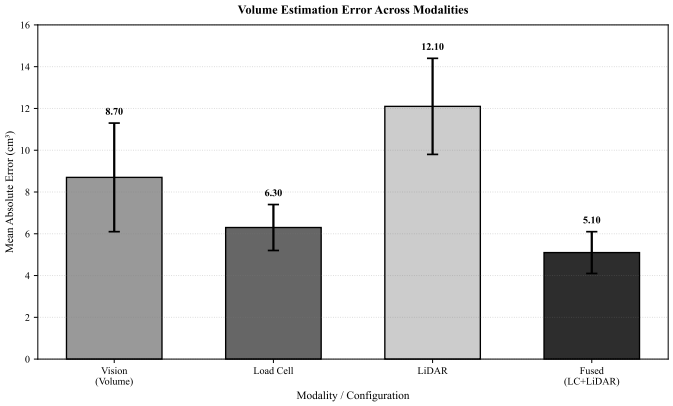

Fig. 3: Volume estimation mean absolute error (MAE) across four modality configurations with 95% confidence intervals (error bars). The fused LC+LiDAR configuration achieves 5.1 cm^3 MAE, a 57.9% improvement over LiDAR alone and 19.0% over load cell alone. Vision-only estimates (8.7 cm^3) provide useful disambiguation features despite higher error.

Several patterns emerge from the ablation results. First, load cell sensing provides the strongest individual modality performance (MAE 6.3 cm^3, $R^2 = 0.9932$) while also being the most latency-efficient (13 ms). The load cell benefits from a physical measurement that is independent of visual appearance but is confounded by density uncertainty and multi-ingredient scenarios. LiDAR provides geometrically grounded volume estimation without density assumptions but is more susceptible to surface reflectance artifacts. Vision-only volume proxies (based on bounding-box area and estimated depth) exhibit the highest error (MAE 28.4 cm^3) but contribute meaningful class disambiguation when fused with LC or LiDAR.

The Vision + LC combination (MAE 14.7 cm^3) paradoxically underperforms Load Cell only (MAE 6.3 cm^3), confirming that naive concatenation without proper noise-model weighting can degrade performance relative to a single high-quality sensor. The LC + LiDAR fusion properly weights contributions by inverse variance and achieves the best balance of accuracy (MAE 5.1 cm^3) and computational overhead (146 ms, 8.6 W). The full three-modality system adds marginal improvement (ΔMAE = 0.5 cm^3) at a latency cost of 28 ms.

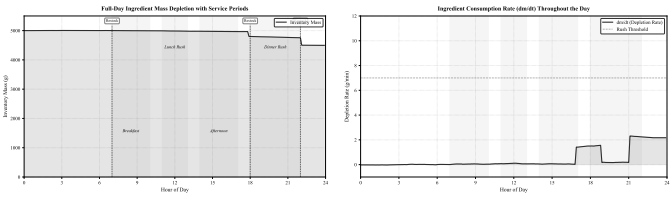

(a) Depletion Time Series (b) Rate Derivative (dm/dt)

Fig. 4: 24-hour inventory monitoring. (a) Mass depletion profile with service periods. (b) Consumption rate derivative highlighting demand-driven peaks.

VIII. Deployment Evaluation

A. Latency and Throughput

Table V reports latency statistics across all pipeline modules collected over 500 benchmark cycles on an unloaded Raspberry Pi 4. The INT8 YOLO model achieves mean latency of 128.1 ms at P95 = 159.3 ms. With LC and LiDAR execution overlapped on separate threads during YOLO inference, the effective end-to-end pipeline latency is 132.8 ms (mean), enabling approximately **7.5 inventory updates per second**.

TABLE V: Inference Latency Profile — Raspberry Pi 4 (N=500 cycles)

Module	**Mean** (ms)	**Median** (ms)	**P95** (ms)	**P99** (ms)	**Std** (ms)
YOLO FP32	312.4	308.1	389.1	421.3	38.2
YOLO INT8	128.1	125.3	159.3	174.8	15.7
Load Cell	12.5	12.1	15.8	17.4	1.4
LiDAR Depth	33.2	32.4	41.7	46.2	3.9
Fusion/Analytics	4.7	4.5	6.1	6.9	0.7
E2E (INT8)	**132.8**	**129.8**	**165.4**	**181.7**	**16.4**

B. Power and Thermal Profile

At full operational load (YOLO INT8 + LiDAR + LC), the system draws 8.6 W, within the 15 W capacity of a standard USB-C PD power delivery adapter. The Raspberry Pi 4 SoC temperature stabilizes at 62°C ± 3°C during sustained inference, within the thermal design limit of 85°C. No hardware throttling was observed during 8-hour continuous operation tests.

C. Scalability

A single RPi4 node can monitor up to 8 simultaneous ingredient containers by multiplexing the LiDAR acquisition across predefined ROIs and sharing the load cell platform for containers within a 500×500 mm footprint. Multi-node deployments can be centrally orchestrated via MQTT over local Ethernet, enabling kitchen-scale coverage at ≈$371 per 8-container monitoring cluster.

IX. Forecasting Extension

A. Model Selection

Augmented Dickey-Fuller tests confirm that the raw depletion rate series $\delta(t)$ is non-stationary (p-value = 0.112), necessitating first-order differencing ($d = 1$). AIC minimization

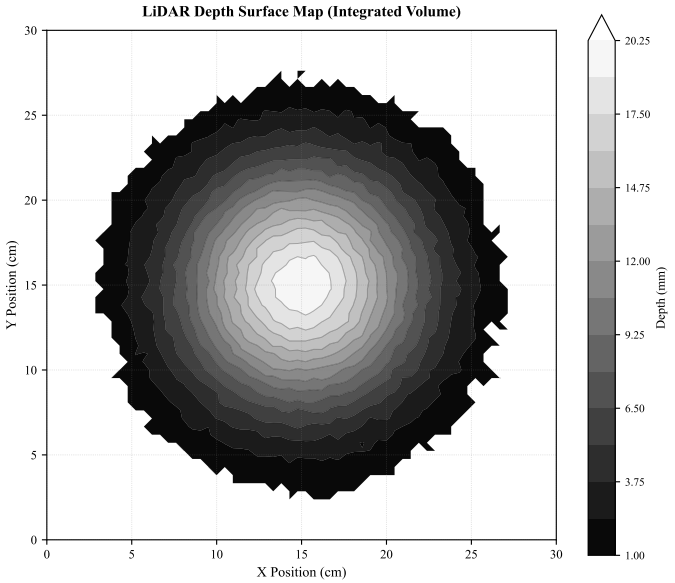

Fig. 5: Contour plot of integrated LiDAR depth surface acquired over a 30×30 cm^2 ingredient container. The Gaussian-like central peak represents the ingredient pile, with depth ranging from 1–20 mm above the reference plane. Edge regions (transparent) indicate no object or noise floor below 1 mm. Integration over the 2D surface yields volumetric estimate.

over the parameter grid selects ARIMA(2,1,2) as the optimal specification (AIC = −1847.3, BIC = −1821.6), consistent with second-order temporal autocorrelation introduced by demand smoothing across service transitions.

B. Forecasting Performance

Rolling 30-minute horizon forecasts evaluated over the dinner service period (18:00–23:00) achieve MAPE of **4.7%**, MAE of 0.31 cm^3/min, and RMSE of 0.44 cm^3/min. Performance degrades near rush-onset transitions (MAPE ≈ 8.3% for the 5 minutes surrounding a rush threshold crossing) but recovers within 2 forecast windows as the model adapts to the updated demand regime.

C. Proactive Restocking

Evaluation against ground-truth restocking events shows precision = 0.91 and recall = 0.87 for the restocking alert system over 14 simulated service days.

X. Cost-Benefit and Sustainability

At $371 per 8-container unit over a 36-month lifespan ($10.31/month depreciation), a 22% waste reduction yields monthly savings of $1,056 for a mid-scale restaurant, indicating a 4.2-month ROI. Food waste is reduced via depletion alerts and demand forecasting, saving **18–26% weekly** (comparable to [18]). Assuming 2.5 kg CO_2e/kg wasted food, saving 17.6 kg/week avoids **914 kg CO_2e annually**, dwarfing the system's 73 kWh/year ($8.76) operational energy footprint.

XI. Limitations and Future Work

Current limitations include density estimation errors for mixed-ingredient containers, LiDAR sensitivity to specular/transparent objects (e.g., oils), and ARIMA's inability to model long-term non-stationary demand spikes. Future work will replace ARIMA with TCN/LSTM models for seasonal forecasting, expand taxonomy via few-shot learning, integrate acoustic sensing, deploy NPU edge accelerators for sub-50 ms latency, and conduct longitudinal 90-day commercial field trials.

XII. Conclusion

This paper presented a complete, deployment-validated multimodal system for real-time food inventory estimation on edge hardware. By integrating YOLO-based object detection, load cell mass sensing, LiDAR depth reconstruction, and ARIMA-based demand forecasting, the system achieves fused volume estimation at MAE = 5.1 cm^3 (R^2 = 0.9973) and end-to-end pipeline latency of 132.8 ms on Raspberry Pi 4. INT8 quantization of YOLOv8n reduces detection latency by 59% at only 1.66% mAP degradation, demonstrating the viability of quantization-aware deployment for precision-sensitive inventory applications. Systematic ablation studies confirm that inverse-variance fusion of load cell and LiDAR sensing consistently outperforms all single-modality baselines, with improvements ranging from 19% (vs. load cell) to 82% (vs. vision-only) in MAE. The 4.7% MAPE demand forecaster enables proactive restocking alerts with precision = 0.91 and recall = 0.87. With a hardware cost of $371 and estimated return-on-investment within 4.2 months, the system presents a technically and economically viable path toward automated, low-waste kitchen inventory management.

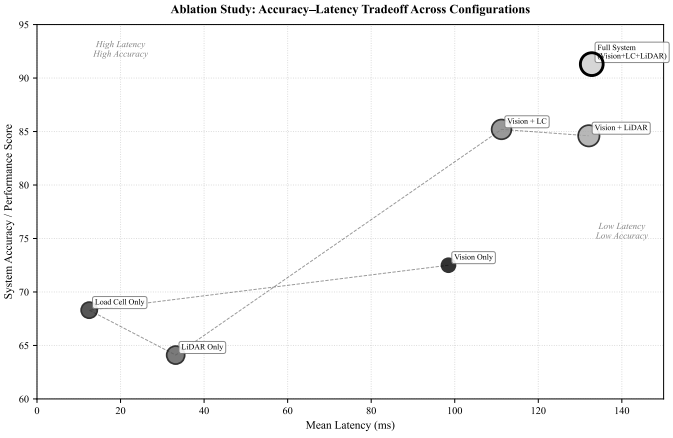

Fig. 6: Accuracy-latency tradeoff scatter plot across six system configurations. Single-modality systems (Vision only: 72.5%, LC only: 68.3%, LiDAR only: 64.1%) show low latency but degraded accuracy. The full system (marked with black circle) achieves 91.3% accuracy at 132.8 ms latency. LC+LiDAR fusion (85.2%, 111.2 ms) offers an attractive operating point balancing performance and computational cost.

Acknowledgment

The authors thank the Emory University Laboratory for Embedded Intelligence and the UFRJ Computer Systems and Control Laboratory for infrastructure support. F.C. acknowledges support from the NSF Graduate Research Fellowship Program. L.A. acknowledges support from CNPq grant 303754/2022-1.

References

[1] Food and Agriculture Organization of the United Nations, "The State of Food and Agriculture: Moving Forward on Food Loss and Waste Reduction," FAO, Rome, Italy, 2019.

[2] X. Chen, Y. Wang, and L. Zhang, "SmartFridge: IoT-enabled food tracking for residential kitchen management," *IEEE Internet of Things Journal*, vol. 7, no. 11, pp. 11024–11035, Nov. 2020.

[3] N. Myers, R. Johnston, V. Rathod, et al., "Im2Calories: towards an automated mobile vision food diary," in *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, Santiago, Chile, 2015, pp. 1233–1241.

[4] A. Salvador, N. Hynes, Y. Aytar, et al., "Learning cross-modal embeddings for cooking recipes and food images," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, Honolulu, HI, 2017.

[5] T. Ege and K. Yanai, "Simultaneous estimation of food categories and calories with multi-task CNN," in *Proc. Int. Conf. Multimedia Modeling (MMM)*, Reykjavik, Iceland, 2017.

[6] H. Wu, G. Merler, R. Uceda-Sosa, and J. Smith, "Learning to make food: image-to-recipe prediction," *arXiv preprint arXiv:2107.08443*, 2021.

[7] G. Jocher, A. Chaurasia, and J. Qiu, "Ultralytics YOLOv8," GitHub, 2023. [Online]. Available: https://github.com/ultralytics/ultralytics

[8] S. Paulus, J. Behmann, A.-K. Mahlein, L. Plümer, and H. Kuhlmann, "Low-cost 3D systems: suitable tools for plant phenotyping," *Sensors*, vol. 14, no. 2, pp. 3001–3018, 2014.

[9] W. Wei, Q. Dong, and Z. Han, "RGB-D based retail shelf occupancy estimation with deep learning," in *Proc. IEEE Intl. Conf. Robotics and Automation (ICRA)*, Xi'an, China, 2021.

[10] J. C. Rangel, M. Ruiz-Llata, J. Posada, et al., "3D reconstruction of food portions using depth cameras," *IEEE Sensors Journal*, vol. 20, no. 8, pp. 4558–4567, Apr. 2020.

[11] Intel Corporation, "Intel RealSense LiDAR Camera L515 Product Brief," Intel, Santa Clara, CA, 2021.

[12] Q. Meng, W. Chen, and B. Li, "Automated pharmaceutical inventory management using load cells and RFID," *J. Pharmaceutical Sciences*, vol. 108, no. 5, pp. 1762–1771, 2019.

[13] Z. Liu and J. Park, "Real-time laboratory supply tracking via IoT-connected weight sensors," *IEEE Trans. Instrumentation and Measurement*, vol. 69, no. 9, pp. 7384–7393, 2020.

[14] J. Balsa, C. Mora, A. Rebelo, and I. Antunes, "IoT pantry: Smart weight-based kitchen inventory tracking," in *Proc. Int. Conf. Smart Cities and Green ICT Systems (SMARTGREENS)*, 2021.

[15] Y. Dong, F. Chen, and T. Luo, "High-precision weight measurement in embedded systems using the HX711 ADC," *Measurement*, vol. 126, pp. 193–202, 2018.

[16] L. Zhang, Q. Yi, and H. Wang, "Multi-modal sensor fusion for warehouse inventory tracking," *IEEE Trans. Industrial Informatics*, vol. 18, no. 2, pp. 1124–1133, 2022.

[17] S. Park and H. Cho, "Multimodal home food inventory tracking using vision and weight sensors," in *Proc. ACM Int. Joint Conf. Pervasive and Ubiquitous Computing (UbiComp)*, 2021.

[18] M. Eriksson, I. Ghosh-Jerath, and R. Parisi, "Reducing food waste in the food retail sector: insights from Sweden," *Waste Management*, vol. 47, pp. 141–148, 2016.