Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Towards Model-Ready AI for Apple Shelf-Life Estimation: A Fruit-Proximal Sensing Framework

Pascal Nsunba Kanuayi
University of the Western Cape https://orcid.org/0009-0005-7995-5487
Jean Frederic Isingizwe Nturambirwe
fisingizwe@uwc.ac.za
University of the Western Cape https://orcid.org/0000-0002-1794-7343
Clement Nyirenda
University of the Western Cape https://orcid.org/0000-0002-4181-0478

Research Article

Keywords: Sustainable agricultural AI; Post-harvest shelf-life estimation; Fruit-proximal sensing; Internet of Things; Model readiness; Cold chain monitoring

Posted Date: November 17th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11742556/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Towards Model-Ready AI for Apple Shelf-Life Estimation: A Fruit-Proximal Sensing Framework

Pascal Nsunba Kanuayi[1,3][0009−0005−7995−5487],
Jean Frederic Isingizwe Nturambirwe[2][0000−0002−1794−7343]*, and
Clement Nthambazale Nyirenda[1,2][0000−0002−4181−0478]

[1] Department of Computer Science, University of the Western Cape, Robert Sobukwe Rd, Bellville, Cape Town, 7535, South Africa
3963848@myuwc.ac.za
[2] eResearch Office, University of the Western Cape, Robert Sobukwe Rd, Bellville, Cape Town, 7535, South Africa
fisingizwe@uwc.ac.za, cnyirenda@uwc.ac.za
[3] African-German Centre for Sustainable and Resilient Food Systems and Applied Agricultural and Food Data Science (UKUDLA), University of the Western Cape, Cape Town, South Africa

Abstract. Reliable apple-shelf-life AI requires exposure histories linked to defined biological outcomes. This paper presents a serviceable, three-dimensional printed smart apple and a framework for preparing fruit-proximal measurements for predictive evaluation. The device combines environmental and inertial sensing, timestamped local storage, and intermittently connected data transfer. A farm-to-packhouse trial recorded 470 BME280, 468 BME680, and 466 inertial observations across a nominal 79-minute journey without gaps exceeding 60 s. Cold-storage analyses include a 10.8-hour Device 1 cooling segment, a continuous 24.8-hour Pi1 record, and a duty-cycled 15.1-hour Pi2 record with 27.3–27.5% observed-interval coverage. These records quantify acquisition performance and channel-specific environmental responses. A six-point model-readiness gate specifies outcome definition, exposure linkage, temporal integrity, independent units, technical quality, and group-level validation. Firmness provides the quantitative endpoint for the proposed remaining-life workflow, with cultivar- and protocol-specific acceptance criteria. The empirical contribution is deployment and data-quality evidence; linked biological prediction is the next evaluation stage. The framework provides a reproducible route from field sensing to grouped comparison against static and temperature-only baselines.

Keywords: Sustainable agricultural AI; Post-harvest shelf-life estimation; Fruit-proximal sensing; Internet of Things; Model readiness; Cold-chain monitoring

* Corresponding author.

1 Introduction

Fresh apples remain biologically active after harvest. Respiration, transpiration, ripening, senescence, mechanical injury, and microbial activity continue throughout handling, transport, and storage. The rate of quality loss depends on the magnitude, duration, and sequence of exposures rather than on a single measurement at dispatch or arrival. Temperature and relative humidity affect physiological deterioration and water loss, while impact and vibration can initiate bruising whose visible consequences emerge later [14, 18]. Consequently, consignment-level assessment benefits from combining nominal storage settings with the magnitude, duration, and sequence of actual exposures.

Connected sensing offers a way to retain this exposure history. However, most cold-chain measurements describe the air at a wall, vehicle, or room-level location. Fruit within a bin, crate, or pallet may experience a different microclimate and a distinct mechanical trajectory. Real-time monitoring literature identifies temperature and humidity as the dominant measurements and documents uneven coverage across commodities and supply-chain stages [5]. For artificial-intelligence (AI) modelling, these records need calibration, deployment context, an explicit biological outcome, and evaluation at the relevant independent unit.

This study builds on the cold-store data-collection and querying research of Nsunba et al. [13], which used the Actor Model to coordinate sensor requests, payload decoding, MongoDB storage and querying. The present paper extends this research direction through an apple-shaped instrument, fruit-proximal transport and storage evaluation, and readiness requirements for linking environmental histories to biological outcomes.

This paper investigates that boundary through a "smart apple": a three-dimensional (3D)-printed, apple-shaped sensing platform intended to travel with commercial apples. It combines environmental and inertial sensing with an embedded computer, local persistence, and opportunistic remote transfer. The research question is:

> How can a fruit-proximal IoT platform be designed and evaluated so that its field exposure records constitute model-ready evidence for defensible apple-shelf-life AI?

Embedded systems provide the instrument, post-harvest science defines the outcome, data engineering preserves exposure provenance, and machine-learning methodology specifies the evaluation unit. Enclosure response, deployment context, and biological sampling therefore jointly determine model readiness.

The contributions are an implemented sensing architecture, reproducible transport and storage evidence, a six-point model-readiness gate, and a grouped evaluation protocol. The field records establish operational feasibility and identify calibration and continuity requirements for prediction with linked real-apple observations.

2 Background and Research Gap

2.1 Post-Harvest Quality and Fruit-Shaped Sensing

Quality includes firmness, colour, soluble solids, acidity, mass loss, bruising, decay and sensory acceptability. The acceptance threshold depends on cultivar, intended use and stakeholder. Remaining shelf life is the interval before a declared endpoint is reached under stated future conditions. Temperature affects deterioration rates, humidity affects water loss, and mechanical exposure can initiate damage. Impact interpretation also depends on fruit maturity, firmness, contact geometry and surface [18].

IoT combines sensing, communication and analytics [15]. Local persistence retains the scientific record through intermittent connectivity. Instrumented spheres identify severe handling operations [20]; Southern African smAvo/smaTo platforms extend this approach with environmental, inertial and contextual measurements in 3D-printed fruit surrogates [3].

Shape, mass, centre of gravity, shell material and thermal response determine how a surrogate represents real fruit. These require controlled characterisation. One device records its own trajectory and microenvironment; multiple batches and journeys supply operational and biological replication. Sampling density supplies temporal resolution. These roles guide both device placement and the eventual predictive evaluation.

2.2 Shelf-Life AI and Digital Twins

Controlled studies show that apple quality and shelf life can be modelled from storage conditions and measured quality attributes [4, 21]. Machine learning can capture nonlinear relationships, but its usefulness depends on the breadth and validity of the underlying experimental data [11]. Digital twins offer a bridge from sensed conditions to a dynamic representation of produce quality [6]. Product-centred twins have been used to quantify quality trade-offs in refrigerated fruit chains [17], and recent work combines sensor histories with data-driven quality-loss models [8, 22].

The central research challenge is to establish a traceable path from apple-shaped, field-deployed sensing to representative exposure histories and apple-specific prediction. In this setting, adjacent sensor rows share the same exposure history and outcome. Randomly distributing them between training and test sets creates dependence and optimistic performance. Blocked validation is preferred when data contain temporal, spatial, or hierarchical structure [16]; more broadly, data leakage is a documented source of irreproducible claims in applied machine learning [10]. This paper treats readiness for inference as an explicit research result.

3 Interdisciplinary Research Design

The study follows design-science logic: construct an artefact for an operational problem, evaluate it in the intended environment, and use the findings to refine

its design [9]. The linked cycles are device construction, field evaluation, and specification of predictive-data requirements.

3.1 Requirements and Evidence Boundaries

The functional requirements were a serviceable apple-shaped device, local environmental and motion sensing, UTC timestamps, device identity, local persistence and network-dependent transfer. Field evaluation examined recovery and operation during handling, transport and cold storage.

Four evidence levels guide evaluation: manufacturability and assembly; timestamped acquisition, persistence and recovery; controlled thermal and mechanical representativeness against real apples; and predictive performance on unseen biological or deployment groups. This paper evaluates the first two levels and specifies the requirements for the latter two.

3.2 Field-Deployment Design

Two operational stages were observed. During farm-to-packhouse transport, one smart apple was positioned at the top centre of a commercial pallet and travelled for approximately 79 minutes. A Wi-Fi modem in the vehicle supported near-real-time transfer, while the device retained local records. This deployment provided a focused, single-device evaluation of integration into routine transport.

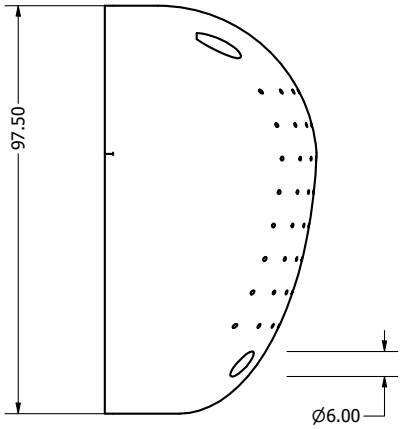
During storage in the single evaluated cold room, three devices were deployed: one at a top-centre position and two at middle-centre positions, with periodic battery replacement. The duty-cycled device, archived as Pi2, occupied a middle-centre position. Device 1, Pi1 and Pi2 retain their archive identities in the analysis. Their temperature histories describe the recorded storage windows and operating configurations. The present evidence concerns one transport journey and these storage windows. Biological model fitting uses the separately supplied laboratory records in the next evaluation stage.

Normal operations continued around instrument insertion, checking and retrieval. Each trace is interpreted within its recorded window, configuration and available maintenance context.

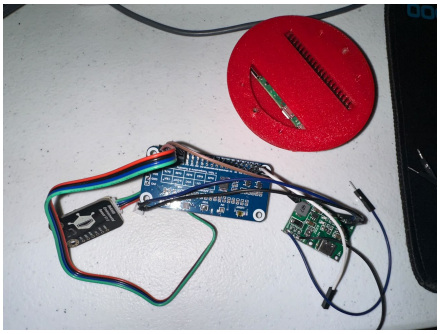
4 Smart Apple Artefact

4.1 Physical and Sensing Architecture

A two-part 3D-printed enclosure holds environmental and inertial sensors, a Raspberry Pi Zero 2 W, a battery and modular supports. Daily binary files provide local persistence; a remote document database provides visibility when connected. The nominal enclosure height is 97.5 mm and its profile radius is 48.65 mm. Joining lips and mounting bosses permit service access. Distributed apertures support air exchange, with response lag, robustness and electronics protection to be assessed together. Supports restrain the boards and battery so that internal movement is controlled.



(a) Dimensioned enclosure profile (mm)



(b) Environmental sensing and power wiring

Fig. 1. Enclosure CAD profile with nominal dimensions and implemented internal electronics. The profile provides an engineering size reference.

Figure 1 shows the enclosure profile and internal electronics. The dimensioned profile and assembly document the implemented enclosure and its nominal size; controlled thermal and mechanical tests will characterize its response against real fruit.

The sensor inventory includes a Waveshare Environment Sensor HAT and successive BME680/BME688-class environmental modules. The HAT integrates a BME280 and an ICM20948 nine-axis inertial sensor [19]; the later BME688 module is designed as a BME680-compatible replacement [7]. The streams contain temperature, relative humidity, gas resistance and inertial channels. The accelerometer range is $\pm 2g$; its approximately 10-second recording cadence defines the sampled-motion resolution. Pressure fields are provisioned in the message schema and were disabled in the analysed deployment configuration. The archived BME680 channel name is retained across the BME68x hardware history.

4.2 Acquisition, Persistence, and Communication

The acquisition application is written in C# and executes enabled sensor loops as asynchronous tasks. Readings are timestamped in UTC and serialised as length-delimited Protocol Buffer messages. Separate daily binary files are maintained for each sensor. The local record is the recoverable scientific source when the remote service is unreachable.

When database output is enabled, corresponding records are inserted into sensor-specific collections through a MongoDB-compatible interface. Timeouts

prevent an unavailable network from blocking acquisition indefinitely, and reconnection attempts are throttled. A hybrid strategy is important because continuous streaming is neither required nor dependable throughout an orchard, vehicle, packhouse, and cold room. The analysed environmental and inertial streams were sampled at approximately 10-second intervals; event-triggered high-rate acquisition is a future extension.

Analysis retains timestamps, source identity, sensor configuration, application revision and available position and service metadata. Reconciliation separates network interruptions from local acquisition gaps.

4.3 Environmental Calibration

Temperature and relative humidity use a single-point calibration against a supplied reference of 26°C and 52% RH. The first 360 observations per logged channel define the nominal first hour at the stated 10-second cadence; the calibration log records sample order. For raw measurement $x_{s,j}$ and reference r_x, the unit-slope correction is

$$x_{s,j}^{\mathrm{corr}} = x_{s,j} + b_{s,x}, \qquad b_{s,x} = r_x - \overline{x}_s. \tag{1}$$

Temperature offsets are -5.795°C for BME280 and $+2.482$°C for BME680; RH offsets are $+20.728$ and -3.116 percentage points, respectively. Calculations use the full-precision coefficients. The same logged-channel coefficients are applied throughout the deployment summaries and figures, with raw values, calibration identity and quality flags retained. The reference experiment defines the fitted condition; field values outside the observed fitting ranges are flagged as extrapolated, and corrected RH outside $[0, 100]$% is retained without clipping. Multi-condition, device-specific reference measurements extend this calibration across the storage domain.

4.4 Power as a Data-Quality Variable

The platform used a 5000 mAh, 3.7 V lithium-polymer battery. It recorded the complete nominal transport window, although depletion occurred during extended operation. System load, overnight charging continuity and initial state of charge are factors for controlled endurance testing. Subsequent adaptations restricted processor demand, enabled Wi-Fi power saving, reduced sampling demand and introduced experimental ESP32-C3-controlled switching. Duty cycling includes orderly filesystem shutdown. Comparative tests record battery, initial voltage, charging history, workload, firmware and duration.

5 Field Findings and Design Learning

Figure 2 documents the actual top-centre and middle-centre sensor placements among commercial apples. During deployment, the closed devices travelled with



(a) Top-centre sensor placement

(b) Middle-centre sensor placement

Fig. 2. Position-labelled smart apple placements among commercial fruit.

the produce rather than being mounted on a vehicle or room wall. Their records are therefore fruit-proximal, but each represents its selected position rather than the complete batch.

The transport trial established insertion, loading, road movement and vehicle-provided connectivity in commercial operation. Storage records extend the evaluation to continuous and duty-cycled acquisition. Local and remote records are reconciled by device, sensor and timestamp to distinguish network outages from acquisition gaps. Battery-service events and operating windows retain their own identities. The duty-cycled middle-centre record supplies position context for the storage-temperature analysis.

6 Quantitative Deployment Results

The analysis uses 18 unique local binary files across four declared windows. The decoder validates Protocol Buffer messages and records recovery locations, accepted counts and source SHA-256 checksums. Timestamp sorting precedes window selection. Coverage is the sum of adjacent positive intervals of at most 60 s divided by the declared window duration; longer gaps retain unobserved exposure. Source identities and per-window counts are verified against the versioned deployment audit. Means and standard deviations summarize the recorded observations within each window. Temperature and RH values below use the calibration in Section 4.3; gas resistance and inertial channels retain their recorded scale.

6.1 Farm-to-Packhouse Exposure

The nominal journey from 10:50 to 12:09 SAST on 25 May 2026 contains 470 BME280, 468 BME680, and 466 ICM20948 observations. Median sampling intervals were 10.069, 10.135, and 10.168 s, respectively, and every stream covered the complete field-note window without an interval exceeding 60 s. Table 1 summarises the fruit-proximal exposure measurements.

Table 1. Descriptive results for the nominal 79-minute transport window; temperature and RH include the single-point calibration.

Stream	Variable (unit)	n	Mean	Range
BME280	Temperature (°C)	470	18.24	16.88–21.38
BME280	Relative humidity (%)	470	64.35	61.52–66.59
BME680	Temperature (°C)	468	22.23	21.23–28.65
BME680	Relative humidity (%)	468	63.85	54.78–65.84
BME680	Gas resistance (kΩ)	468	11.22	8.00–12.55
ICM20948	Resultant g-force (g)	466	1.013	0.453–1.577

Figure 3 shows the environmental transition. Both temperature channels declined early and then stabilised, while their humidity channels changed in opposing directions before reaching stable levels. The shared timing records a common transition. Corrected BME680 mean temperature is 3.99°C above BME280, while their mean RH values differ by 0.50 percentage points. Resultant g-force averaged 1.013 g and reached 1.577 g, capturing the sampled mechanical profile of the same journey. These approximately 10-second samples retain gravity in the resultant and characterize sampled motion; brief shocks require higher-rate acquisition.

6.2 Cold-Environment Transition

Device 1 supplied a shared continuous BME280/BME680 segment from approximately 11:27 to 22:13 SAST on 2 June 2026, a duration of 10.8 h. BME280 temperature ranged from 1.85 to 13.92°C and averaged 2.59°C; BME680 temperature ranged from 6.70 to 22.91°C and averaged 7.49°C. Both channels rose briefly near the start, then cooled and remained relatively steady (Figure 4a). Mean corrected RH was 84.19% and 87.28%, respectively. The BME280 summary retains 119 values above 100% RH with physical-bounds flags. The trajectory demonstrates sustained acquisition during the transition into cold storage.

6.3 Additional Deployment Records and Acquisition Coverage

Table 2 extends the Device 1 analysis with Pi1 and Pi2. Pi1 spans 13:58 SAST on 5 June to 14:46 on 6 June and provides near-complete coverage in all three streams. Pi2 spans 00:23 to 15:32 SAST on 5 June (source files dated 4–5 June in UTC). Its enabled two-minute acquisition period and shutdown configuration produce an intermittent record. A further 175.6-minute interruption contributes to its 27.3–27.5% coverage. Cumulative exposure features integrate observed intervals only. The records describe different dates and operating configurations within the same cold room.

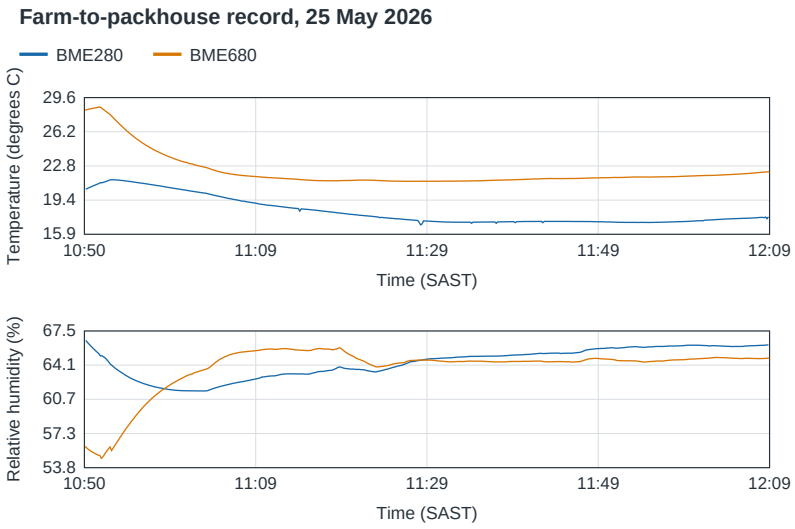


Fig. 3. Temperature and relative-humidity records during the nominal farm-to-packhouse journey (SAST, UTC+2), with the single-point calibration applied.

6.4 Storage-Temperature Variation and Channel Response

Figure 4 shows the corrected temperature histories. After the initial cooling transition, Pi1 retains a relatively steady record with short fluctuations. Its full-window BME280 and BME680 ranges are 1.23–14.32 and 9.45–23.67°C, respectively. The duty-cycled Pi2, located at middle centre, records a warmer and more variable history: 11.24–23.95°C for BME280 and 17.66–43.41°C for BME680, with standard deviations of 1.55 and 5.02°C. Its gaps remain visible rather than being interpolated. These histories retain position, date and operating mode as joint deployment context.

Corrected BME680 mean temperature exceeds BME280 by 4.91°C in Device 1, 8.02°C in Pi1 and 9.40°C in Pi2. The additive calibration shifts channel levels while preserving within-channel variation. Bosch identifies PCB temperature, self-heating and ambient temperature as contributors to BME680 output [2, Section 1.5]. Co-located reference measurements across storage conditions will characterize channel response, electronics heating and enclosure airflow. Gas resistance retains its electrical-response interpretation pending condition- and gas-specific calibration.

7 The Model-Readiness Framework

The field study motivates a gate between sensor deployment and predictive modelling. Applied before model fitting, the gate defines the evidence required for a

Table 2. Storage windows in the evaluated cold room. Temperature and RH means include the single-point calibration; coverage excludes gaps exceeding 60 s.

Archive	Stream	n	Span (h)	Coverage (%)	Mean T (°C)	Mean RH (%)
Device 1	BME280	3,853	10.776	100.0	2.59	84.19
Device 1	BME680	3,825	10.776	100.0	7.49	87.28
Device 1	ICM20948	3,236	10.776	85.1	–	–
Pi1	BME280	8,878	24.802	100.0	1.90	75.33
Pi1	BME680	8,820	24.802	100.0	9.92	65.63
Pi1	ICM20948	8,764	24.802	100.0	–	–
Pi2	BME280	1,504	15.137	27.3	20.39	64.23
Pi2	BME680	1,502	15.137	27.5	29.79	41.15
Pi2	ICM20948	1,493	15.137	27.3	–	–

meaningful numerical performance claim. Table 3 gives the six criteria and their quality-assurance purpose.

The archive provides reproducible decoding, source provenance, verified sampling windows and quantified continuity. Table 3 also requires calibration, physical-device/position linkage, biological endpoints and usable independent groups. Each is assessed separately before sensor-linked fitting.

7.1 Outcome Definition and Exposure Linkage

The proposed workflow uses a quantitative firmness endpoint and a complementary probability of remaining acceptable through a requested horizon. Here, *quality* denotes measured attributes, *acceptability* denotes passing a declared criterion, and *marketability* denotes the wider commercial decision, including decay and appearance. *Remaining shelf life* denotes days from prediction time to loss of the selected acceptability criterion under stated future conditions.

For the firmness workflow, loss occurs when measured flesh firmness F_i reaches a documented limit $F_{\min}$ on the same instrument scale:

$$t_{F,i} = \inf\{\tau : F_i(\tau) \le F_{\min}\}, \qquad RSL_i(t) = \max(0, t_{F,i} - t). \tag{2}$$

Three published endpoint scenarios are retained: 55 N desired firmness for Gala Schniga SchniCo Red under simulated shipping and 25°C shelf-life conditions, measured on peeled flesh using an 11-mm penetrometer [12]; 53 N for 'Luiza' storage-life assessment; and 62 N for its market-delivery criterion [1]. Each scenario retains its cultivar, treatment, protocol and decision purpose. Applying one requires a compatible response scale and a recorded threshold choice before fitting. Firmness-based remaining life answers a specific texture-quality question; decay, colour and sensory rejection remain separately recorded outcomes.

Cao et al. provide a distinct sensory definition: individual weighted scores below 60 indicate rejection, and the reported endpoint is when 25–50% of fruit are unacceptable [4]. This supports explicit endpoint selection rather than treating

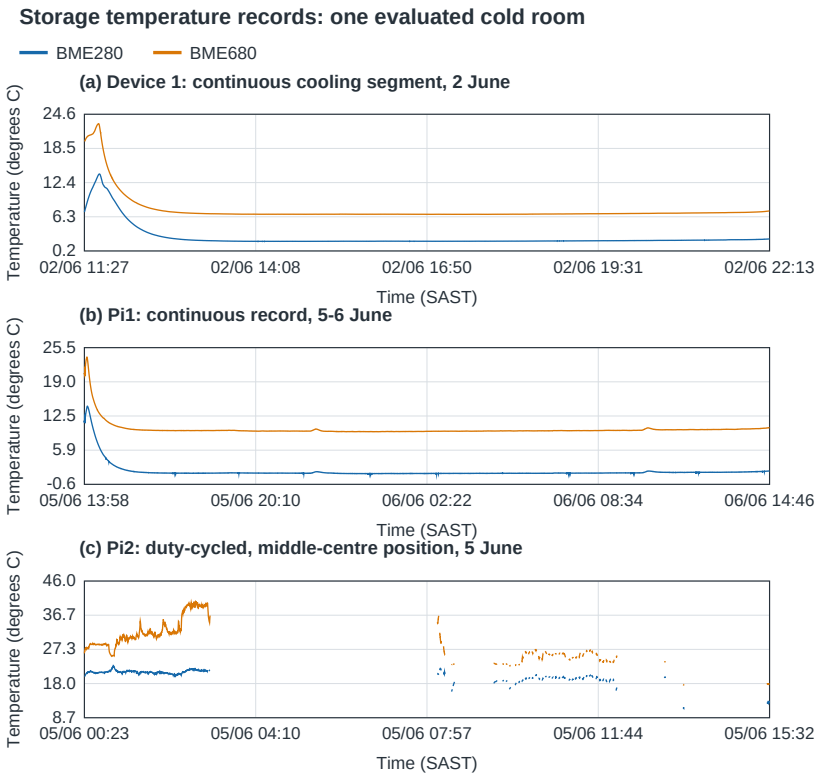


Fig. 4. Single-point-corrected temperature records from the evaluated cold room (SAST, UTC+2). Panels retain their individual dates and time axes. Lines break at gaps exceeding 60 s; Pi2 occupied a middle-centre position.

every measured quality attribute as the target. For a batch-level sensory model, one rejection fraction must be fixed in advance.

Periodic assessments bracket a threshold crossing between the last acceptable and first unacceptable visit. The target is interval-censored, or uses a predeclared interpolation rule with the visit interval retained as uncertainty. Fruit still acceptable at the last visit contribute right-censored follow-up. The supplied laboratory "Shelf life (d)" field denotes time under simulated retail conditions; the final visit day is not a failure label. Each observation must retain cultivar, batch, method, units, replicate and assessment time and be linked to the relevant device and exposure history. Destructive firmness tests use separate fruit at each visit. Features at prediction time t use only measurements recorded on or before t.

Table 3. Model-readiness gate for fruit-proximal shelf-life AI.

Criterion	Required evidence	Quality-assurance purpose
Defined outcome	Primary quality/marketability endpoint, method, threshold, and future storage condition	Gives shelf life a reproducible biological and operational definition.
Traceable linkage	Outcome linked to apple/sample group, batch, device, position, and exposure window	Aligns predictors and targets through a documented chain of custody.
Temporal integrity	Features use only measurements available at prediction time	Ensures that evaluation represents prospective prediction.
Independent units	Counts of apples, batches, and deployments distinguished from sensor rows	Provides an accurate effective sample size.
Technical quality	Calibration, clock, completeness, revision, and maintenance status documented	Separates environmental signal from instrument and operational effects.
Group-level validation	At least one usable independent group can be withheld	Measures transfer to a new batch or deployment.

7.2 Quality Assurance and Feature Construction

Processing begins with an immutable manifest of raw local and remote files, checksums, device identifiers, record counts, and configuration versions. Quality checks identify duplicates, non-monotonic clocks, range exceptions, long constant runs, acquisition intervals, and revision changes. Expected counts are calculated within verified operating windows; documented maintenance is reported as a separate deployment event. Any interpolation used for visualisation is declared and limited to scientifically justified short intervals.

Environmental streams are aligned no finer than the slowest relevant sampling rate. Candidate thermal features include elapsed time, mean and maximum temperature, variability, cooling or warming rate, cumulative degree-hours outside a justified band, excursion count, maximum excursion duration, and recovery time. Humidity features include mean, variability, and time outside a justified range. Resultant g-force is calculated from the three recorded acceleration axes in g:

$$a_m = \sqrt{a_x^2 + a_y^2 + a_z^2}. \tag{3}$$

Candidate mechanical features then include peak magnitude, event count above a calibrated threshold, jerk, duration, and stage-specific vibration energy. Gas resistance is used only when operating mode, revision, and calibration justify a biological interpretation.

Features are calculated by post-harvest stage before aggregation across the exposure history. Availability indicators retain information about power, connectivity, and maintenance events. Observed values remain unchanged, and any de-

clared preprocessing, scaling, imputation, or feature selection is estimated within the training folds only.

7.3 Leakage-Resistant Evaluation

The independent unit is the entity represented by the target: an apple, quality-assessment group, batch, or deployment. All rows belonging to the same unit remain in one partition, and grouped splitting accounts for the shared exposure history and outcome of temporally adjacent measurements. Where sufficient groups exist, one or more complete groups form a final test set and the remainder support grouped cross-validation. Where only leave-one-group-out validation is feasible, variability across held-out groups is reported and no single favourable fold is selected as the headline result.

Once assessment intervals supply the declared endpoint, machine-learning predictions will be compared with its laboratory-derived remaining-life target in days, retaining censoring where applicable. Two predictive baselines test incremental value: a static estimate based on elapsed time and declared nominal life, and a temperature-only estimate based on cumulative time–temperature exposure. Candidate algorithms will progress from linear and regularised regression to support-vector regression and tree ensembles as the number and range of independent groups support additional complexity.

For observed or predeclared point-valued remaining-life targets, mean absolute error (MAE) expresses prediction error in days:

$$\mathrm{MAE} = \frac{1}{n}\sum_{i=1}^{n} |\widehat{RSL}_i - RSL_i|. \tag{4}$$

Root mean square error indicates sensitivity to large errors, while mean signed error exposes systematic optimism or conservatism. For threshold classification, sensitivity, specificity, precision, and F1-score are reported with the class distribution. Overestimation of remaining life to the declared endpoint is analysed separately because its consequence differs from conservative underestimation.

A candidate is selected when it improves meaningfully on both baselines for unseen groups, uses information available at prediction time, and has uncertainty compatible with the intended decision. Model complexity is proportional to the incremental predictive value. Interpretation uses coefficients or training-safe feature importance to examine responses to temperature, humidity, elapsed time, and mechanical events; these are reported as predictive associations.

8 Discussion

8.1 Contribution and Evaluated Scope

The study establishes a serviceable apple-shaped instrument, recoverable field logging, and quantitative acquisition evidence from commercial operations. The

79-minute journey provides a complete short transport segment. Device 1 and Pi1 provide sustained cooling and storage records, while Pi2 quantifies the coverage consequence of intermittent acquisition. Their analytical value is in the documented duration, channel response and missingness structure.

The framework integrates these engineering findings with a supervised-learning data contract: defined endpoint, linked history, calibrated channels and independent evaluation groups. Sampling density determines temporal resolution; biological and deployment replication determine the scope of prediction. The contribution is the traceable connection between these requirements and observed field-data quality. Predictive performance belongs to the subsequent linked-outcome evaluation specified in Section 7.3.

8.2 Next Evaluation and Practical Use

Current modelling work includes implemented long short-term memory (LSTM) and gated recurrent unit (GRU) encoders for ordered temperature and relative-humidity histories. Observation masks and elapsed-time inputs retain missingness and irregular acquisition intervals. The prediction workflow combines the history representation with a requested horizon in days and declared future storage conditions to estimate acceptability probability. Study-specific fitting and evaluation are being prepared by linking identified sensor histories to compatible laboratory outcomes and follow-up intervals. The planned comparison will assess LSTM/GRU candidates and feature-based models against static and temperature-only baselines on held-out groups, following Section 7.3. Preprocessing is fitted on training groups, and validation groups guide model selection. Complementary firmness predictions will be assessed against laboratory observations. Calibration checks cover the storage temperature/RH domain, enclosure response and hardware revisions; censoring retains follow-up that ends before the declared endpoint.

After predictive evaluation, possible uses include prioritising inspection, reviewing cooling after an excursion and allocating stock by estimated remaining life. Each use requires an action threshold, responsible actor, response window and acceptable error rate. The locally serviceable, store-and-forward design addresses intermittent connectivity in Southern African supply chains. Sustainability assessment will relate resulting operational benefits to power, maintenance and instrument costs.

9 Conclusion

This paper presented a smart-apple platform and a framework for preparing fruit-proximal field histories for shelf-life AI. Reproducible analysis of 18 source files confirms continuous acquisition across a nominal 79-minute journey and substantial storage records, including 24.8 h of near-complete Pi1 acquisition and a quantified duty-cycled Pi2 record. Channel-specific measurements, gap-aware exposure accounting and source provenance make the engineering evidence

auditable. The defined firmness endpoint and six-point readiness gate specify the remaining linkage, calibration and group-level comparison required for biological prediction.

Acknowledgments. This work is based on research supported in part by the National Research Foundation of South Africa (Ref. No. SRUG22051310444). The opinions, findings, conclusions, or recommendations expressed are those of the author(s) alone, and the NRF accepts no liability whatsoever in this regard.

This study was funded within the framework of the UKUDLA project, which is funded by the Federal Foreign Office (AA), the Federal Ministry of Research, Technology and Space (BMFTR), the Federal Ministry of Agriculture, Food and Regional Identity (BMLEH/BLE), and the German Academic Exchange Service (DAAD).

Disclosure of Interests. The authors declare that they have no competing interests relevant to the content of this article.

References

1. Argenta, L.C., Anese, R.d.O., Thewes, F.R., Wood, R.M., Nesi, C.N., Neuwald, D.A.: Maintenance of 'luiza' apple fruit quality as affected by postharvest practices. Revista Brasileira de Fruticultura **44**(4), e905 (2022). https://doi.org/10.1590/0100-29452022905
2. Bosch Sensortec: BME680 Datasheet. Bosch Sensortec (2024), https://www.bosch-sensortec.com/media/boschsensortec/downloads/datasheets/bst-bme680-ds001.pdf, bST-BME680-DS001-09, revision 1.9, February 2024; accessed 8 October 2026
3. Broekman, A., Steyn, W.J., Steyn, J.L., Bill, M., Korsten, L.: smAvo and smaTo: A fruity odyssey of smart sensor platforms in southern africa. HardwareX **8**, e00156 (2020). https://doi.org/10.1016/j.ohx.2020.e00156
4. Cao, M., Wang, D., Qiu, L., Ren, X., Ma, H.: Shelf life prediction of 'royal gala' apples based on quality attributes and storage temperature. Horticultural Science and Technology **39**(3), 343–355 (2021). https://doi.org/10.7235/HORT.20210031
5. da Costa, T.P., Gillespie, J., Cama-Moncunill, X., Ward, S., Condell, J., Ramanathan, R., Murphy, F.: A systematic review of real-time monitoring technologies and its potential application to reduce food loss and waste. Sustainability **15**(1), 614 (2023). https://doi.org/10.3390/su15010614
6. Defraeye, T., et al.: Digital twins are coming: Will we need them in supply chains of fresh horticultural produce? Trends in Food Science & Technology **109**, 245–258 (2021). https://doi.org/10.1016/j.tifs.2021.01.025
7. DFRobot: Gravity BME688 Environmental Sensor (SEN0617). DFRobot (2026), https://wiki.dfrobot.com/sen0617/, accessed 10 August 2026
8. Elashmawy, R., Doron, M., Kanjilal, R., Brecht, J.K., Uysal, I.: The digital cold chain: Sensor-driven product quality with AI. Postharvest Biology and Technology **230**, 113714 (2025). https://doi.org/10.1016/j.postharvbio.2025.113714
9. Hevner, A.R., March, S.T., Park, J., Ram, S.: Design science in information systems research. MIS Quarterly **28**(1), 75–105 (2004)
10. Kapoor, S., Narayanan, A.: Leakage and the reproducibility crisis in machine-learning-based science. Patterns **4**(9), 100804 (2023). https://doi.org/10.1016/j.patter.2023.100804
11. Li, D., Bai, L., Wang, R., Sun, Y.: Research progress of machine learning in extending and regulating the shelf life of fruits and vegetables. Foods **13**(19), 3025 (2024). https://doi.org/10.3390/foods13193025
12. Małachowska, M., Tomala, K.: Apple quality during shelf-life after long-term storage and simulated transport. Agriculture **13**(11), 2045 (2023). https://doi.org/10.3390/agriculture13112045
13. Nsunba, P., Isingizwe Nturambirwe, J.F., Nyirenda, C.N.: Streamlining real-time sensor data collection and querying processes from environmental parameters in a food cold store. In: 2025 IST-Africa Conference (IST-Africa). pp. 1–11. IEEE (2025). https://doi.org/10.23919/IST-Africa67297.2025.11060485
14. Onwude, D.I., Chen, G., Eke-emezie, N., Kabutey, A., Khaled, A.Y., Sturm, B.: Recent advances in reducing food losses in the supply chain of fresh agricultural produce. Processes **8**(11), 1431 (2020). https://doi.org/10.3390/pr8111431
15. Pal, A., Kant, K.: Smart sensing, communication, and control in perishable food supply chain. ACM Transactions on Sensor Networks **16**(1), 1–41 (2020). https://doi.org/10.1145/3360726

16. Roberts, D.R., Bahn, V., Ciuti, S., Boyce, M.S., Elith, J., Guillera-Arroita, G., Hauenstein, S., Lahoz-Monfort, J.J., Schröder, B., Thuiller, W., Warton, D.I., Wintle, B.A., Hartig, F., Dormann, C.F.: Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. Ecography **40**(8), 913–929 (2017). https://doi.org/10.1111/ecog.02881
17. Shrivastava, C., Berry, T., Cronje, P., Schudel, S., Defraeye, T.: Digital twins enable the quantification of the trade-offs in maintaining citrus quality and marketability in the refrigerated supply chain. Nature Food **3**(6), 413–427 (2022). https://doi.org/10.1038/s43016-022-00497-9
18. Van Zeebroeck, M., Van Linden, V., Ramon, H., De Baerdemaeker, J., Nicolaï, B.M., Tijskens, E.: Impact damage of apples during transport and handling. Postharvest Biology and Technology **45**(2), 157–167 (2007). https://doi.org/10.1016/j.postharvbio.2007.01.015
19. Waveshare Electronics: Environment Sensor HAT: Technical Documentation. Waveshare Electronics (2026), https://www.waveshare.com/wiki/Environment_Sensor_HAT, accessed 10 August 2026
20. Zapp, H.R., Ehlert, S.H., Brown, G.K., Armstrong, P.R., Sober, S.S.: Advanced instrumented sphere (IS) for impact measurements. Transactions of the ASAE **33**(3), 955–960 (1990). https://doi.org/10.13031/2013.31423
21. Zhang, Y., Zhu, D., Ren, X., Shen, Y., Cao, X., Liu, H., Li, J.: Quality changes and shelf-life prediction model of postharvest apples using partial least squares and artificial neural network analysis. Food Chemistry **394**, 133526 (2022). https://doi.org/10.1016/j.foodchem.2022.133526
22. Zou, Y., Wu, J., Meng, X., Wang, X., Manzardo, A.: Digital twin integration for dynamic quality loss control in fruit supply chains. Journal of Food Engineering **397**, 112577 (2025). https://doi.org/10.1016/j.jfoodeng.2025.112577