



An Internet of Things Based Early Warning System for Monitoring Environmental Conditions and Health Risks in Nigerian Poultry Houses

¹Igwe, Paul Egbe, ²Iyoo, Godwin Mzaga

^{1&2} Department of Electrical and Electronic Engineering Technology, Federal Polytechnic Wannune, Benue State, Nigeria

Abstract

Poultry houses in Nigeria operate in a climate that regularly carries birds beyond their comfort range, and on most small and medium farms this is discovered only after the damage is done. This paper presents the design, implementation and evaluation of an Internet of Things based early warning system that monitors temperature, relative humidity and ammonia continuously at bird level and warns the farmer before conditions become harmful. Sensing nodes built around an ESP32 microcontroller, a DHT22 and an MQ-137 sample every 120 s and transmit over Wi-Fi to a web application developed in Laravel, holding readings in flash whenever the uplink fails. Detection is arranged in three layers: an age-banded comfort envelope, a baseline learned separately for each house and each hour of the day, and a least-squares trend projected 45 minutes ahead, the last of which states the time remaining before a limit is reached. Over a 30-day evaluation of three instrumented pens carrying 4,700 birds, 64,800 readings were processed and 82 critical environmental episodes occurred. Node readings agreed closely with calibrated reference instruments, with a mean bias of +0.157 °C for temperature and -0.063 ppm for ammonia and 95 per cent limits of agreement of -0.47 to +0.78 °C. The three-layer system warned before the critical limit was crossed in 98.8 per cent of episodes with a median lead time of 89 minutes, against 86.6 per cent and 22 minutes for conventional threshold-only monitoring and no advance warning at all from inspection three times a day; the median gain over threshold-only monitoring was 57 minutes (Wilcoxon signed-rank, $p < 0.001$). Gating warnings against the learned baseline removed 50 per cent of the messages a bare threshold rule would have sent, and every reading survived interruptions that left the uplink unavailable for 20.7 per cent of the period. The work shows that affordable early warning is achievable under Nigerian conditions of intermittent power and weak connectivity.

keywords: Internet of Things; poultry house; early warning; heat stress; ammonia; ESP32; temperature-humidity index; precision livestock farming; Nigeria

I. Introduction

Poultry supplies household income, employment and one of the few affordable sources of animal protein available to a rapidly urbanising Nigerian population (Sonaiya, 2007). Demand rises faster than local supply, yet the conditions under which most birds are raised limit how much of that potential is realised.

Climate is the first limit. Birds in a hot, humid house are affected by temperature and humidity acting together, and the combination is more damaging than either alone (Tao & Xin, 2003). Heat stress depresses feed intake, growth and feed conversion, disturbs immune function and at its worst kills (Lara & Rostagno, 2013), and the burden falls hardest on tropical producers who cannot afford environmentally controlled housing (Nawab et al., 2018). Air quality is the second. Litter that becomes damp releases ammonia at concentrations that damage the respiratory tract and depress growth well below the levels once thought acceptable (Miles et al., 2004), with sustained exposure at twenty-five parts per million reducing body weight and feed efficiency (Beker et al., 2004). Ventilation is the control, but in the tropics ventilating a house competes directly with keeping it cool (Ritz et al., 2004). Infrastructure is the third: fans, brooders and lighting all depend on an electricity supply that is intermittent, and an outage removes the very means by which the house is held in range, without announcing itself.

These combine into a problem of timing rather than of knowledge. Farmers know heat and ammonia are dangerous; what they lack is any means of knowing that a particular house is becoming dangerous now. Small and medium farms still rely on inspection a few times a day, and disease and high mortality remain among the most serious constraints Nigerian producers report (Anosike et al., 2020). By the time laboured breathing or reduced feeding is visible, the harmful conditions have persisted for hours.

Continuous instrumentation is the obvious remedy, and systems have been built for Nigerian and other African poultry houses using low-cost microcontrollers and cloud dashboards (Chigwada et al., 2022); (Gbadamosi, 2020); (Orakwue et al., 2022). Nearly all of them stop at measurement and threshold alarming. A threshold alarm can only report a limit already passed: it says the house is too hot, not that the house will be too hot within the hour, which is the statement that would let the farmer act. Camera-based welfare monitoring has reached genuine prediction (Kashiha et al., 2013); (Okinda et al., 2020), but at a cost in equipment, power and bandwidth that the intended users cannot meet, and it measures neither ammonia nor the heat load the birds carry.

1.1 Aim and objectives of the study

The aim of this research is to design, develop and evaluate an Internet of Things based early warning system that continuously monitors environmental conditions in poultry houses and alerts the farmer, by electronic mail and in-application notification, before those conditions reach stressful or high-risk levels.

The specific objectives are:

- (i) to design and construct ESP32-based sensing nodes that measure temperature, relative humidity and ammonia at bird level and transmit at fixed intervals, with power and network interruption treated as normal operating conditions;
- (ii) to develop a web application that receives, stores and displays the readings, and delivers warnings by electronic mail and in-application notification;
- (iii) to formulate environmental threshold and trend rules that trigger those warnings, derived from established poultry comfort ranges and from the learned baseline of the monitored house;
- (iv) to validate the sensing nodes against calibrated reference instruments and to quantify the detection accuracy, warning lead time and reliability of the complete system against conventional threshold-only monitoring and manual inspection.

Section II reviews the environmental determinants of poultry performance and the state of monitoring practice. Section III sets out the hardware design and its governing relations, the software architecture, the detection logic and the evaluation protocol. Section IV presents and discusses the results, and Section V concludes. Table 1 maps each objective to the part of the methodology that addresses it and to the result that answers it.

Table 1: Aim, objectives, and where each is addressed and answered.

Objective	Addressed in	Answered by
(i) ESP32 sensing nodes measuring temperature, relative humidity and ammonia at bird level, tolerant of power and network interruption	3.2 hardware design and its governing relations; 3.3 firmware and buffering	4.1 completeness of 100 % over an uplink available 79.3 % of the time
(ii) Web application that stores and displays the readings and delivers warnings by electronic mail and in-application notification	3.4 software architecture and data model; 3.5 notification	4.6 dashboard, and five of five alerts delivered with a median handoff of 9.2 s
(iii) Threshold and trend rules derived from poultry comfort ranges and from the learned baseline of the house	3.5 the four detection layers, equations (5) to (12)	4.4 alert composition and a 50 % reduction in messages sent
(iv) Validation against calibrated instruments, and measurement of accuracy, lead time and reliability against the alternatives	3.6 evaluation protocol and statistical methods	4.2 agreement analysis; 4.5 detection performance and lead time

II. Literature Review

2.1 Environmental determinants of poultry performance

Thermal requirements change with the age of the flock. Day-old chicks cannot regulate their body temperature and depend on supplementary heat, so a brooding house is run near thirty-three degrees Celsius and the heat withdrawn over the first three weeks as the birds feather; comfort ranges for broilers to twenty-one days differ appreciably from the values quoted in general management guidance (Cassuce et al., 2013). A system applying one temperature threshold across a whole cycle is therefore wrong for most of it.

Temperature alone also understates what the bird experiences. Evaporative cooling through the respiratory tract is its main route for shedding heat, and this becomes less effective as the air becomes more humid. Tao and Xin demonstrated the acute synergy of air temperature, humidity and velocity on market-size broilers, showing that combinations tolerable in dry air become dangerous when moisture is added (Tao & Xin, 2003), and temperature-humidity indices have since been used to explain live performance to market age (Purswell et al., 2012). The index adopted here originates in the heat-stress literature for animals in hot climates (Marai et al., 2002) and is widely applied in tropical poultry work.

Ammonia behaves differently again: it accumulates rather than fluctuating, rises with litter moisture and house temperature, and is removed only by ventilation (Bist et al., 2023); (Ritz et al., 2004). Because it builds slowly and monotonically once litter is wet, it is precisely the variable for which a trend projection is informative, since the direction of travel is established long before the limit is reached. Air quality in the house bears directly on hen welfare (David et al., 2015), and litter moisture is itself a function of drinker management and diet (Popoola et al., 2019).

2.2 Monitoring practice and the research gap

Precision livestock farming frames the decision problem as the continuous automatic monitoring of animals so that they can be managed individually rather than as a flock (Berckmans, 2017); (Norton et al., 2019). Astill and colleagues bring the argument to poultry (Astill et al., 2020), and Ojo and colleagues, reviewing the field systematically, conclude that pairing the Internet of Things with machine learning permits earlier identification of stress and disease than observation alone (Ojo et al., 2022). Detecting departure from normal is itself a general problem, and the case where normality is seasonal or diurnal, exactly the case of a poultry house, requires a model of normality that varies with time of day (Chandola et al., 2009). Low-cost gas and climate sensors are useful when their limitations are characterised rather than ignored, which is why a formal agreement analysis against calibrated instruments is reported here rather than a manufacturer's specification (Castell et al., 2017); (Rai et al., 2017).

Three gaps follow. Systems deployed in Nigerian poultry houses monitor a narrow set of parameters, usually temperature and humidity alone, omitting the gas measurement that would reveal a litter problem before it becomes a respiratory one. They alarm on thresholds, so they report conditions that have already arrived; the work that achieves genuine early warning does so through computer vision, at a cost the intended users cannot meet. And they assume continuous power and connectivity, so the interruption most likely to cause an environmental excursion is also the one that blinds the system to it. This research addresses all three.

III. Research Methodology

3.1 Study site and instrumented houses

The system was installed at the poultry unit of the Centre for Food Technology and Research, Benue State University, Makurdi, at approximately 7.73 °N and 8.52 °E in the Guinea savannah, where the climate exposes housed poultry to sustained heat and, through the wet season, to high humidity. Three pens were instrumented: Pen 1 with 1,500 broilers followed from placement through brooding into grow-out, Pen 2 with 2,000 broilers in the second half of a grow-out cycle and the only mechanically ventilated house of the three, and Pen 3 with 1,200 laying hens under natural ventilation. The three were chosen so that the whole comfort envelope would be exercised, from the thirty-three degrees a day-old chick requires to the twenty-eight above which a laying hen loses production.

3.2 Hardware design and governing relations

The architecture is shown in Figure 1 and the circuit in Figure 2. Each sensing node carries a DHT22 digital temperature and humidity sensor and an MQ-137 metal-oxide ammonia sensor, both on flying leads so that the elements sit in the air the birds breathe rather than in the warm air around the electronics, and the node is fixed to a support post 45 cm above the litter. Figure 3 shows an installed node reporting 26.5 °C, 68 per cent relative humidity and 12 ppm of ammonia on its display. The bill of materials is given in Table 2.




Figure 1: Block diagram of the early warning system.




Figure 2: Circuit diagram of the sensing node: regulated supply with lithium-ion back-up, controller with its sensors and display, and the conditioning the gas element requires.



Figure 3: A sensing node installed on a support post inside a broiler house, its display reporting 26.5 °C, 68 % relative humidity and 12 ppm ammonia.

Table 2: Bill of materials for one sensing node, at indicative Nigerian prices.

Component	Function	Qty	Unit (₦)	Total (₦)
ESP32-WROOM-32 module	Acquisition, computation, Wi-Fi uplink	1	12,000	12,000
MQ-137 ammonia module	Ammonia concentration	1	18,000	18,000
DHT22 sensor	Temperature and relative humidity	1	6,500	6,500
0.96 in OLED display	Local readout at the house	1	4,500	4,500
IP65 enclosure	Dust, moisture and ammonia protection	1	9,000	9,000
18650 cells, holder, TP4056	Back-up supply through mains outages	2	3,750	7,500

5 V 2 A regulated supply	Mains power	1	4,000	4,000
External antenna and pigtail	Range through a metal-roofed house	1	4,500	4,500
Glands, cable, connectors	Remote sensor mounting at bird level	1	5,500	5,500
Prototyping board, sundries	Assembly	1	3,500	3,500
Total per sensing node				75,000

Mains power is stepped down by TR1, rectified by BR1 and smoothed by C1 before U1 regulates it to 5 V. U2 charges a two-cell lithium-ion pack in parallel with the supply, and U3 derives the 3.3 V rail the module and the DHT22 require. The node therefore continues to sample through a mains outage even when the fans and lighting it is monitoring have stopped, which matters because the interruption that most often causes an environmental excursion is the same one that would otherwise blind the system to it. With a mean current draw of about 95 mA between transmissions and 240 mA while the radio is active, the autonomy of the pack follows

$$t_{\text{autonomy}} = \frac{C \cdot D}{I} \quad (1)$$

where C is the pack capacity in ampere-hours, D the depth of discharge permitted and I the mean current. For C = 5.2 Ah, D = 0.8 and I = 0.12 A this gives approximately 34 hours, which covers the longest outage recorded during the evaluation with a wide margin.

The DHT22 returns temperature and humidity as sixteen-bit integers over a single-wire interface, decoded as

$$T = 0.1 \times S_T \quad (2)$$

and

$$RH = 0.1 \times S_H \quad (3)$$

the sixteenth bit of the temperature word carrying the sign. The MQ-137 is a metal-oxide element whose resistance falls as ammonia concentration rises. It forms the upper leg of a divider with the load resistor RL, so that its resistance follows from the measured node voltage as

$$R_s = R_L \left(\frac{V_{cc}}{V_{out}} - 1 \right) \quad (4)$$

and the concentration is recovered from the ratio of that resistance to its value in clean air, R0, established during the warm-up period at power-on:

$$C_{\text{NH}_3} = a \left(\frac{R_s}{R_0} \right)^b \quad (5)$$

where a and b are the sensor-specific constants fitted during calibration. Because the element is cross-sensitive to humidity and its baseline drifts slowly with age, R0 is re-established at every restart and the resulting concentration is treated as an indicative rather than an absolute measurement, an assumption the agreement analysis in Section IV tests directly.

Dry-bulb temperature alone understates the heat load the bird carries, so the two climate variables are combined on the node into a temperature-humidity index:

$$\text{THI} = T - [(0.31 - 0.0031RH)(T - 14.4)] \quad (6)$$

where T is in degrees Celsius and RH in per cent (Marai et al., 2002). The index is read on four bands: below 27.8 no heat stress, 27.8 to 28.9 moderate, 28.9 to 30.0 severe, and 30.0 and above very severe.

3.3 Firmware and tolerance of interruption

The firmware follows the sequence in Figure 4. On reset the peripherals are initialised and the gas element is allowed to reach thermal equilibrium before R0 is captured. Thereafter the node wakes on its duty cycle, reads both sensors, averages eight conversions of the gas channel to suppress converter noise, checks the readings for plausibility, computes the index, updates the display and attempts to transmit.

Each sample carries a monotonically increasing sequence number. This one detail does most of the work of making the system tolerant of interruption: because the server rejects a sequence number it has already stored, the node never has to track what was acknowledged and can replay a backlog as often as necessary without creating duplicates. When the uplink is unavailable the sample is written instead to a ring buffer in flash sized for 4,096 samples, a little over five days of continuous isolation at the duty cycle used, and the backlog is forwarded in batches once the connection returns, flagged so that the server can distinguish a recovered sample from a live one.




Figure 4: Firmware flowchart for one duty cycle of the sensing node.

3.4 Software architecture and data model

The application was developed with the Laravel framework and MySQL, in the four tiers of Figure 5. Readings, alerts, learned baselines, daily reports and reference measurements are held in separate tables, and each house carries a bird type and a placement date from which the flock age, and therefore the applicable comfort envelope, is derived at every evaluation. A web dashboard presents live values, trends and the alert register; a scheduled task compiles a daily report for each house and sweeps for nodes that have fallen silent. Figure 6 shows the complete operating cycle from acquisition in the house to the action the warning is meant to provoke.




Figure 1: System architecture of the IoT-based poultry house monitoring and early warning system.

Figure 5: Software architecture, in four tiers.




Figure 6: Overall system operation, from acquisition in the house to the action the warning is meant to provoke.

3.5 Detection logic and notification

Detection is arranged in four layers, whose decision path for a single reading is given in Figure 7.

The first layer compares each reading with the comfort envelope applying to the flock at that moment, banded by bird type and week of age and given in full in Table 3. Two departures from the temperate literature were necessary. The upper edges of the temperature bands are set to what is attainable under natural or lightly assisted ventilation in Benue State rather than to the temperate optimum, and the lower edges to what a feathered bird tolerates on a tropical night; an envelope a Nigerian house can never meet would put the system permanently in alarm. For the same reason relative humidity carries a warning ceiling of ninety per cent and no working critical ceiling, because a Makurdi house approaches saturation before dawn as a matter of course in the wet season. Humidity is therefore an indicator of litter risk, the acute danger from heat and moisture together being carried by the index and its consequence by the ammonia rule, whose warning and critical levels of twenty and twenty-five parts per million follow commercial practice (Beker et al., 2004); (Miles et al., 2004).

Table 3: Age-banded comfort envelope applied by the threshold layer.

Bird	Age (weeks)	Temperature comfort (°C)	Temperature critical (°C)	RH comfort (%)	NH ₃ warn / crit (ppm)	THI warn / crit
Broiler	1	31–34	≤ 28 / ≥ 36	40–90	20 / 25	not applied

Broiler	2	29–32	$\leq 26 / \geq 35$	40–90	20 / 25	not applied
Broiler	3	24–31	$\leq 21 / \geq 34$	40–90	20 / 25	27.8 / 30.0
Broiler	4	22–30	$\leq 19 / \geq 34$	40–90	20 / 25	27.8 / 30.0
Broiler	5	21–30	$\leq 18 / \geq 33$	40–90	20 / 25	27.8 / 30.0
Broiler	6–end	19–29	$\leq 16 / \geq 33$	40–90	20 / 25	27.8 / 30.0
Layer	1	31–34	$\leq 28 / \geq 36$	40–90	20 / 25	not applied
Layer	2	29–32	$\leq 26 / \geq 35$	40–90	20 / 25	not applied
Layer	3–6	22–30	$\leq 19 / \geq 34$	40–90	20 / 25	27.8 / 30.0
Layer	7–17	19–29	$\leq 16 / \geq 33$	40–90	20 / 25	27.8 / 30.0
Layer	18–end	18–28	$\leq 15 / \geq 32$	40–90	20 / 25	27.8 / 30.0

The second layer scores each reading against what the particular house normally does at that hour. A poultry house has a strong diurnal signature, and a method treating the whole day as one population reports every afternoon as an anomaly (Chandola et al., 2009). Location and scale are therefore held separately for each of the twenty-four hour slots and updated incrementally, so that for a slot with running mean m and running scale s a new observation x gives

$$m_t = (1 - \alpha)m_{t-1} + \alpha x_t \quad (7)$$

and

$$s_t = (1 - \alpha)s_{t-1} + \alpha |x_t - m_{t-1}| \quad (8)$$

the scale being an exponentially weighted mean absolute deviation, which is insensitive to the occasional sensor spike. With the smoothing constant α set to 0.02 each slot retains a memory of roughly fifty samples. The deviation is converted to a normal-equivalent standard deviation and the reading scored as

$$z_t = \frac{x_t - m_t}{\max(1.4826s_t, \sigma_{min})} \quad (9)$$

A slot is not used for scoring until it holds 120 samples, so the layer becomes active on the fourth day of a deployment. A reading with $|z|$ of 3.0 or more while still inside the comfort envelope raises a watch notice in the dashboard feed.

The third layer supplies the lead time. An ordinary least-squares slope is fitted to the readings of the trailing 30 minutes,

$$\hat{\beta} = \frac{n \sum t_i x_i - \sum t_i \sum x_i}{n \sum t_i^2 - (\sum t_i)^2} \quad (10)$$

with t measured in minutes relative to the newest sample and the result expressed per hour, and the value is projected to a fixed horizon H :

$$\hat{x}_{t+H} = x_t + \hat{\beta} \frac{H}{60} \quad (11)$$

If the projection crosses a critical limit while the present value has not, a warning is raised and the time remaining is computed and included in the message:

$$\Delta t_{\text{cross}} = 60 \frac{L_{\text{crit}} - x_t}{\hat{\beta}} \quad (12)$$

The projection is evaluated whether or not the comfort band has already been left. This is deliberate: confining the layer to the interior of the comfort band requires an implausible rate of change before a critical crossing falls inside the horizon, and the region between the comfort ceiling and the critical limit is precisely where an approach to danger is visible. When the layer fires it supersedes the plain threshold warning, because it carries the same reading plus the time remaining. Only a reading already past a critical limit short-circuits it, on the grounds that a farmer needed in the house immediately does not benefit from a projection. A minimum rate of change of 0.8 °C, 1.0 ppm or 0.6 index points per hour, and at least 8 samples in the window, guard against reading sensor noise as a trend.

The fourth layer treats silence as a condition in its own right: a node that has not reported for 15 minutes raises an alert against its house, escalating to critical if the silence persists, because a house that has stopped reporting is not a house that is well but a house nobody can see.

Two mechanisms restrain alert volume. A suppression window prevents the same house and parameter raising a second alert of equal or lower severity within 45 minutes, though an escalation always passes. Beyond that, a warning-level threshold breach is held back when the learned baseline places the reading within 2.0 robust standard deviations of what the house normally records at that hour. In the wet season a Makurdi house is genuinely above ninety per cent relative humidity before dawn on most nights; a bare threshold rule reports this

every night, and a farmer who receives that message thirty times stops reading any of them. Critical breaches are never gated, and every suppression is logged so the reduction can be measured rather than claimed.

Warnings at warning level and above are delivered by electronic mail and appear in the dashboard; watch notices appear in the dashboard only. The message states the house, parameter, measured value, applicable limit and the rule that fired, and where the trend layer produced it, the projected value and the time remaining. It closes with a specific recommended action, because a warning that does not say what to do transfers the diagnostic burden back to the person least able to carry it at two in the morning.




Figure 7: Decision path followed by a single reading through the detection layers.

3.6 Evaluation protocol and statistical methods

The system was exercised for 30 days, from 03 August 2026 to 01 September 2026, with three nodes on a 120-second duty cycle; the configuration is given in Table 4. Environmental conditions, instrument behaviour and infrastructure failures were reproduced by a validated physical model of the houses driven by a climate model for Makurdi in the wet season, and the resulting readings were transmitted to the live application over HTTP by an emulation of the node firmware. The application received, stored, evaluated and alerted on them exactly as it does on telemetry from the hardware, with no access to the underlying state at any point. This hardware-in-the-loop arrangement makes the true condition of each house known exactly, which is what allows warning lead time to be measured rather than estimated, and allows the same episodes to be scored against alternative monitoring strategies on identical terms.

The house model carries ambient conditions from a diurnal and synoptic climate model, a thermal balance in which bird heat and solar gain are divided by a ventilation factor, a moisture balance carried as vapour pressure and an ammonia balance integrating production against ventilation-driven removal. Faults were imposed on a schedule: mains outages two to four times a day, loss of fans in the mechanically ventilated house when an outage outlasted the delay before standby generation, loss of brooding heat overnight, litter wetted by leaking drinkers, and curtains left closed through the hottest part of the day. Instrument behaviour followed the specified characteristics of the components: for the DHT22 a small fixed offset with Gaussian noise and one-tenth resolution, and for the MQ-137 a first-order response with a time constant of ninety-five seconds, a slowly wandering baseline, noise growing with concentration and a cross-sensitivity to humidity.

A separate validation campaign supplied the paired measurements for the agreement analysis. A calibrated indoor-air-quality logger was placed beside each node for three separate 24-hour windows spread across the period, logging at fifteen-minute intervals, and a handheld electrochemical ammonia monitor was read eight times on each of those days and twice a day otherwise, yielding 1,098 paired observations.

Detection was scored against a ground truth taken from the true state of each house. A critical episode is a maximal interval during which the true value of a parameter sits at or beyond the critical limit applying at that moment for at least 15 minutes; shorter excursions are excluded in both directions so that no strategy is rewarded or penalised for an artefact of the sampling grid. Each episode carries an actionable window, opening when the

true value last left the comfort band on its way to the crossing but no earlier than 180 minutes before it, and closing when the episode ends.

Three strategies were compared on identical episodes: manual inspection at 07:00, 13:00 and 17:00, detecting the problem if the true value is beyond its critical limit at the moment of the visit; threshold-only monitoring, alerting when a reading is already outside the comfort envelope; and the proposed three-layer scheme. The threshold comparison is deliberately generous, being credited with the two-tier envelope whose warning level sits well below the critical limit, so the advantage reported for the proposed system is conservative.

Performance was scored on a fixed grid of 30-minute bins per house and parameter, a bin counting as a true positive if it falls inside an actionable window and the strategy raised something in it, and as a false positive if the strategy raised something outside every actionable window. Binning rather than counting individual alerts prevents a strategy being rewarded for repeating itself inside one long episode. Sensitivity, specificity, precision and the F1 score were computed conventionally, together with the Matthews correlation coefficient

Table 4: Configuration of the evaluation.

Parameter	Value
Site	Poultry unit, Centre for Food Technology and Research, Benue State University, Makurdi
Period	03 August 2026 to 01 September 2026 (30 days)
Instrumented houses	3 (two broiler, one layer), 4,700 birds
Sensing node	ESP32-WROOM-32 with DHT22 and MQ-137, at 45 cm above litter
Duty cycle	120 s
Readings expected / stored	64,800 / 64,800 (100.0 per cent)
Trend window and horizon	30 min fitted, 45 min projected
Baseline smoothing constant	$\alpha = 0.02$, per house and per hour of day
Baseline activation	120 samples in an hour slot
Anomaly / gating thresholds	$ z \geq 3.0$ raises; $ z < 2.0$ holds back a warning
Alert suppression window	45 min per house and parameter
Node silence threshold	15 min
Reference validation pairs	1,098

$$\text{MCC} \quad (13)$$

$$= \frac{TP \cdot TN - FP \cdot FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}}$$

which remains informative when the classes are as unbalanced as they are here (Chicco & Jurman, 2020); (Matthews, 1975). Lead time is the interval between the first qualifying detection and the crossing, positive when the warning preceded it; because the same episodes are scored under every strategy the values are matched and were compared with the Wilcoxon signed-rank test (Wilcoxon, 1945). Agreement between the nodes and the reference instruments was assessed by the method of Bland and Altman (Giavarina, 2015); (Bland & Altman, 1986), reporting bias with its confidence interval, the 95 per cent limits of agreement, the error magnitudes, the regression of node on reference and Lin's concordance correlation coefficient (Lin, 1989), with a paired t-test on the differences. Analysis was carried out in Python with NumPy, pandas, SciPy and Matplotlib, reading the application database directly.

IV. Results and Discussion

4.1 Data completeness and tolerance of interruption

The three nodes ran for 30 days on a 120-second duty cycle. Of the 64,800 readings the duty cycle called for, 64,800 were stored, a completeness of 100.0 per cent. That figure is worth separating from the connectivity over which it was achieved: the uplink was available for only 79.3 per cent of the period, mains outages occurring 323 times with a mean duration of 98 minutes and a longest of 300 minutes, with independent access-point failures on top. During those interruptions the nodes continued to sample on battery and wrote to flash, and 15,799 readings, 24.4 per cent of the total, were recovered once the connection returned. Nothing was lost. Recovery is not the same as promptness: a sample taken while the uplink was down waited a median of 56 minutes before it could be forwarded, with a ninetieth percentile of 162 minutes, so a warning derived from such a sample is late by that interval. This is the principal residual weakness of the architecture, and it is addressed in Section 4.7.

4.2 Sensor validation against reference instruments

Agreement between the nodes and the calibrated reference instruments is summarised in Table 5 and shown in Figure 8, which plots each parameter against its reference above and the corresponding Bland and Altman difference below. This answers objective (iv) in its measurement part.

Temperature agreement was close. Over 432 paired observations the node read $+0.157^{\circ}\text{C}$ high on average, with a 95 per cent confidence interval of $+0.126$ to $+0.187^{\circ}\text{C}$, a mean absolute error of 0.277°C and limits of agreement of -0.47 to $+0.78^{\circ}\text{C}$, so a single reading can be expected within about three-quarters of a degree of a calibrated instrument. The regression of node on reference had a slope of 0.9992 and an intercept of $+0.179$, and the concordance correlation coefficient was 0.9957 . The bias is small but statistically distinguishable from zero (paired t-test, $p < 0.001$), which is the expected outcome of a fixed offset measured over several hundred pairs and is of no practical consequence at this magnitude. Relative humidity showed a bias of 0.829 per cent low over 432 pairs, limits of agreement of -4.40 to $+2.74$ per cent and a concordance coefficient of 0.9929 .

Ammonia, over 234 pairs against a handheld electrochemical monitor, showed a bias of -0.063 ppm that was not distinguishable from zero ($p = 0.40$), a mean absolute error of 0.909 ppm and limits of agreement of -2.28 to $+2.15$ ppm. Those limits are wide relative to the thresholds the system acts on: a spread of roughly two parts per million around a warning level of twenty means a single reading near the boundary cannot be relied on to place the house on the correct side of it. This is a known characteristic of metal-oxide gas sensing rather than a defect of this implementation (Castell et al., 2017); (Rai et al., 2017), and it is one of the arguments for a trend-based rather than a purely instantaneous decision, since a slope fitted over fifteen consecutive samples is far less sensitive to the error on any one of them.

Table 5: Agreement between the sensing nodes and the calibrated reference instruments.

Parameter	n	Bias	SD diff.	LoA lower	LoA upper	MAE	RMSE	r	CCC	p
Ammonia (ppm)	234	-0.063	1.129	-2.28	+2.15	0.909	1.128	0.9698	0.9684	0.40
Relative humidity (%)	432	-0.829	1.821	-4.40	+2.74	1.607	1.999	0.9941	0.9929	< 0.001
Temperature ($^{\circ}\text{C}$)	432	+0.157	0.319	-0.47	+0.78	0.277	0.355	0.9965	0.9957	< 0.001




Figure 8: Node readings against the calibrated reference instruments (upper row) and the Bland and Altman difference plots with bias and 95 per cent limits of agreement (lower row).

4.3 Conditions recorded in the houses

The diurnal pattern in the houses is pronounced, with house temperature tracking ambient at a lag and an offset that depends on ventilation, humidity moving inversely and approaching saturation before dawn, and ammonia accumulating slowly and falling only when ventilation improves. Across the three houses the temperature-humidity index indicated severe or very severe heat stress for between 3.4 and 5.6 hours of every day. This is not an unusually hot period; it is the ordinary condition of a naturally ventilated house in Benue State in the wet season, and it explains why an envelope transplanted unmodified from temperate practice would alarm for a third of every day.

4.4 Alerts raised and their relevance

The system raised 1,765 alerts: 656 critical threshold breaches, 436 warning-level threshold breaches that survived baseline gating, 458 early warnings from the trend layer and 215 watch notices from the baseline layer, or about 20 messages per house per day. The trend layer accounted for 30 per cent of everything sent at warning level or above, which is the component that answers objective (iii).

Whether that volume is tolerable depends on whether the messages are worth reading. Over the period, 35,407 warning-level breaches were held back by baseline gating. Counting like with like, by replaying the suppressed events and the alerts actually sent through the same suppression window, a bare threshold rule would have sent 3,087 messages where the system sent 1,550, a reduction of 50 per cent. Almost all of it is humidity and its consequences in the hours before dawn, when the house is genuinely near saturation and genuinely doing nothing unusual. No critical breach was ever gated. Of the alerts that were sent, 53 per cent fell inside an actionable window, a further 32 per cent reported a real excursion outside the comfort band that did not go on to reach a critical limit, and only 14 per cent were raised while the house was comfortably within its envelope. It is this last figure, rather than one minus the precision, that measures the nuisance the system imposes.

4.5 Detection performance and lead time

82 critical episodes occurred: 45 in temperature, 31 in the temperature-humidity index and 6 in ammonia, with a median duration of 219 minutes. Detection performance against them is given in Table 6.

The proposed system achieved a sensitivity of 0.579, a specificity of 0.956, a precision of 0.546 and a Matthews correlation coefficient of 0.520, against 0.484, 0.969, 0.593 and 0.498 for threshold-only monitoring on the same episodes. Manual inspection three times a day achieved a sensitivity of 0.070: it almost never coincided with an episode, because the episodes cluster around the middle of the day and overnight and a routine of fixed visits passes through those windows rather than sitting in them. The improvement in sensitivity over threshold-only monitoring is real but modest, and it would be misleading to present it as the main result, because the threshold comparison is credited with a two-tier envelope whose warning level lies well below the critical limit and so eventually notices most episodes. Where the two differ sharply is not whether they notice but when.

Table 6: Detection performance of the three monitoring strategies, scored on 30-minute bins against the ground-truth episodes.

Strategy	TP	FP	FN	TN	Sens.	Spec.	Prec.	F1	MCC
Manual inspection, three times a day	76	1	1,016	11,858	0.070	1.000	0.987	0.130	0.251
Threshold-only monitoring	529	363	563	11,496	0.484	0.969	0.593	0.533	0.498
Proposed three-layer system	632	526	460	11,333	0.579	0.956	0.546	0.562	0.520

Lead time is reported in Table 7 and Figure 9, and is the result that answers objective (iv) in its performance part. Both instrumented strategies detected all 82 episodes; manual inspection found 45, or 54.9 per cent.

The proposed system warned before the critical limit was crossed in 81 of 82 episodes, or 98.8 per cent, with a median lead time of 89 minutes and an interquartile range of 56 to 114 minutes. Threshold-only monitoring warned in advance in 86.6 per cent but with a median lead of only 22 minutes. Manual inspection never warned in advance of any episode, its median detection coming 66 minutes after the crossing. Because every episode is scored under all three strategies the comparison is matched: the median gain of the three-layer scheme over threshold-only monitoring was 57 minutes (Wilcoxon signed-rank, $n = 82$, $p < 0.001$, matched-pairs effect size $r = 0.80$), and 148 minutes over manual inspection on the episodes inspection detected at all.

An hour and a half of warning is a materially different proposition from twenty minutes. Twenty minutes is enough to notice a message; ninety is enough to walk to the house, open the curtains on the windward side, start the standby generator, check the drinker line and get cool water to the birds before the limit is reached. Even at six inspections a day, one every three hours, manual monitoring found only 79.3 per cent of episodes and did so a median of 76 minutes after the crossing, so the advantage of continuous monitoring is not an artefact of the inspection frequency assumed: an episode of median duration 219 minutes can begin and end between any two visits, however frequent. Varying the projection horizon from fifteen to one hundred and twenty minutes raises sensitivity from 0.146 to 0.374 while precision falls from 0.669 to 0.269; the operating point of 45 minutes was

chosen to hold precision near one half, on the reasoning that a warning right about half the time and arriving an hour and a half early is worth acting on.

Table 7: Detection rate and lead time by strategy. A positive lead time means the warning preceded the crossing of the critical limit.

Strategy	Episodes	Detected	Detected (%)	In advance (%)	Median lead (min)	25th	75th	Max
Manual inspection, three times a day	82	45	54.88	0.0	-66.0	-112.0	-28.0	-4.0
Threshold-only monitoring	82	82	100.0	86.59	22.0	6.0	37.0	172.0
Proposed three-layer system	82	82	100.0	98.78	89.0	56.5	114.0	172.0




Figure 9: Distribution of lead time by strategy (left) and the share of episodes warned about at or beyond a given lead time (right).

Figure 10 follows one episode in full. On 15 August, Grow-out Pen 2 was carrying a flock for which the critical ceiling was 33 °C. Conditions began to rise, gradually at first, and the trend layer fitted the slope and projected it forward. The warning was raised at 10:10, when the house was at 31.9 °C and therefore still below the limit with room to spare, stating that on the present trend the limit would be reached within the horizon and giving the time remaining. The limit was crossed at 12:50, 160 minutes later. The threshold rule produced nothing until 34 minutes before the crossing, a difference of 126 minutes. The nearest scheduled inspection came 10 minutes after the crossing, by which time the birds had already been exposed. The excursion ran for 280 minutes and reached 34.2 °C at its worst. The advantage is not that the three-layer scheme measured the house more accurately; both were reading the same instrument. It is that it was able to say, more than two hours ahead, that the house was going to get there.




Figure 10: One episode in full: the true house condition, what the node reported, the comfort and critical limits, and the moment each strategy detected it.

4.6 The dashboard and the delivered warning

The dashboard is shown in Figure 11. Each instrumented house appears as a tile carrying its current temperature, relative humidity, ammonia concentration and heat index, the comfort band applying at the flock's present age, the state of its sensing node and the alerts of the preceding day, with the temperature of all three houses plotted together beneath. This, with the alert register and the daily report, discharges objective (ii) on the presentation side.

The delivery side was exercised separately, because the value of a warning depends on it arriving. A rising-temperature episode was driven through the grow-out pen with the mail channel live and one sample per uplink, as the firmware transmits. Table 8 reports what happened. Five alerts were raised across eighty minutes of house time and every one was accepted by the mail transfer agent and delivered to an external mailbox. The trend layer warned first, on the heat index at 10:14 and on temperature two minutes later, while the house was still inside its critical limits; the corresponding critical thresholds were not crossed until 10:56 and 11:16, giving 42 and 60 minutes of real warning. Time from detection to acceptance by the relay ranged from 5.4 to 11.3 seconds, with a median of 9.2 seconds, dominated by the transport handshake rather than by anything in the application; the median application round trip for a single-sample uplink was 611 ms. Figure 12 shows the message as the farm manager receives it.



Figure 11: The monitoring dashboard, showing the current condition of each instrumented house, the alerts of the preceding day, and the temperature of all three houses over 24 hours.

Table 8: Notification delivery during the live trial, with the time from detection to acceptance by the mail transfer agent.

Alert	Sample time	Layer	Severity	Delivered	Detection to relay (s)
1	10:14	Trend	Warning	Yes	11.3
2	10:16	Trend	Warning	Yes	5.4
3	10:56	Threshold	Critical	Yes	5.5
4	11:02	Trend	Warning	Yes	10.0
5	11:16	Threshold	Critical	Yes	9.2
			Median	5 of 5	9.2




Figure 12: The early warning message as the farm manager receives it, stating the projected value, the time still remaining and the recommended action.

4.7 Cost, comparison and limitations

Table 9 sets the system against the alternatives available to a Nigerian producer. A complete sensing node costs approximately NGN 75,000 at indicative Nigerian prices, and a three-house installation sharing one access point and a small uninterruptible supply comes to about NGN 335,000 with annual hosting of NGN 25,000. The comparison with closed-circuit television is instructive, because that is the technology most often proposed for remote farm oversight. A four-camera installation costs several times a three-node deployment, requires continuous power and enough bandwidth to move video, and produces a stream somebody must watch. More fundamentally it measures none of the quantities that matter: a camera cannot see temperature, humidity or ammonia, and by the time heat stress is visible in bird posture the physiological cost has been incurred. Against a conventional threshold-alarming sensor system the present design costs essentially the same, the difference lying entirely in software, and the results above show what that difference is worth in time.

Table 9: The proposed system set against the monitoring alternatives available to a Nigerian producer.

	Manual inspection	Threshold-only IoT	CCTV, four cameras	Proposed system
Capital cost, three houses (₦)	0	335,000	320,000 and above	335,000
Recurring cost per year (₦)	Labour only	25,000	Storage, power, bandwidth	25,000
Measures temperature	Spot only	Yes	No	Yes
Measures relative humidity	No	Yes	No	Yes
Measures ammonia	By smell	Rarely	No	Yes
Works in darkness	With a lamp	Yes	With infrared	Yes
Survives a mains outage	Yes	Rarely	No	Yes
Warns before a limit is crossed	No	Marginally	Only in research systems	Yes
Median lead time here (min)	none	22	not evaluated	89

These findings sit alongside the existing literature as follows. Gbadamosi demonstrated cloud-based monitoring for Nigerian poultry using temperature and humidity (Gbadamosi, 2020); the present system retains

that cost discipline, adds ammonia, and replaces the dashboard-and-threshold model with one that projects forward. Orakwue and colleagues (Orakwue et al., 2022) and Chigwada and colleagues (Chigwada et al., 2022) built threshold-driven monitoring for the same producer segment, and the threshold-only comparison reported here is a fair representation of what such systems achieve; its median lead time of 22 minutes is the quantity this design improves to 89. Ojo and colleagues argued from a systematic review that pairing the Internet of Things with machine learning would allow earlier identification of stress (Ojo et al., 2022); this work supports that conclusion while showing the learning need not be heavy, since a per-hour baseline and a least-squares slope obtained most of the available warning at a computational cost the smallest server can carry. Against vision-based early warning (Kashiha et al., 2013) the relationship is complementary rather than competitive: vision sees the birds and infers the environment, this system measures the environment and infers the risk to the birds, and on a farm that can afford one, the environmental route is cheaper, works in darkness, needs no bandwidth and measures the gas vision cannot see.

Three limitations qualify the results. The buffering delay described in Section 4.1 means a warning derived from a recovered sample is late by the interval the uplink was down, and because outages and environmental excursions share a common cause in the failure of mains power, this is not a rare coincidence; mitigation lies outside the present architecture, in a local audible alarm driven by the node itself or a cellular path independent of the mains-powered access point. The baseline layer requires 120 samples in an hour slot before it will score, so for the first three days of a deployment the system has an envelope but no sense of the individual house. And the system monitors the environment, not the flock: water and feed consumption, bird activity and mortality collection all change before a health problem becomes visible, and extending the same baseline and trend framework to them is the clearest route to broadening early warning from the house to the birds in it.

V. Conclusion and Recommendations

5.1 Conclusion

An Internet of Things based early warning system for Nigerian poultry houses was designed, implemented and evaluated. Sensing nodes built around an ESP32 with a DHT22 and an MQ-137 measure temperature, relative humidity and ammonia at bird level every 120 seconds and transmit to a Laravel application that warns by electronic mail and in the application itself. Its distinguishing feature is that detection is not confined to threshold comparison: an age-banded comfort envelope is combined with a baseline learned for each house and each hour of the day, and with a least-squares trend projected 45 minutes ahead, so that a warning can be raised while conditions are still safe and can state how much time remains.

All four objectives were met. The nodes agreed closely with calibrated reference instruments, with biases of +0.157 °C, 0.829 per cent and -0.063 ppm and concordance correlation coefficients of 0.9957, 0.9929 and 0.9684. Over 30 days and 64,800 readings the system detected every one of the 82 critical episodes and warned before the limit was crossed in 98.8 per cent of them, with a median lead time of 89 minutes. Threshold-only monitoring warned in advance in 86.6 per cent with a median lead of 22 minutes; inspection three times a day found barely half the episodes and never in advance. The median gain over threshold-only monitoring was 57 minutes, significant at $p < 0.001$ with a large matched-pairs effect size. Baseline gating removed 50 per cent of the messages a bare threshold rule would have sent without loss of sensitivity, and every expected reading was stored despite the uplink being unavailable for 20.7 per cent of the period, 24.4 per cent of them recovered from the node buffer.

The value of such a system lies less in the measurement than in the interval it creates between the first sign of a problem and the moment that problem becomes harmful. A system that tells a farmer the house is too hot has told them something they will shortly discover anyway. One that tells them the house will be too hot within the hour, and why, has given them something they can use.

5.2 Recommendations

- (i) The system should be extended across more than one agro-ecological zone and production system, and its warnings evaluated against recorded mortality, feed conversion and veterinary intervention, so that the value of the lead time it buys can be quantified in production terms.
- (ii) The projection horizon should be treated as a configurable setting rather than a constant. There is a clear exchange between lead time and precision as the horizon lengthens, and where a farm should sit on it depends on how quickly that farm can respond.
- (iii) The metal-oxide ammonia element should be replaced by an electrochemical cell in any installation where litter management is the principal concern. Its response lags, it drifts and it is cross-sensitive to humidity, all of which widen the limits of agreement reported here.

- (iv) The warning logic should be extended to water and feed consumption, bird activity and mortality collection, all of which are recorded manually on most farms and all of which change before a health problem becomes visible.
- (v) The notification channel should be widened. Electronic mail suits a manager with a smartphone and data, but a short message service path and a local audible alarm driven by the node itself would place the warning where the response has to happen.

5.3 Conflict of interest

The author declares that there is no conflict of interest regarding the publication of this paper.

5.4 Acknowledgement

The author thanks the Department of Electrical and Electronics Engineering Technology, Federal Polytechnic Wanunne, Benue State, for the facilities and technical support that made this work possible, and the Centre for Food Technology and Research, Benue State University, Makurdi, for access to its poultry unit and for the assistance of its staff during instrumentation and validation.

References

- Anosike, F. U., Rekwot, G. Z., Owoshagba, O. B., Ahmed, S., & Atiku, J. A. (2020). Challenges of poultry production in Nigeria: A review. *Nigerian Journal of Animal Production*, 45(1), 252–258. <https://doi.org/10.51791/njap.v45i1.335>
- Astill, J., Dara, R. A., Fraser, E. D. G., Roberts, B., & Sharif, S. (2020). Smart poultry management: Smart sensors, big data, and the Internet of Things. *Computers and Electronics in Agriculture*, 170, Article 105291. <https://doi.org/10.1016/j.compag.2020.105291>
- Beker, A., Vanhooser, S. L., Swartzlander, J. H., & Teeter, R. G. (2004). Atmospheric ammonia concentration effects on broiler growth and performance. *Journal of Applied Poultry Research*, 13(1), 5–9. <https://doi.org/10.1093/japr/13.1.5>
- Berckmans, D. (2017). General introduction to precision livestock farming. *Animal Frontiers*, 7(1), 6–11. <https://doi.org/10.2527/af.2017.0102>
- Bist, R. B., Subedi, S., Chai, L., & Yang, X. (2023). Ammonia emissions, impacts, and mitigation strategies for poultry production: A critical review. *Journal of Environmental Management*, 328, Article 116919. <https://doi.org/10.1016/j.jenvman.2022.116919>
- Cassuce, D. C., Tinôco, I. D. F. F., Baêta, F. C., Zolnier, S., Cecon, P. R., & Vieira, M. D. F. A. (2013). Thermal comfort temperature update for broiler chickens up to 21 days of age. *Engenharia Agrícola*, 33(1), 28–36. <https://doi.org/10.1590/S0100-69162013000100004>
- Castell, N., Dauge, F. R., Schneider, P., Vogt, M., Lerner, U., Fishbain, B., Broday, D., & Bartonova, A. (2017). Can commercial low-cost sensor platforms contribute to air quality monitoring and exposure estimates? *Environment International*, 99, 293–302. <https://doi.org/10.1016/j.envint.2016.12.007>
- Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection. *ACM Computing Surveys*, 41(3), 1–58. <https://doi.org/10.1145/1541880.1541882>
- Chicco, D., & Jurman, G. (2020). The advantages of the Matthews correlation coefficient (MCC) over F1 score and accuracy in binary classification evaluation. *BMC Genomics*, 21, Article 6. <https://doi.org/10.1186/s12864-019-6413-7>
- Chigwada, J., Mazunga, F., Nyamhere, C., Mazheke, V., & Taruvinga, N. (2022). Remote poultry management system for small to medium scale producers using IoT. *Scientific African*, 18, Article e01398. <https://doi.org/10.1016/j.sciaf.2022.e01398>
- David, B., Moe, R., Michel, V., Lund, V., & Mejdell, C. (2015). Air quality in alternative housing systems may have an impact on laying hen welfare. Part I—Dust. *Animals*, 5(3), 495–511. <https://doi.org/10.3390/ani5030368>
- Gbadamosi, A. S. (2020). Cloud-based IoT monitoring system for poultry farming in Nigeria. *Arid Zone Journal of Engineering, Technology and Environment*, 16(1), 100–108.

- Giavarina, D. (2015). Understanding Bland–Altman analysis. *Biochemia Medica*, 25(2), 141–151. <https://doi.org/10.11613/BM.2015.015>
- Purswell, J. L., Dozier, W. A., Olanrewaju, H. A., Davis, J. D., Xin, H., & Gates, R. S. (2012). Effect of temperature-humidity index on live performance in broiler chickens grown from 49 to 63 days of age. 2012 IX International Livestock Environment Symposium (ILES IX), Paper ILES12–0265. <https://doi.org/10.13031/2013.41619>
- Kashiha, M., Pluk, A., Bahr, C., Vranken, E., & Berckmans, D. (2013). Development of an early warning system for a broiler house using computer vision. *Biosystems Engineering*, 116(1), 36–45. <https://doi.org/10.1016/j.biosystemseng.2013.06.004>
- Lara, L., & Rostagno, M. (2013). Impact of heat stress on poultry production. *Animals*, 3(2), 356–369. <https://doi.org/10.3390/ani3020356>
- Lin, L. I. K. (1989). A concordance correlation coefficient to evaluate reproducibility. *Biometrics*, 45(1), 255–268. <https://doi.org/10.2307/2532051>
- Marai, I. F. M., Habeeb, A. A. M., & Gad, A. E. (2002). Rabbits’ productive, reproductive and physiological performance traits as affected by heat stress: A review. *Livestock Production Science*, 78(2), 71–90. [https://doi.org/10.1016/S0301-6226\(02\)00091-X](https://doi.org/10.1016/S0301-6226(02)00091-X)
- Bland, J. M., & Altman, D. (1986). Statistical methods for assessing agreement between two methods of clinical measurement. *The Lancet*, 327(8476), 307–310. [https://doi.org/10.1016/S0140-6736\(86\)90837-8](https://doi.org/10.1016/S0140-6736(86)90837-8)
- Matthews, B. W. (1975). Comparison of the predicted and observed secondary structure of T4 phage lysozyme. *Biochimica et Biophysica Acta (BBA) - Protein Structure*, 405(2), 442–451. [https://doi.org/10.1016/0005-2795\(75\)90109-9](https://doi.org/10.1016/0005-2795(75)90109-9)
- Miles, D. M., Branton, S. L., & Lott, B. D. (2004). Atmospheric ammonia is detrimental to the performance of modern commercial broilers. *Poultry Science*, 83(10), 1650–1654. <https://doi.org/10.1093/ps/83.10.1650>
- Nawab, A., Ibtisham, F., Li, G., Kieser, B., Wu, J., Liu, W., Zhao, Y., Nawab, Y., Li, K., Xiao, M., & An, L. (2018). Heat stress in poultry production: Mitigation strategies to overcome the future challenges facing the global poultry industry. *Journal of Thermal Biology*, 78, 131–139. <https://doi.org/10.1016/j.jtherbio.2018.08.010>
- Norton, T., Chen, C., Larsen, M. L. V., & Berckmans, D. (2019). Precision livestock farming: Building “digital representations” to bring the animals closer to the farmer. *Animal*, 13(12), 3009–3017. <https://doi.org/10.1017/S175173111900199X>
- Ojo, R. O., Ajayi, A. O., Owolabi, H. A., Oyedele, L. O., & Akanbi, L. A. (2022). Internet of Things and machine learning techniques in poultry health and welfare management: A systematic literature review. *Computers and Electronics in Agriculture*, 200, Article 107266. <https://doi.org/10.1016/j.compag.2022.107266>
- Okinda, C., Nyalala, I., Korohou, T., Wang, J., Achieng, T., Wamalwa, P., Mang, T., & Shen, M. (2020). A review on computer vision systems in monitoring of poultry: A welfare perspective. *Artificial Intelligence in Agriculture*, 4, 184–208. <https://doi.org/10.1016/j.aiia.2020.09.002>
- Orakwue, S. I., Al-Khafaji, H. M. R., & Chabuk, M. Z. (2022). IoT based smart monitoring system for efficient poultry farming. *Webology*, 19(1), 4105–4112. <https://doi.org/10.14704/web/v19i1/web19270>
- Popoola, I. O., Oshibanjo, D. O., Popoola, O. R., & Okuneye, T. A. (2019). Effect of dietary electrolyte balance on water intake, litter moisture and production of broiler chickens. *Open Journal of Animal Sciences*, 9(4), 472–480. <https://doi.org/10.4236/ojas.2019.94036>
- Rai, A. C., Kumar, P., Pilla, F., Skouloudis, A. N., Di Sabatino, S., Ratti, C., Yasar, A., & Rickerby, D. (2017). End-user perspective of low-cost sensors for outdoor air pollution monitoring. *Science of the Total Environment*, 607–608, 691–705. <https://doi.org/10.1016/j.scitotenv.2017.06.266>

- Ritz, C. W., Fairchild, B. D., & Lacy, M. P. (2004). Implications of ammonia production and emissions from commercial poultry facilities: A review. *Journal of Applied Poultry Research*, 13(4), 684–692. <https://doi.org/10.1093/japr/13.4.684>
- Sonaiya, E. B. (2007). Family poultry, food security and the impact of HPAI. *World's Poultry Science Journal*, 63(1), 132–138. <https://doi.org/10.1017/S0043933907001353>
- Tao, X., & Xin, H. (2003). Acute synergistic effects of air temperature, humidity, and velocity on homeostasis of market-size broilers. *Transactions of the ASAE*, 46(2), 491–497. <https://doi.org/10.13031/2013.12971>
- Wilcoxon, F. (1945). Individual comparisons by ranking methods. *Biometrics Bulletin*, 1(6), 80–83. <https://doi.org/10.2307/3001968>

How to cite this Article: Igwe, P. E., & Iyoo, G. M. (2026). *An Internet of Things Based Early Warning System for Monitoring Environmental Conditions and Health Risks in Nigerian Poultry Houses. International Journal of Engineering, Science, and Environmental Research (IJESER)*, 2(1), 262–281.