IOT-BASED ENVIRONMENTAL STRESS INDEX AND LLM DIAGNOSTICS FOR LAYING QUAIL EGG PRODUCTION DURING PANCAROBA IN SIWATU VILLAGE, BATANG REGENCY, CENTRAL JAVA



Conduct by:

Name : Supriyanto
Student ID : 22523294

PROGRAM STUDI INFORMATIKA – PROGRAM SARJANA
FAKULTAS TEKNOLOGI INDUSTRI
UNIVERSITAS ISLAM INDONESIA
2026

SUPERVISOR ENDORSEMENT PAGE

IOT-BASED ENVIRONMENTAL STRESS INDEX AND LLM DIAGNOSTICS FOR LAYING QUAIL EGG PRODUCTION DURING PANCAROBA IN SIWATU VILLAGE, BATANG REGENCY, CENTRAL JAVA

THESIS

Conduct By:

Name : Supriyanto

Student ID : 22523294

Yogyakarta, August 3st, 2026

Advisor,

(Mukhammad Andri Setiawan, S.T., M.Sc., Ph.D.)

EXAMINER ENDORSEMENT PAGE

IOT-BASED ENVIRONMENTAL STRESS INDEX AND LLM DIAGNOSTICS FOR LAYING QUAIL EGG PRODUCTION DURING PANCAROBA IN SIWATU VILLAGE, BATANG REGENCY, CENTRAL JAVA

THESIS

Has been defended in front of the examiners as one of the requirements to obtain a Bachelor of Informatics degree from the Undergraduate Program in Informatics at the Faculty of Industrial Technology, Universitas Islam Indonesia

Yogyakarta, 3rd August 2026

Chair

Mukhammad Andri Setiawan, S.T., M.Sc., Ph.D. _____

Examiner 1

Ahmad Fathan Hidayatullah, S.T., M.Cs., Ph.D. _____

Examiner 2

Irving Vitra Paputungan, S.T., M.Sc., Ph.D _____

Acknowledged by,

Head of Undergraduate Program in Informatics

Faculty of Industrial Technology

Universitas Islam Indonesia

(DHOMAS HATTA FUDHOLI, S.T., M.Eng., Ph.D.)

AUTHENTICITY STATEMENT

The undersigned:

Name : Supriyanto

Student ID : 22523294

Final project with title:

IOT-BASED ENVIRONMENTAL STRESS INDEX AND LLM DIAGNOSTICS FOR LAYING QUAIL EGG PRODUCTION DURING PANCAROBA IN SIWATU VILLAGE, BATANG REGENCY, CENTRAL JAVA

Stating that all components and contents in this final project are my own work. If in the future it is proven that some parts of this work are not my own work, the final project submitted as my own work is ready to be withdrawn and ready to bear any risks and consequences.

Thus this statement letter is made, hopefully it can be used properly.

Yogyakarta, August 3st, 2026



(Supriyanto)

DEDICATION

Bismillahirrahmanirrahim

All praise and gratitude belong to Allah Swt, for it is through His abundant grace, mercy, and guidance that this thesis has been brought to completion. May peace and blessings be forever upon the Prophet Muhammad saw., the finest example for all of humanity.

With all sincerity, the author dedicates this work to the most beloved people Bapak, Ibu, Mas, and Mbak who have always been a place of refuge through every step of this journey. This thesis was born from prayers that never ceased, offered in every prostration, and above all, from the quiet supplications of a mother in the stillness of the last third of the night, asking nothing but the very best for her child. Your love, sacrifice, and unwavering support are the greatest force that has carried the author to this very moment. May this work stand as a small testament of devotion and a heartfelt expression of gratitude. This achievement was never the author's alone it belongs equally to your prayers, your love, your blessings, and above all, to the will of Allah Swt.

This work is also dedicated to the author's own self. Thank you for holding on. Thank you for fighting through every obstacle and refusing to surrender when the road grew difficult. Thank you for every ounce of patience chosen over despair, every tear quietly wiped away, and every flicker of determination kept alive through the hardest of days. The journey was never easy but you made it through.

This thesis, entitled *"IoT-Based Environmental Stress Index and LLM Diagnostics for Laying Quail Egg Production During Pancaroba in Siwatu Village, Batang Regency, Central Java,"* was born from a deeply personal story. Since the author's years in senior high school, a quail farming business was built and developed from the ground up from nothing, through sheer determination and hard work and it was precisely from that self-built farm that the author found the means to pursue an education and arrive at this very moment. It is only fitting, then, that the knowledge gained through this journey be given back to those who made it possible. Through this research, the author hopes to offer a meaningful contribution to small and medium-scale farmers, particularly in mitigating the impact of environmental changes on livestock health and productivity. It is the author's deepest hope that the findings of this work may one day be felt and embraced widely, and that it may, in some small way, ease the burdens carried by farmers across every corner of this nation.

MOTTO

“For indeed, with hardship [will be] ease. Indeed, with hardship [will be] ease. So when you have finished [your duties], then stand up [for worship].”
(QS. Al-Insyirah: 5-7)

“Indeed, Allah will not change the condition of a people until they change what is in themselves”
(QS. Ar-Ra'd: 11)

“Sebaik-baik manusia adalah yang paling bermanfaat bagi manusia lainnya.”
(HR. Ahmad, ath-Thabrani, ad-Daruqutni)

“Hidup yang tidak dipertaruhkan tidak akan pernah dimenangkan”
Sutan Sjahrir

FOREWORD

Bismillahirrahmanirrahim.

Alhamdulillahirabbil 'alamin, all praise and gratitude belong to Allah Swt. for His abundant grace, mercy, and guidance that have allowed the author to complete this final thesis report. This report was prepared as one of the requirements to obtain a Bachelor's degree in the Informatics Study Program. The author acknowledges that throughout the process of completing this report, numerous obstacles and challenges had to be faced and overcome. However, through the support, prayers, and guidance of many wonderful people, this report has finally been brought to completion.

On this occasion, with the utmost sincerity, the author would like to express the deepest gratitude to:

1. Allah Swt., for the gift of health, strength, and ease that have enabled the author to complete this final thesis report.
2. The author's most beloved Bapak, Mae, Mas, and Mbak, whose prayers have never once ceased, whose support has always been given with the most sincere and wholehearted love, and whose sacrifices both moral and material have been the very foundation that carried the author through every year of study and to the completion of this thesis. Everything that has been achieved belongs as much to you as it does to the author.
3. Mr. Mukhammad Andri Setiawan, as the supervising lecturer, who has generously given his time, energy, and thought, and who with great patience has provided direction, guidance, and knowledge to the author throughout the completion of this report.
4. Mr. Wisnu Manupraba and Mrs. Novi Setiani, who guided and mentored the author during the internship at Javan Cipta Solusi. Thank you for the direction, motivation, and the beautiful memories that were created together. Gratitude is also extended to all mentors and colleagues at Javan Cipta Solusi for the invaluable experiences shared.
5. The author's closest friends Ahmad, Azis, Ihya, Novera, Mifta, Aan, Mu'is, Ibnu, Raya, Umay, Amel, Fazil, Salma, Zaky, Nang, Abe, Zaidan, Nisa, Edo, Hana, and all other friends who cannot be mentioned one by one who have accompanied the author through every step of this academic journey, sharing both the joys and the struggles, and who have never stopped offering encouragement and motivation

throughout the completion of this thesis. The togetherness and bonds forged with all of you are among the most cherished things the author carries from these years of study.

6. The entire Comdev family and PPK Ormawa Sinera, for the sense of belonging, the organizational experiences, and the spirit that was always shared with the author throughout the years of study.
7. Afa, Naja, Gabriel, Adi, Adam, Bagas, Fani, Novi, Ima, and all members of Forsimba batch 2022 as well as the entire Forsimba family, who have written so many stories, experiences, and beautiful memories into the author's life throughout this period of study.
8. The entire UNISI Football family, for the wealth of experiences, the sense of camaraderie, and the invaluable lessons in teamwork and solidarity that will never be forgotten.

The author acknowledges that this final thesis report is far from perfect. Therefore, any constructive criticism and suggestions are greatly welcomed for the sake of improvement in the future. It is the author's sincere hope that this report may prove beneficial to its readers and contribute to the advancement of knowledge, particularly in the fields of technology and animal husbandry.

Yogyakarta, August 3st, 2026



(Supriyanto)

ABSTRACT

Laying quail (*Coturnix japonica*) production in Batang Regency, Central Java faces recurring stress during the pancaroba season the wet-to-dry transitional period marked by unpredictable temperature swings, persistently high humidity, and ammonia accumulation inside cages. Small-scale farmers in the region rely on visual inspection and smell to assess cage conditions, meaning that production drops are discovered after the fact rather than anticipated. This research designed, built, and tested an integrated monitoring system combining Internet of Things sensors, a composite Environmental Stress Index (ESI), and a Claude Sonnet Large Language Model (LLM) diagnostic layer, deployed at a working laying quail farm in Siwatu Village, Wonotunggal District, Batang Regency, Central Java.

The system used an ESP32 microcontroller connected to temperature, humidity, ammonia, and light intensity sensors that transmitted readings every 30 minutes to a cloud backend. Each night the backend computed a daily ESI score a weighted composite of thermal load, ammonia concentration, and light adequacy classified the result into Good, Needs Attention, or Critical, then sent the summary to Claude Sonnet to generate a structured diagnosis covering cage status, identified problems, and recommended actions. Results were delivered to the farmer through a mobile application on his Android phone.

Field deployment across 30 valid days covering approximately 1,800 laying quail showed that environmental conditions throughout the pancaroba window stayed consistently above safe thresholds, with the daily ESI averaging 0.53 and not a single day qualifying as Good 83.3% of days fell in the Needs Attention band and 16.7% in Critical. Statistical analysis produced a significant negative correlation between the daily ESI and daily egg production ($r = -0.361$, $p = 0.050$), confirming that the composite index tracks laying performance under real field conditions. The LLM diagnostic layer achieved 100% sensitivity for Critical days and consistently erred toward greater caution rather than less, a pattern clinically safer than conventional threshold-only monitoring. A documented intervention episode near the end of the observation period showed a 35% reduction in ammonia concentration following an Emergency-level alert, providing direct field evidence that the system prompted timely corrective action. All seven functional test scenarios passed.

Keywords: Internet of Things, Environmental Stress Index, Large Language Model, laying quail, pancaroba season, egg production, Batang Regency.

GLOSSARY

ESP32	a low-power microcontroller by Espressif Systems with integrated Wi-Fi and Bluetooth capabilities, used as the main processing unit of the IoT system in this research.
BH1750	a digital ambient light intensity sensor that measures illuminance in lux units, communicates via the I2C protocol, and is used in this research to monitor environmental light levels.
DHT22	a high-accuracy digital temperature and humidity sensor using a single-wire communication protocol, used in this research to monitor environmental conditions in real-time.
DS3231	a high-precision Real-Time Clock (RTC) module that uses the I2C protocol to store and provide accurate time information even when the device is powered off.
MQ137	a gas sensor that detects ammonia (NH_3) concentration in the air, used in this research to monitor air quality in livestock enclosures or specific environments.
Relay 2 Channel	an electronic module containing two electromechanical switches controlled by a microcontroller to automatically turn high-voltage electrical devices on or off.
Java Spring Boot	a Java-based backend development framework that simplifies server-side application creation with minimal configuration, used as the API backend in this research.
React Native	a cross-platform mobile application development framework by Meta that uses JavaScript, allowing a single codebase to run on both Android and iOS.
PostgreSQL	an open-source relational database management system using SQL, used in this research as the primary storage for sensor data and user data.
VPS(Vitual Private Server)	a virtual server running on a physical machine through virtualization technology, used in this research to host the Java Spring Boot backend and make it publicly accessible.

TABLE OF CONTENTS

TITLE i
SUPERVISOR ENDORSEMENT PAGE ii
EXAMINER ENDORSEMENT PAGE iii
AUTHENTICITY STATEMENT iv
DEDICATION v
MOTTO vi
FOREWORD vii
ABSTRACT ix
GLOSSARY x
TABLE OF CONTENTS xi
LIST OF TABEL xiv
LIST OF FIGURES xv
CHAPTER I INTRODUCTION 1
1.1 Background 1
1.2 Problem Formulation 3
1.3 Scope and Limitations 4
1.4 Research Objectives 4
1.5 Research Benefits 5
1.6 Research Methodology 5
1.7 Writing Systematics 7
CHAPTER II LITERATURE REVIEW 9
2.1 Japanese Quail and Its Sensitivity to the Cage Environment 9
2.2 Pancaroba and the Compounding of Environmental Stress 10
2.3 The Internet of Things in Precision Livestock 10
2.4 Hardware Used in the Sensing Tier 11
2.5 Software Stack 13
2.6 *Large Language Model* (LLM) and Claude Sonnet 14
2.7 Composite Indices and the *Temperature-Humidity Index* (THI) 14
2.8 The Environmental Stress Index 14
2.9 Methods Used to Evaluate the System 16
2.10 Where the Problem Sits in the Existing Work 16
2.11 Choosing Between the Available Approaches 19
2.12 Conceptual Framework 19
CHAPTER III RESEARCH METHODOLOGY 20
3.1 Analysis phase 20
3.1.1 Research site and observation period 20
3.1.2 Requirements analysis 21
3.1.3 Data Collection Design 23
3.2 System Design Phase 24
3.2.1 System Architecture Design 24
3.2.2 Hardware Wiring Design 26
3.2.3 Firmware Logic Design 27
3.2.4 Database Design 28
3.2.5 REST API Endpoint Design 30
3.2.6 ESI Computation Design 35
3.2.7 LLM Diagnostic Pipeline Design 36
3.2.8 Mobile Application UI Design 37

3.3 Execution Phase 38
3.3.1 Hardware assembly and installation 38
3.3.2 Backend deployment 39
3.3.3 Mobile application deployment 40
3.3.4 Field data collection execution 40
3.4 Testing Phase 41
3.4.1 Black-box functional testing 42
3.4.2 Statistical analysis procedure 42
3.4.3 LLM consistency analysis 43
3.4.4 Expert validation protocol 44
3.4.5 Case study documentation 46
CHAPTER IV RESULT AND DISCUSSION 47
4.1 Deployment outcome and data collection summary 47
4.1.1 What was deployed and what ran 47
4.2 Environmental characteristics during pancaroba 48
4.2.1 Descriptive statistics across the 30 valid days 48
4.2.2 Daily evolution of each parameter 49
4.2.3 Composite thermal load through THI 50
4.3 Environmental Stress Index distribution and behaviour 51
4.3.1 Descriptive statistics across the 30 valid days 51
4.3.2 Distribution shape 52
4.3.3 Sub-score contributions 52
4.3.4 Anatomy of the five Critical days 54
4.4 Relationship between ESI and egg production 57
4.4.1 Egg production during the observation window 58
4.4.2 Pearson correlation: composite ESI vs egg count 59
4.4.3 Channel-wise correlations to identify the strongest predictor 60
4.4.4 What this means for the ESI framework 61
4.5 LLM diagnostic performance 62
4.5.1 Agreement between ESI category and LLM cage status 62
4.5.2 Asymmetric precaution the direction of every disagreement 63
4.5.3 Representative LLM outputs across a range of conditions 64
4.5.4 Why the LLM over-detected 65
4.5.5 Productivity did not differ significantly between consistency groups 65
4.6 Documented intervention case study 20 to 24 May 65
4.6.1 What the sensors and system reported 66
4.6.2 The peak and the recovery 66
4.6.3 What the recovery does and does not prove 67
4.7 Secondary finding: LLM over-escalation on rainfall days 67
4.7.1 Three days where the LLM and the birds disagreed 68
4.7.2 What actually happened the rain signature 68
4.7.3 What this reveals about the pipeline 69
4.8 System functionality verification 69
4.8.1 Black-box functional test results 70
4.8.2 Interview-based validation 71
4.9 Synthesis and limitations 72
4.9.1 ESI and LLM as complementary layers 73
4.9.2 Limitations 73
CHAPTER V CONCLUSIONS AND RECOMMENDATIONS 76
5.1 Conclusions 76

5.2 Recommendations 78

REFERENCE 81
APPENDIX 84

LIST OF TABEL

Table 2.1 Hardware specifications 13
Table 2.2 ESI bands and the corresponding action recommendation 16
Table 2.3 Component thresholds and their literature sources 16
Table 2.4 Comparison of this study with existing related work 18
Table 2.5 Comparison of interpretation approaches 19
Table 3.1 Summary of interview findings from the requirements analysis phase 22
Table 3.2 Selected parameters and their physiological justification 23
Table 3.3 Data collection summary 24
Table 3.4 GPIO pin mapping on the ESP32 26
Table 3.5 Database table descriptions 29
Table 3.6 Authentication and registration 31
Table 3.7 User Profile 31
Table 3.8 Sensor data IOT 31
Table 3.9 Livestock Management 32
Table 3.10 Feed stock management 32
Table 3.11 Feed Usage 32
Table 3.12 Egg income 33
Table 3.13 Egg sales 33
Table 3.14 Quail reduction 33
Table 3.15 Operating cost 33
Table 3.16 Dashboard 34
Table 3.17 AI diagnosis LLM pipeline 34
Table 3.18 Community forum 34
Table 3.19 Expected fields in the LLM JSON response 37
Table 3.20 Mobile application tabs and their content 38
Table 3.21 Server environment specifications 39
Table 3.22 Data collection log summary 41
Table 3.23 Black-box functional test scenarios 42
Table 3.24 Interpretation criteria for correlation strength 43
Table 3.25 LLM vs ESI classification comparison framework 44
Table 3.26 Interview instrument 1 pre-deployment farmer 44
Table 3.27 I Interview instrument 2 animal husbandry student 45
Table 3.28 Interview instrument 3 post-deployment farmer 45
Table 4.1 Data collection totals 47
Table 4.2 Descriptive statistics of the four monitored channels 48
Table 4.3 ESI classification distribution across 30 days 51
Table 4.4 ESI classification distribution across 30 days 54
Table 4.5 Counterfactual ESI with ammonia held at the non-Critical mean 55
Table 4.6 Critical days compared with Needs Attention days 56
Table 4.7 Channel-wise Pearson correlation with daily egg production 60
Table 4.8 Confusion matrix ESI category vs LLM cage status 62
Table 4.9 Direction of ESI-LLM disagreements 63
Table 4.10 Representative LLM diagnostic outputs 64
Table 4.11 Case study timeline, 20–24 May 2026 66
Table 4.12 Days where the LLM over-escalated 68
Table 4.13 Black-box functional test results 70

LIST OF FIGURES

Figure 2.1 IoT three-layer architecture 11
Figure 3.1 Research site, Siwatu village 20
Figure 3.2 IoT-ESI-LLM system architecture 25
Figure 3.3 Sensor unit wiring 26
Figure 3.4 Firmware main loop flowchart 27
Figure 3.5 Lighting control logic flowchart 28
Figure 3.6 Production database ERD 29
Figure 3.7 ESI computation flowchart 35
Figure 3.8 LLM diagnostic pipeline 36
Figure 4.1 Daily cage environmental profile 49
Figure 4.2 Daily THI with stress thresholds 50
Figure 4.3 Daily ESI values for the 30-day observation window 51
Figure 4.4 ESI distribution histogram 52
Figure 4.5 ESI sub-score composition 53
Figure 4.6 Egg production vs ESI (dual axis) 58
Figure 4.7 ESI vs egg count scatter plot 59
Figure 4.8 Correlation matrix of environmental parameters and egg production 60
Figure 4.9 LLM vs ESI agreement matrix 62
Figure 4.10 LLM consistency and egg production 63
Figure 4.11 Case study 20–24 May 66
Figure 4.12 Rainfall day sensor profiles 68

CHAPTER I
INTRODUCTION

1.1 Background

Quail (*Coturnix japonica*) farming represents one of the poultry sectors that plays a significant role in providing animal protein and serves as a source of household income for smallholder farmers in Indonesia, particularly in the Central Java region. Quail offer several advantages, including a relatively short production cycle, efficient space and feed requirements, and high laying capacity that can reach hundreds of eggs per bird per year (Badan Pusat Statistik, 2025)**.** These characteristics make quail a strategic commodity for smallholder farmers in rural areas seeking to enhance economic productivity with limited capital.

However, quail productivity is highly dependent on cage environmental conditions, which must remain within specific comfort zones. Previous studies have shown that the thermal comfort zone for laying quail lies within a temperature range of 20–25 °C and a relative humidity of 50–70% (El-Tarabany, 2016; Alagawany et al., 2017)**,** while the ammonia concentration in cage air should not exceed 10 ppm to maintain respiratory health (Naseem & King, 2018)**.** Environmental conditions outside these ranges consistently reduce egg production by 6–21% depending on severity and duration, weaken the immune system, and impair reproductive performance through decreased feed intake and disrupted ovarian function (Alagawany et al., 2017).

The challenges of maintaining cage environmental stability become increasingly complex during the seasonal transition period locally known as pancaroba. The pancaroba period in Central Java is characterized by extreme and irregular weather changes, including sudden afternoon rainfall events that can cause temperature drops of several degrees within hours, humidity saturation, and unstable ventilation conditions that trigger ammonia accumulation inside enclosed housing (Alagawany et al., 2017; Naseem & King, 2018). The specific environmental dynamics of pancaroba in the Batang coastal region were further confirmed during the field deployment described in Chapter IV of this study. These rapid and persistent environmental fluctuations impose cumulative physiological stress on laying quail, which ultimately leads to a direct decline in egg productivity.

Smallholder quail farmers in Siwatu Village, Wonotunggal District, Batang Regency, Central Java, still rely heavily on visual inspection and empirical experience to evaluate cage conditions. As a result, corrective actions are typically taken only after a drop in egg production

is observed the following day, by which point the impact of environmental stress has already occurred. Limited instrumentation and technical expertise make it difficult for farmers to monitor cage conditions objectively and responsively, particularly during the pancaroba period, which demands rapid responses to microclimate changes.

The development of Internet of Things (IoT) technology has opened new opportunities for the development of automated and continuous environmental monitoring systems for poultry cages. Several previous studies have developed IoT prototypes for poultry farm monitoring using low-cost microcontrollers such as the ESP32, integrated with various environmental sensors (Lashari et al., 2018; Gunawan et al., 2021)**.** However, most existing systems only display sensor data per parameter on separate dashboards with fixed threshold-based alerting. This approach still requires farmers to mentally integrate multiple values a cognitively demanding task when several variables shift simultaneously during pancaroba.

The composite index approach has long been used in animal husbandry and environmental health, with the Temperature-Humidity Index (THI) serving as a classic example widely applied in ruminant and poultry farming (Lara & Rostagno, 2013). In the field of digital agriculture, the composite index approach has been shown to improve data interpretability for non-expert users by combining multiple sensor parameters into a single score normalized to the 0–1 range, allowing faster decision-making compared to monitoring each parameter separately (Lara & Rostagno, 2013). Nevertheless, no composite environmental index has yet been specifically developed for laying quail that integrates thermal stress, air quality, and light intensity into a single metric.

In addition to the limitations in data representation, conventional IoT systems have also not been able to translate sensor data into contextual and easily understandable action recommendations for farmers. Recent advances in artificial intelligence, particularly Large Language Models (LLMs), offer an opportunity to bridge this gap by generating natural-language diagnostics based on processed sensor data (Menezes et al., 2025; Yuan et al., 2024). The integration of LLMs into the IoT pipeline enables the transformation of quantitative data into actionable diagnostic narratives, allowing farmers to take corrective actions before the impact of environmental stress manifests in declining egg production

Based on the literature review, the integration of a composite environmental index with LLM-based diagnostics for quail cage monitoring during the seasonal transition period has not been empirically addressed. Most existing studies still focus on one of the components, namely IoT systems with threshold-based alerting or LLM exploration in isolation from field sensor

systems. Furthermore, the empirical validation of the relationship between a composite environmental index and quail egg productivity under actual pancaroba conditions in the field remains very limited.

These observations indicate a research gap that needs to be addressed. Therefore, this study aims to design, develop, and evaluate an IoT-based environmental monitoring system integrated with a Large Language Model for quail farming during the seasonal transition period. The proposed system consists of a three-layer architecture, namely a data acquisition layer using an ESP32 microcontroller with a DHT22 sensor (temperature and humidity), an MQ-137 sensor (ammonia), a BH1750 sensor (light intensity), a DS3231 Real-Time Clock module, and a 2-Channel Relay for lighting control; a processing layer based on Java Spring Boot with a PostgreSQL database hosted on a Virtual Private Server (VPS); and an interpretation layer that uses Claude Sonnet through the Anthropic API to generate diagnostics. The processed data is then presented to farmers through a React Native mobile application. As the main metric, this system introduces the Environmental Stress Index (ESI), a composite index on a 0–1 scale that combines THI, ammonia, and light intensity components with literature-based weighting. The findings of this study are expected to provide empirical contributions to the development of decision-support systems for smallholder quail farmers facing environmental challenges during the pancaroba period.

1.2 Problem Formulation

Based on the background described above, the problem formulations of this study are as follows:

a. How to design and implement an IoT-based quail cage environmental monitoring system integrated with a Large Language Model to automatically generate cage condition diagnostics?
b. How to formulate the Environmental Stress Index (ESI) as a composite index that combines the parameters of temperature, humidity, ammonia concentration, and light intensity for laying quail farming?
c. How is the relationship between the daily ESI value and the level of quail egg production during the seasonal transition period (pancaroba) in Batang Regency, Central Java?
d. How is the diagnostic performance of the Claude Sonnet Large Language Model in interpreting environmental sensor data and generating action recommendations for farmers through a React Native-based mobile application?

1.3 Scope and Limitations

To keep the research focused and systematic, the following scope and limitations are established:

a. The research object is limited to one laying quail enclosure containing approximately 1,800 quail in Siwatu Village, Wonotunggal District, Batang Regency, Central Java.
b. The observation period is conducted over 30 consecutive days during the seasonal transition period (pancaroba), namely from 26 April to 25 May 2026.
c. The monitored environmental parameters include air temperature, relative humidity, ammonia concentration, and light intensity, with a sampling interval of 30 minutes.
d. The hardware used consists of an ESP32 microcontroller, a DHT22 sensor, an MQ-137 sensor, a BH1750 sensor, a DS3231 Real-Time Clock module, and a 2-Channel Relay module for cage lighting control.
e. The system backend is developed using the Java Spring Boot framework with a PostgreSQL database deployed on a Virtual Private Server (VPS).
f. The mobile application is developed using the React Native framework and is intended for devices running on the Android operating system.
g. The Large Language Model used is Claude Sonnet, accessed through the Anthropic API, with input in the form of daily summaries generated from aggregated sensor data.
h. This study does not address aspects of quail genetics, feed composition, veterinary health management, or marketing strategies for egg products.

1.4 Research Objectives

The objectives of this study are as follows:

a. To design and implement an IoT-based quail cage environmental monitoring system integrated with the Claude Sonnet Large Language Model to generate automated diagnostics based on sensor data.
b. To formulate and implement the Environmental Stress Index (ESI) as a composite index that combines temperature, humidity, ammonia, and light intensity parameters for laying quail farming.
c. To examine the empirical relationship between daily ESI values and the number of quail eggs produced during the seasonal transition period at the research site.

d. To evaluate the diagnostic performance of the Claude Sonnet Large Language Model in generating actionable recommendations that can be delivered to farmers through a React Native mobile application.

1.5 Research Benefits

a. Theoretical Benefits.

This study contributes to the development of research on IoT systems integrated with LLMs for precision livestock farming, particularly by introducing the Environmental Stress Index as a composite index empirically validated in the context of laying quail in a tropical climate during the seasonal transition period. As detailed in the comparative analysis in Section 2.10 (Table 2.4), this contribution is distinct from prior IoT–poultry monitoring work in three respects: it proposes the first composite environmental index formulated specifically for laying quail, it validates that index against field-measured egg production, and it documents two previously unreported behavioural patterns of an LLM operating inside an agricultural IoT diagnostic pipeline.

b. Practical Benefits.

The results of this study can serve as a decision-support tool for smallholder quail farmers in monitoring cage environmental conditions objectively and responsively, enabling corrective actions to be taken before the impact of environmental stress manifests in declining egg production.

c. Academic Benefits.

This study can serve as a reference for future research aiming to develop IoT systems based on LLMs in the agricultural and livestock domains in tropical regions, particularly those that integrate composite environmental indices with diagnostics based on generative artificial intelligence.

1.6 Research Methodology

The methodology employed in this study follows the Research and Development (R&D) approach with the ASET model (Analysis, System Design, Execution, Testing), consisting of several systematic stages. The stages of the research methodology are as follows:

a. Literature Review

This stage is conducted to gather references related to quail physiology, environmental stress indices for poultry, IoT systems for precision livestock farming, and the application of Large Language Models in the agricultural domain. The literature review aims to

establish a theoretical foundation and identify relevant research gaps. The sources used include scientific journals, conference proceedings, and other academic publications published between 2021 and 2026.

b. Requirements Analysis

The requirements analysis stage is conducted to determine the hardware components, software, and environmental parameters relevant to the research objectives. At this stage, interviews are also conducted with the partner farmer in Siwatu Village and an animal husbandry expert to validate the physiological parameters and thresholds to be used in formulating the index.

c. System Design

At this stage, the three-layer system architecture is designed, consisting of a data acquisition layer based on the ESP32 with DHT22, MQ-137, BH1750, DS3231, and 2-Channel Relay; a processing layer based on Java Spring Boot with a PostgreSQL database deployed on a VPS; and an interpretation layer that utilizes Claude Sonnet through the Anthropic API. This stage also includes the design of the database structure, REST API endpoints, the Environmental Stress Index formula, and the user interface design of the React Native mobile application.

d. System Implementation

The implementation stage is carried out by developing ESP32 firmware for sensor data acquisition and transmission to the backend, developing the Spring Boot backend along with daily aggregation logic and ESI computation, integrating Claude Sonnet through the Anthropic API for daily diagnostics, and building the React Native mobile application with five main tabs (Dashboard, Cage, Community, Medical, and Account).

e. Testing and Data Collection

System testing is conducted in the field at the partner cage for 30 consecutive days during the pancaroba period. Functional testing of the system is performed using the black-box testing method, while environmental data collection is conducted continuously every 30 minutes. Daily egg production data is recorded manually by the farmer each morning.

f. Data Processing and Analysis

The collected sensor data is processed through daily aggregation to generate daily ESI values. The relationship between daily ESI and egg production is analyzed using Pearson correlation at a significance level of $\alpha = 0.05$. Additional analyses are conducted to

evaluate the consistency between ESI classification and LLM diagnostics, as well as documented case studies of extreme environmental events.

g. Conclusion

The final stage of the research is drawing conclusions based on the quantitative and qualitative analyses that have been conducted. The conclusions are used to answer the research problem formulations and provide recommendations for system development and future research.

1.7 Writing Systematics

This section explains the writing structure of the research report as a general overview to help readers understand the flow of discussion. This report is organized into five interrelated chapters, with the following details:

CHAPTER I INTRODUCTION

This chapter discusses the research background, problem formulation, scope and limitations, research objectives, research benefits, research methodology, and writing systematics related to the development of an IoT-based quail cage environmental monitoring system integrated with the Claude Sonnet Large Language Model

CHAPTER II LITERATURE REVIEW

This chapter outlines the theories and concepts that form the foundation of the research, including the physiological characteristics of laying quail, cage environmental parameters, the Temperature-Humidity Index (THI), Internet of Things technology, the ESP32 microcontroller, environmental sensors (DHT22, MQ-137, BH1750, DS3231), the Java Spring Boot framework, the PostgreSQL database, the React Native framework, and Large Language Models with a particular focus on the Claude Sonnet API.

CHAPTER III RESEARCH METHODOLOGY

This chapter explains in detail the stages of the research methodology used, including system requirements analysis, three-layer architecture design, formulation of the Environmental Stress Index, database design, API design, mobile application interface design, and system testing scenarios in the field.

CHAPTER IV RESULTS AND DISCUSSION

This chapter presents the results of system implementation, environmental data collected over 30 days, analysis of the ESI value distribution, Pearson correlation analysis between ESI

and egg production, evaluation of Claude Sonnet's diagnostic performance, a case study of farmer intervention, and discussion of the findings obtained.

CHAPTER V CONCLUSIONS AND RECOMMENDATIONS

This chapter contains the conclusions that answer the research problem formulations and provides recommendations that can be considered for system development and future research in the domain of precision livestock farming based on IoT and LLM.

CHAPTER II
LITERATURE REVIEW

This chapter reviews the body of knowledge underpinning the development of an IoT-based environmental monitoring system that is coupled with a Large Language Model for quail farming during the seasonal transition. The review starts by defining the technical terms that appear in the title and the research questions, before moving on to the supporting theories, the field context in which the problem occurs, and a comparison of the technological alternatives that were considered. A conceptual framework at the end of the chapter ties these threads together.

2.1 Japanese Quail and Its Sensitivity to the Cage Environment

Japanese quail (*Coturnix japonica*) occupy a small but economically meaningful niche in Indonesian smallholder poultry. The birds mature quickly, need little floor space, eat modestly, and a single hen can produce in the order of three hundred eggs a year (Badan Pusat Statistik, 2025). For a household in rural Central Java with limited capital, this combination is attractive a small flock can begin returning income within weeks of arrival

The same physiology that produces these returns, however, leaves the birds with very little margin for environmental error. Reported comfort bands sit between 20 and 25 °C with relative humidity of fifty to seventy percent (El-Tarabany, 2016; Alagawany et al., 2017). Once ambient temperature settles in the upper twenties, feed intake drops, body weight declines, and egg production falls studies report reductions of 6–21% depending on the intensity and duration of heat exposure (Alagawany et al., 2017). El-Tarabany (2016) showed that Japanese quail exposed to high THI levels (76–80) had significantly lower fertility rates and poorer egg quality compared to birds kept below THI 70. The mechanism that ultimately suppresses egg output runs through the hypothalamic–pituitary–gonadal axis reduced feed intake, elevated corticosterone, and impaired ovarian function rather than through any single visible symptom; a farmer watching the cage will usually see only the consequence on the next morning's tray. The mechanism that ultimately suppresses egg output therefore runs through the hypothalamic–pituitary–gonadal axis rather than through any single visible symptom a farmer watching the cage will usually see only the consequence on the next morning's tray.

Air chemistry inside the housing matters in much the same way. Welfare guidelines treat ten parts per million as the working ceiling for ammonia, and exposure beyond twenty-five

parts per million is associated with acute damage to the respiratory epithelium (Naseem & King, 2018). Photoperiod is the third axis. Layers require fourteen to sixteen hours of illumination per day, balanced by a genuine dark interval long enough to allow melatonin synthesis (Naeem et al., 2018).

2.2 Pancaroba and the Compounding of Environmental Stress

The Indonesian term *pancaroba* names the changeover between the wet and dry seasons. In coastal Central Java the period is recognisable from its weather signature: a hot morning, an abrupt afternoon thunderstorm, temperature falling five to seven degrees in under two hours, then prolonged saturation as the cloud cover persists. Ventilation becomes unpredictable, and damp manure decomposes faster, releasing ammonia at concentrations that would not occur in either of the stable seasons that flank it (Naseem & King, 2018).

What makes the period difficult for layer cages is not any one of these effects in isolation but their overlap. Birds experience a transient shock from the temperature drop while still carrying the chronic burden of elevated humidity and accumulating ammonia. The cumulative physiological burden from simultaneous thermal stress, high humidity, and ammonia exposure is what suppresses reproductive performance and egg output (Alagawany et al., 2017; El-Tarabany, 2016), and it is the reason farmers in the Batang area commonly identify pancaroba as the lowest-producing window of the year a view that the participating farmer in this study echoed during the pre-installation interview.

2.3 The Internet of Things in Precision Livestock

The Internet of Things refers to networks of physical objects that carry sensors, run small amounts of software, and exchange data through ordinary internet connectivity. In livestock contexts, the relevance is concrete: cages can be instrumented to record environmental conditions continuously and at low cost (Lashari et al., 2018).

Implementations in poultry over the last five years have shown that the hardware side of this is workable on small budgets. Lashari et al. (2018) and Gunawan et al. (2021) both report functioning cage telemetry from low-cost microcontrollers paired with consumer-grade sensors. What these systems typically do not solve is the interpretation problem on the farmer's side. Most published dashboards show one chart per parameter and trigger an alert when any single value crosses a fixed cut-off a perfectly reasonable design that nonetheless leaves the

operator with the cognitive task of integrating multiple readings during precisely the window when several of them are drifting at once.

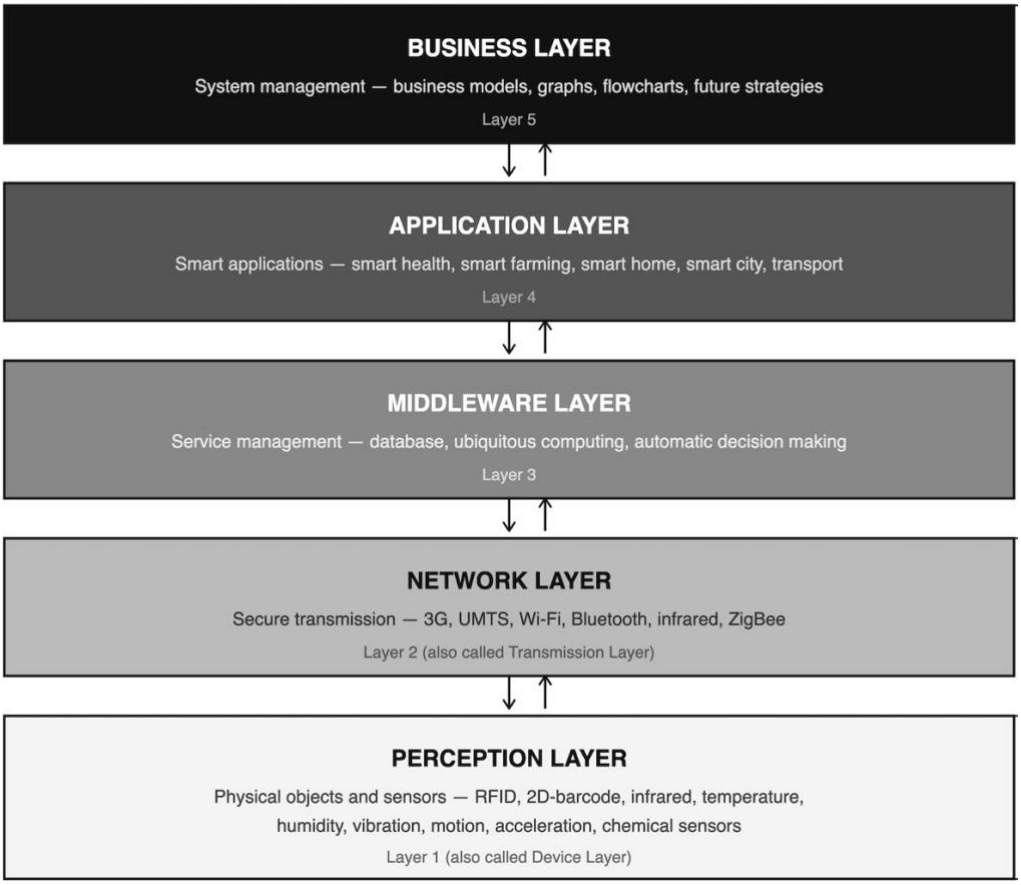


Figure 2.1 IoT three-layer architecture
Source: Adapted from Lashari et al. (2018)

2.4 Hardware Used in the Sensing Tier

The instrumented cage in this study draws on six commodity components. The choices were driven by availability in the local electronics market, total parts cost under one million rupiah, and prior reports of stable behaviour under tropical humidity.

a. ESP32

The ESP32 is a low-power microcontroller from Espressif Systems that integrates Wi-Fi and Bluetooth on a single chip. It is the workhorse of the cage installation reading every sensor at the half-hourly interval and posting the resulting payload to the backend over the farm's Wi-Fi. The dual-core 32-bit processor running at 240 MHz provides more headroom than this application requires, but the price point and pin count make it the practical default for projects of this scale.

b. DHT22

The DHT22 reports air temperature and relative humidity through a single-wire digital interface, with a temperature range of −40 to 80 °C at ±0.5 °C and humidity coverage across the full 0–100 percent band at ±2 percent. The sensor is adequate for cage monitoring in absolute terms, although extended exposure to saturated humidity which pancaroba routinely produces sits at the edge of its calibrated range and is a known limitation discussed later in the report.

c. MQ-137

The MQ-137 is a metal-oxide sensor that returns an analog voltage proportional to ambient ammonia concentration. Its sensitivity covers the 5–500 ppm window, which brackets both the welfare ceiling and the acute-injury threshold relevant to layer housing. The sensor requires a warm-up period and is sensitive to temperature drift, both of which are handled in firmware rather than through external compensation circuitry.

d. BH1750

The BH1750 is a digital ambient-light sensor communicating over I²C, with a measurement range from 1 to 65,535 lux. It serves two roles in this installation: contributing the photometric input to the stress index and acting as the trigger for the automatic lighting controller during overcast afternoons.

e. DS3231

The DS3231 is a temperature-compensated real-time clock module on the I²C bus. Its function is straightforward but important every reading is timestamped at source, independently of whether the Wi-Fi link to the backend is live at the moment of measurement. Without this, brief network outages would produce ambiguous chronology in the dataset.

f. 2-Channel Relay

The relay module contains two electromechanical switches rated for 250 V AC at 10 A, driven from microcontroller GPIO. In this build only one channel is in use, switching the cage lamp circuit under firmware control. The second channel was wired but reserved for a ventilation fan that the partner farm has not yet installed.

Table 2.1 Hardware specifications

Component	Function	Key Parameters	Interface
ESP32	Microcontroller	Dual-core 32-bit, 240 MHz, Wi-Fi, Bluetooth	Wi-Fi / BLE
DHT22	Temperature, humidity	−40 to 80 °C; 0–100 % RH	Single-wire
MQ-137	Ammonia	5–500 ppm	Analog
BH1750	Illuminance	1–65,535 lux	I²C
DS3231	RTC	±2 ppm drift	I²C
Two-channel relay	Lamp actuator	250 V AC, 10 A per channel	Digital GPIO

2.5 Software Stack

The processing and presentation tiers run on commodity software, deployed to a single virtual server in a Jakarta-region data centre.

a. Java Spring Boot

The backend is written in Java with Spring Boot, chosen partly because the author was already familiar with the framework from earlier coursework and partly for its built-in support for the operations this project needs dependency injection, REST endpoints, scheduled tasks for the nightly aggregation job, and Spring Data JPA for the database layer. The framework does not contribute much novelty, but it removes friction.

b. PostgreSQL

PostgreSQL stores the raw sensor records and the derived daily summaries. The relational model fits the data shape, the advanced aggregation functions handle the nightly rollup cleanly, and the JDBC integration with Spring Data is mature. There was no specific requirement for any feature outside the standard SQL surface.

c. React Native

The mobile client is built with React Native through the Expo toolchain. The framework allows a single JavaScript codebase to ship to both Android and iOS (React Native Community, 2024). The five-tab layout Dashboard, Cage, Community, Medical, Account was decided during the requirements phase and is described more fully in Chapter III.

d. Virtual Private Server (VPS)

A virtual private server runs the Spring Boot service and the PostgreSQL instance on the same host. The arrangement is not what a production deployment would look like, but it is sufficient for a thirty-day pilot with one cage, and it avoids exposing the database to the public internet.

2.6 *Large Language Model* (LLM) and Claude Sonnet

A Large Language Model is a neural network trained on large text corpora to produce contextual natural-language output. Their relevance to monitoring systems in agriculture comes from the interpretation side of the problem turning numerical readings into farmer-readable narrative. Recent work has begun to demonstrate this in pest management and crop advisory (Yuan et al., 2024), and Menezes et al. (2025) survey the broader trajectory toward livestock.

Claude Sonnet is one such model, developed by Anthropic and accessed through the Anthropic API. The API accepts a prompt and returns text that can be constrained in format through a system message. In this study, Claude Sonnet receives the daily aggregated summary from the backend and returns a structured JSON object containing cage status, condition narrative, observed issues, recommended actions, and a tentative disease prediction. Raw sensor data is never sent to the model; what reaches it has already been compressed by the backend into approximately fifteen representative values per day.

2.7 Composite Indices and the *Temperature-Humidity Index* (THI)

Combining several environmental measurements into a single normalised number has a long history in occupational health (Moran et al., 2001) and in livestock thermal-stress research (Mader et al., 2010). The most widely used example in poultry is the Temperature-Humidity Index, originally proposed by Thom in 1959 and later validated against laying performance in Japanese quail by El-Tarabany (2016) and by de Oliveira Castro et al. (2023).

The THI is calculated using equation (2.1) below.

$$\mathrm{THI} = (0.8 \times \mathrm{T}) + [100\mathrm{RH} \times (\mathrm{T} - 14.4)] + 46.4 \qquad (2.1)$$

where T is the daily mean temperature in degrees Celsius and RH is the daily mean relative humidity expressed as a percentage. Reported anchor values for laying poultry place the comfort boundary at 68 and the danger threshold at 82 (Lara & Rostagno, 2013).

2.8 The Environmental Stress Index

The Environmental Stress Index extends the THI framework rather than replacing it. The index combines thermal load, ammonia, and illuminance into a single scalar bounded between zero and one, with the explicit purpose of giving the farmer a number that can be acted on without further mental arithmetic.

Each component is converted into a clipped linear sub-score against literature-derived thresholds. The THI sub-score uses equation (2.2):

$$\mathrm{STHI} = \mathrm{clip}(82 - 68\mathrm{THI} - 68, 0, 1) \quad (2.2)$$

The ammonia sub-score follows the same form, anchored at the welfare limit of 10 ppm and the acute-injury threshold of 25 ppm equation (2.3):

$$\mathrm{SNH3} = \mathrm{clip}(25 - 10\mathrm{NH3} - 10, 0, 1) \quad (2.3)$$

The light sub-score captures deficient illumination during active hours, taking the fifteen-lux floor as the cut-off. Naeem et al. (2018) tested Japanese quail at multiple light intensities and found that 15 lux during the laying period produced the highest egg output among all treatment groups equation (2.4):

$$\mathrm{SLight} = \mathrm{clip}(1515 - \mathrm{Lux}, 0, 1) \quad (2.4)$$

The composite is then a weighted sum, given by equation (2.5):

$$\begin{aligned} \mathrm{ESI} &= 0.45 \times \mathrm{STHI} + 0.35 \times \mathrm{SNH3} + 0.20 \times \mathrm{SLightESI} \\ &= 0.45 \backslash\mathrm{times\ S_\{THI\}} + 0.35 \backslash\mathrm{times\ S_\{NH3\}} \\ &+ 0.20 \backslash\mathrm{times\ S_\{Light\}ESI} \\ &= 0.45 \times \mathrm{STHI} + 0.35 \times \mathrm{SNH3} + 0.20 \times \mathrm{SLight} \end{aligned} \quad (2.5)$$

The weights are not optimised from data; they are taken from the relative physiological emphasis the layer-poultry literature places on each stressor. Thermal-humidity loading carries the highest weight at 0.45 because it folds two parameters together and acts most directly on the reproductive axis. Ammonia carries 0.35 because of the respiratory consequences of exposure above the welfare ceiling. Illuminance carries 0.20, lower because the lighting infrastructure already exists in most operating cages and tends to fail less often than the other two pathways. Three operational bands convert the continuous index into a farmer-facing label (Table 2.2), and Table 2.3 summarises the threshold sources.

Table 2.2 ESI bands and the corresponding action recommendation

ESI range	Label	Recommended action
0.00 – 0.30	Good	No intervention; continue routine practice
0.31 – 0.60	Needs Attention	Increase ventilation and observation frequency
0.61 – 1.00	Critical	Cleaning and corrective action without delay

Table 2.3 Component thresholds and their literature sources

Component	Safe Zone	Danger threshold	Source
THI	≤ 68	≥ 82	Lara & Rostagno (2013)
Ammonia	≤ 10 ppm	≥ 25 ppm	Naseem & King (2018)
Iluminance	≥ 15 lux	< 15 lux	Naeem et al. (2018)

2.9 Methods Used to Evaluate the System

Two evaluation methods are applied in this study, each addressing a different question.

a. Pearson correlation

The Pearson product-moment coefficient measures the strength and direction of a linear relationship between two continuous variables. Its value ranges from −1 to +1; the sign indicates direction and the magnitude indicates strength. The two-tailed test is used at $\alpha = 0.05$ to assess whether the daily ESI is statistically associated with the daily egg count over the thirty-day observation window.

b. Black-box testing

Black-box testing evaluates a system's external behaviour without reference to its internal implementation (Sommerville, 2016). The procedure is to enumerate the functions of the system, define expected outputs for representative inputs, and check that the system produces those outputs. For this project the relevant functions span the entire pipeline, from sensor packet reception, through nightly aggregation and ESI computation, to the LLM call and the rendering of the result on the mobile dashboard.

2.10 Where the Problem Sits in the Existing Work

Several research groups have built IoT systems for poultry housing in the last five years, but none of the ones the present author was able to locate combine the three properties that this study requires together: a composite environmental score, an LLM-based interpretation layer, and a field-validated link to production data.

Lashari et al. (2018) and Gunawan et al. (2021) both build cage-monitoring instruments and demonstrate that the telemetry works. Neither paper aggregates the readings into a single farmer-facing metric, and neither links the readings to subsequent egg yield. The result is a

working dashboard but a dashboard that still asks the farmer to integrate several variables mentally.

The composite index approach itself is well established in livestock science. The Temperature-Humidity Index has been used for decades to condense thermal load into a single number that correlates with animal performance (Lara & Rostagno, 2013). El-Tarabany (2016) validated THI specifically for Japanese quail, showing that birds exposed to THI values above 76 had significantly lower fertility and egg quality than birds kept below THI 70. The principle combining correlated stressors into one scalar that a non-specialist can read at a glance transfers directly to cage monitoring. What the existing THI work does not include is ammonia or light intensity, both of which are relevant to laying quail in enclosed housing and which vary independently of the thermal load during pancaroba.

On the AI side, the integration of language models with IoT data is recent but moving. PestGPT (Yuan et al., 2024) demonstrates that an LLM can take aggregated sensor input and produce useful natural-language guidance. Menezes et al. (2025) survey the trajectory and identify livestock applications as a clear gap. None of these papers address layer quail, and none reports field validation against production data.

The remaining gap, as the present study reads it, has three components: there is no composite index built and validated for quail; there is no public report linking such an index to laying performance under real conditions; and there is no working pipeline that combines such an index with LLM-generated recommendations on the farmer's phone.

Beyond the engineering act of connecting IoT sensors to a Large Language Model a step several prior systems have already demonstrated at the telemetry level (Lashari et al., 2018; Gunawan et al., 2021) the most significant scientific contribution of this research lies in three original elements that together move the work from a monitoring prototype to a validated decision-support system.

First, a domain-specific composite index. Existing composite indices such as the Temperature-Humidity Index (Lara & Rostagno, 2013; El-Tarabany, 2016) capture thermal load alone and do not incorporate ammonia or light intensity, both of which vary independently of temperature during pancaroba. The Environmental Stress Index (ESI) proposed in this study is, to the author's knowledge, the first composite index that combines these three stressor dimensions with literature-derived weighting specifically for laying quail.

Second, an empirical, field-measured link between the index and actual production output. Prior IoT poultry-monitoring systems (Lashari et al., 2018; Gunawan et al., 2021)

demonstrate that telemetry works but stop short of linking readings to yield. This study correlates the ESI to daily egg production across 30 consecutive field days (Pearson $r = -0.361$, $p = 0.0498$), which is what allows the system to function as a decision-support tool rather than a passive dashboard.

Third, two previously undocumented behavioural findings about an LLM operating inside an IoT diagnostic pipeline: (a) an asymmetric precaution bias, in which every disagreement between the LLM's classification and the ESI erred toward over-caution rather than under-caution (Section 4.5.2), and (b) a temporal-aggregation blind spot, in which daily-averaged input caused the LLM to misclassify short, intense weather events as sustained environmental failure (Section 4.7). Neither pattern has previously been reported in the IoT-LLM agricultural literature (Yuan et al., 2024; Menezes et al., 2025), which makes them an original methodological contribution independent of the quail application.

Table 2.4 Comparison of this study with existing related work

Study	Target	Parameters	Composite Index	LLM Diagnostics	Field-Validated Link to Production
Lashari et al. (2018)	Poultry (general)	Temperature, humidity	None	No	No telemetry demo only
Gunawan et al. (2021)	Poultry (general)	Temperature, humidity, feed	None	No	No telemetry demo only
Lara & Rostagno (2013)	Ruminant / poultry (general)	Temperature, humidity	Yes (THI, thermal only)	No	No theoretical index
El-Tarabany (2016)	Japanese quail	Temperature, humidity	Yes (THI, thermal only)	No	Partial links THI to fertility/egg quality in lab; no live IoT
Yuan et al. (2024) – PestGPT	Crop pest (non-livestock)	Aggregated sensor inputs	No	Yes (GPT-based)	No not a livestock/production outcome
Menezes et al. (2025)	Livestock (survey)	N/A (review paper)	N/A	Reviewed, not deployed	No identifies gap, does not fill it
This study (2026)	Laying quail (Coturnix japonica)	Temp, humidity, NH_3, light	Yes ESI (3 stressors, literature-weighted)	Yes Claude Sonnet	Yes Pearson $r = -0.361$, $p = 0.0498$, $n = 30$ field days

2.11 Choosing Between the Available Approaches

For the interpretation layer, three approaches were genuinely considered before settling on the one used in this build. Table 2.5 summarises the comparison.

Table 2.5 Comparison of interpretation approaches

Approach	Strength	Weakness	Training data	Output form
Threshold alerting	Simple, transparent, cheap (Gunawan et al., 2021)	No interaction handling between parameters	None	Per-parameter alarm
Classical machine learning	Learns patterns from history (Menezes et al., 2025)	Needs large labelled datasets; opaque to farmers	Thousands of samples	Class label or numeric prediction
Large language model	Produces actionable natural-language advice (Yuan et al., 2024)	Risk of hallucination; depends on external API	Low (zero- or few-shot)	Structured narrative

2.12 Conceptual Framework

The conceptual framework that ties the chapter together is shown in Figure 2.9. The diagram reads from top to bottom problem, gap, proposed system, the three variable groups that make up that system, the evaluation procedures, and the intended deliverable.

The starting problem is the one introduced in Chapter I: smallholder quail farmers in Central Java face cumulative environmental stress during pancaroba and have no instrumentation to anticipate the resulting drop in production. Section 2.10 identified the gap that follows from the existing literature the absence of a composite quail-specific index that is linked both to LLM interpretation and to field-measured productivity.

The proposed system addresses that gap with a three-stage pipeline. The input variables are the four environmental readings produced by the cage sensors. The process variables are the daily aggregation in Spring Boot and PostgreSQL, the ESI computation, and the call to Claude Sonnet through the Anthropic API. The output variables are the ESI classification, the LLM diagnosis, the recommended action, and the rendering of all three on the React Native dashboard. Evaluation rests on the Pearson test for the index-yield relationship, on black-box testing for the system as a whole, and on expert review of representative LLM outputs. The deliverable is an empirically validated decision-support system narrow enough to fit one cage and one season explicitly a pilot, with multi-site validation identified as the next step.

CHAPTER III
RESEARCH METHODOLOGY

This chapter describes how the IoT-ESI-LLM system was built and tested. The work followed a Research and Development approach organised around the ASET model, which breaks a project into four phases: Analysis, System design, Execution, and Testing. Each phase is covered in one main section of this chapter, and the sub-sections within each phase are ordered roughly in the sequence the work was actually done.

3.1 Analysis phase

The analysis phase ran from early March to mid-April 2026. Its purpose was to decide where the system would be deployed, what it needed to measure, and how the field data would be collected once the hardware was running. Three activities made up this phase: selecting and characterising the research site, interviewing the farmer and an animal husbandry specialist to pin down requirements, and designing the data collection protocol.

3.1.1 Research site and observation period

The field deployment took place at a commercial quail operation in Siwatu village, Wonotunggal district, Batang regency, Central Java. The farm is owned and operated by a single household. At the start of data collection the active flock numbered approximately 1,800 laying hens housed in one production cage. The location of the research site is shown in Figure 3.1.



Figure 3.1 Research site, Siwatu village

The site was chosen for three practical reasons. First, the farmer had already expressed dissatisfaction with how difficult it was to keep track of cage conditions during the rainy-to-dry changeover and was willing to let an outsider mount electronics above the birds for a month. Second, the farm sits in the coastal hinterland of north Central Java, an area where the pancaroba signature is strong: hot mornings, sudden afternoon storms, humidity that stays near saturation for hours after the rain stops. Third, the cage is served by a domestic Wi-Fi router whose signal reaches the laying area, which meant the ESP32 could transmit to the internet without any extra networking equipment.

The observation window was set at thirty consecutive days, from 26 April to 25 May 2026. This window was chosen deliberately to coincide with the tail end of the wet season and the beginning of the dry season. Weather records for the Batang coast show that late April through late May is the period when afternoon rainfall events are most unpredictable and when temperature swings within a single day are at their widest. By placing the trial inside this window, the system was tested under what the farmer described as the worst month of the year for egg output.

A thirty-day window is short. It is long enough to capture the environmental variability that the index is designed to summarise, and long enough to run a Pearson test with $n = 30$ data points, but it is not long enough to make claims about seasonal or annual patterns. The study is framed as a pilot for this reason: the goal is to test whether the system works and whether the ESI-production association exists, not to prove that it holds in all conditions.

3.1.2 Requirements analysis

Before any hardware was purchased, two rounds of interviews were conducted to establish what the system needed to do.

The first interview was with the participating farmer, held at the farm on 8 March 2026. The conversation lasted about forty-five minutes and was semi-structured. The farmer was asked to describe his daily routine for checking cage conditions, what signs he looks for when something is wrong, how quickly he usually notices a drop in egg count, and what kind of information would be useful if it appeared on his phone. His answers confirmed several assumptions from the literature but also added detail that the literature does not cover. He checks the cage twice a day, morning and late afternoon. He judges temperature by how the

birds are behaving (panting, huddling) rather than by reading a thermometer, because he does not own one. He notices ammonia by smell but cannot quantify it. He has no way to measure light intensity. When eggs drop, he finds out the next morning, counts the difference, and then tries to figure out what went wrong the day before. He said he would trust a phone notification if it told him the cage was in bad shape, but only if the number behind it made sense to him.

The second interview was with an animal husbandry specialist from the Faculty of Animal Science at Universitas Gadjah Mada, conducted on 15 March 2026 by video call. The specialist was asked to validate the choice of environmental parameters and the threshold values that would later be used in the ESI formula. He confirmed that temperature, humidity, ammonia, and light intensity are the four parameters with the strongest documented links to laying performance in Japanese quail. He recommended the THI comfort boundary at 68 and the danger boundary at 82, the ammonia safe limit at 10 ppm and injury threshold at 25 ppm, and the minimum light floor at 15 lux during active hours. He also suggested that a dark period between 21:00 and 03:00 would be appropriate for the breed and season. These thresholds are consistent with the published values reviewed in Chapter II, but hearing them confirmed by a practitioner working with Indonesian farms gave additional confidence that the values would be reasonable under local conditions.

The findings from both interviews are summarised in Table 3.1.

Table 3.1 Summary of interview findings from the requirements analysis phase

Source	Date	Key findings
Farmer (Siwatu)	8 Mar 2026	No instruments in cage; judges conditions by bird behaviour and smell; discovers production drops one day late; wants a single number on his phone that tells him whether to act
Animal husbandry specialist	15 Mar 2026	Confirmed four parameters (temperature, humidity, ammonia, light); validated THI thresholds (68/82), ammonia thresholds (10/25 ppm), light floor (15 lux); recommended dark window 21:00–03:00

Based on these interviews, four environmental parameters were selected for monitoring. The selection and the reason behind each parameter are listed in Table 3.2.

Table 3.2 Selected parameters and their physiological justification

Parameter	Sensor	Why it matters for laying quail
Air temperature	DHT22	Directly affects feed intake and reproductive physiology; sustained heat reduces egg production by 6–21% depending on severity and duration (Alagawany et al., 2017; El-Tarabany, 2016)
Relative humidity	DHT22	Amplifies thermal stress by blocking evaporative cooling through the respiratory tract; feeds into the THI formula
Ammonia concentration	MQ-137	Damages respiratory epithelium above 25 ppm; reduces feed intake above 10 ppm (Naseem & King, 2018)
Light intensity	BH1750	Governs photoperiod; environmental factors outside the thermoneutral zone disrupt reproductive physiology and reduce egg quality (El-Tarabany, 2016; Naeem et al., 2018).

3.1.3 Data Collection Design

The sampling interval was set at thirty minutes. The reasoning was practical rather than theoretical: during pancaroba, conditions inside the cage can shift from sunny and hot to rain-cooled and saturated in under two hours. A thirty-minute interval captures at least three or four readings across a typical afternoon storm event, which is enough to see the shape of the change. A shorter interval would have generated more data but would also have increased the load on the backend and shortened the battery life of any backup power source. A longer interval risked missing the rapid transitions that make pancaroba monitoring difficult in the first place.

At thirty-minute intervals, each sensor produces 48 readings per day. With four parameters monitored continuously over thirty days, the expected total is approximately 5,760 raw records. In practice the actual count will be slightly lower because of occasional Wi-Fi dropouts or sensor read failures, but the thirty-minute cadence was maintained by the firmware regardless of network availability: readings that failed to transmit were buffered locally on the ESP32 and posted on the next successful connection.

Egg production was recorded manually. Each morning the farmer counted the eggs collected from the cage and entered the number into the mobile application through the Kandang (Cage) tab. This produced one egg-count record per day, for a total of thirty daily records over the observation window.

Qualitative data was collected through two additional interviews: one with the farmer after the system had been running for two weeks (mid-deployment check), and one with both the farmer and the specialist after the full thirty days were complete (post-deployment review). The purpose of the mid-deployment check was to catch usability problems early enough to fix them. The purpose of the post-deployment review was to gather the farmer's assessment of

whether the system changed how he managed the cage, and to have the specialist evaluate a sample of the LLM diagnostic outputs for scientific accuracy.

Table 3.3 Data collection summary

Data type	Source	Frequency	Expected total
Temperature (°C)	DHT22 via ESP32	Every 30 min	~1,440 records
Relative humidity (%)	DHT22 via ESP32	Every 30 min	~1,440 records
Ammonia (ppm)	MQ-137 via ESP32	Every 30 min	~1,440 records
Light intensity (lux)	BH1750 via ESP32	Every 30 min	~1,440 records
Egg count	Farmer (manual entry)	Once daily (afternoon)	30 records
Qualitative interview	Farmer, specialist	Mid-deployment, post-deployment	2 sessions

3.2 System Design Phase

The system design phase started immediately after the requirements analysis closed in mid-April 2026. By that point the parameter list, the threshold values, and the rough idea of what the farmer wanted on his phone were all settled. The job of this phase was to turn those decisions into something that could be built and deployed in time for the 26 April observation window.

The phase is broken into eight activities. The order roughly tracks how the work was done. However, in practice there was overlap the database schema and the firmware logic were drafted in parallel, and the mobile app wireframes were already being sketched. At the same time, the API endpoints were being defined.

3.2.1 System Architecture Design

The system uses a three-layer architecture. Each layer runs on different hardware and talks to the next layer through a defined protocol. The architecture is shown in Figure 3.2.

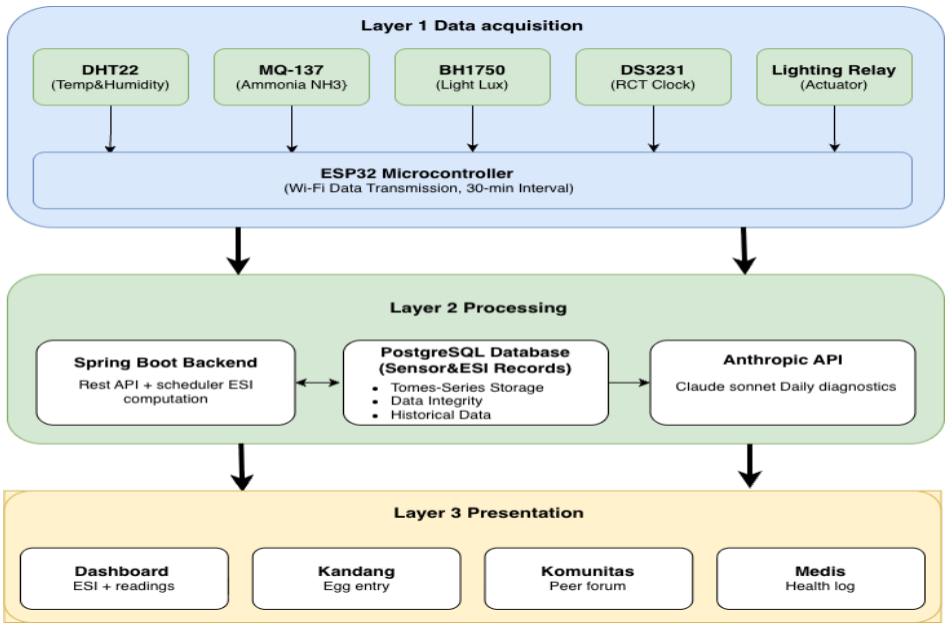


Figure 3.2 IoT-ESI-LLM system architecture

The data acquisition layer sits inside the cage. It consists of the ESP32 microcontroller, the four environmental sensors, the DS3231 RTC, and the relay module. This layer reads sensor values every thirty minutes, timestamps them, packages them into JSON, and posts them over Wi-Fi to the backend.

The processing layer runs on a virtual private server hosted in Jakarta. It receives the sensor packets through a REST API endpoint, writes them to a PostgreSQL database, runs the nightly aggregation job at 22:00, computes the daily ESI, and calls the Anthropic API to obtain the Claude Sonnet diagnostic. The processing layer also exposes a second set of API endpoints for the mobile application to read.

The presentation layer is the React Native mobile application installed on the farmer's Android phone. It pulls the latest sensor readings, the daily ESI value, and the LLM diagnostic from the backend, and presents them through five tabs. It also pushes notifications to the phone when the daily ESI crosses into the Critical band.

Communication between layers uses HTTP. The ESP32 posts to the backend over standard HTTPS. The mobile application reads from the backend over the same protocol. The backend calls the Anthropic API over HTTPS as a client. No layer talks to another layer except through these defined endpoints.

3.2.2 Hardware Wiring Design

The wiring schematic for the sensor unit is shown in Figure 3.3. The ESP32 sits at the centre and connects to every other component.

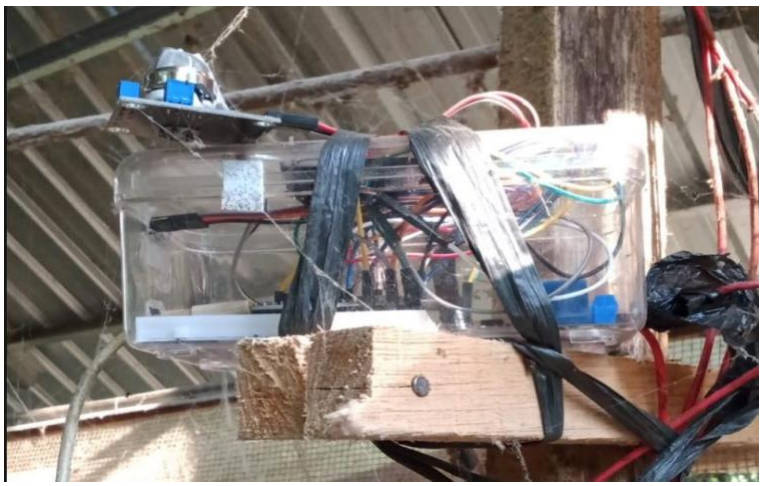


Figure 3.3 Sensor unit wiring

The pin mapping was decided based on which ESP32 pins are safe for which signal types. The DHT22 uses a single-wire digital protocol, so it was assigned to GPIO 4. The MQ-137 outputs an analog voltage, so it was assigned to GPIO 34, one of the ADC1 pins. The BH1750 and DS3231 both use I²C, so they share GPIO 21 (SDA) and GPIO 22 (SCL) I²C is a shared bus, so this is fine. The relay was assigned to GPIO 26 for the first channel and GPIO 27 for the second, though only the first channel is in use. The full mapping is listed in Table 3.4

Table 3.4 GPIO pin mapping on the ESP32

Components	Signal	ESP32 pin	Notes
DHT22	Digital single-wire	GPIO 4	10 kΩ pull-up to 3.3 V
MQ-137	Analog voltage	GPIO 34	ADC1, input only
BH1750	I²C SDA	GPIO 21	Shared with DS3231
BH1750	I²C SCL	GPIO 22	Shared with DS3231
DS3231	I²C SDA	GPIO 21	Shared with BH1750
DS3231	I²C SCL	GPIO 22	Shared with BH1750
Relay Channel 1	Digital out	GPIO 26	Drives lamp circuit
Relay Channel 2	Digital out	GPIO 27	Reserved (unused)

Power for the ESP32 and the sensors comes from a 5 V USB adapter plugged into the farm's mains supply. The DHT22, BH1750, and DS3231 all draw their 3.3 V from the ESP32's onboard regulator. The MQ-137 requires 5 V for its heater element and is fed directly from the USB rail through a separate line. The relay module also runs on 5 V and switches the 220 V AC lamp circuit through a mechanical contact the high-voltage side is fully isolated from the microcontroller.

3.2.3 Firmware Logic Design

The ESP32 firmware was written in Arduino C++. The main loop runs every thirty minutes through a combination of delay() and the DS3231 timer. The overall flow is shown in Figure 3.4.

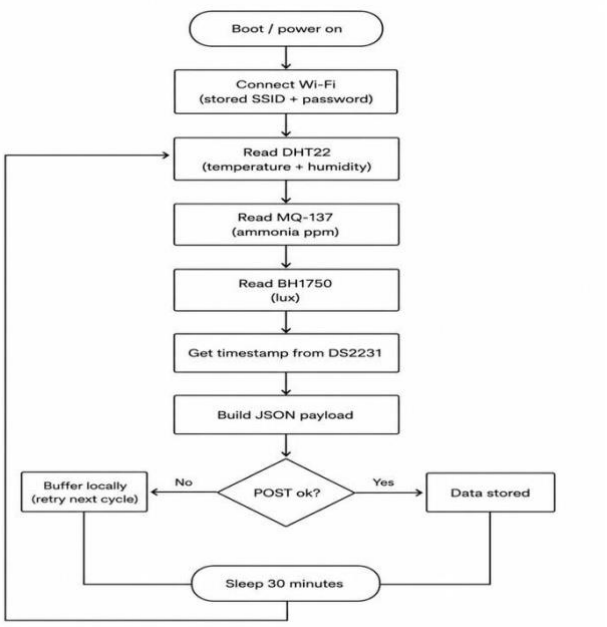


Figure 3.4 Firmware main loop flowchart

On boot, the firmware connects to the farm's Wi-Fi using stored credentials. Once connected, it enters the main loop. Each iteration reads all four sensors in sequence, requests a timestamp from the DS3231, and builds a JSON object with the four readings and the timestamp. The JSON object is posted to the backend through an HTTP POST to /api/sensor-data. If the post succeeds, the firmware sleeps for thirty minutes. If it fails, the JSON object is appended to a local buffer in flash memory and the firmware retries on the next cycle.

The lighting control logic runs separately and is checked every minute, not every thirty. The logic has two rules. First, between 21:00 and 03:00 local time, the lamp is always off regardless of light level this is the obligatory dark window for melatonin synthesis. Second, outside the dark window, the lamp turns on if the BH1750 reading is below 15 lux and turns off if it rises above 25 lux. The 10-lux gap between on and off thresholds prevents the relay from rapidly toggling when light is hovering near 15 lux. The control flow is shown in Figure 3.5.

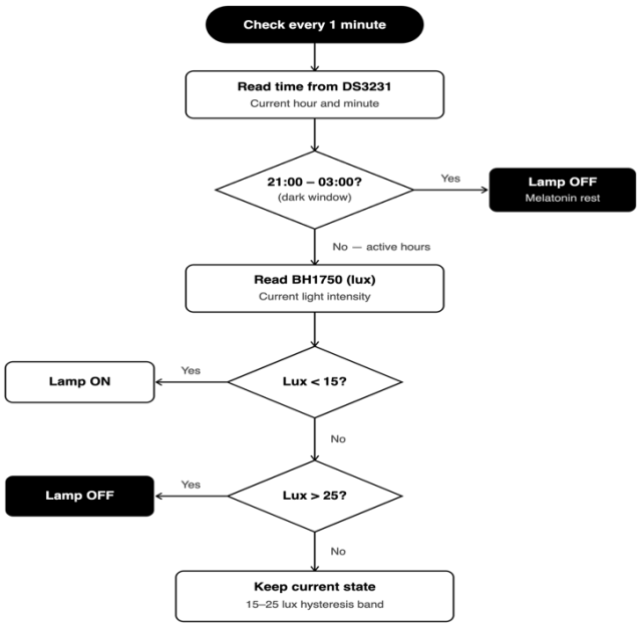


Figure 3.5 Lighting control logic flowchart

3.2.4 Database Design

The system's PostgreSQL database contains fourteen tables organised around a single central entity: the users table. All farm management records belong to a user through a user_id foreign key. The two IoT-specific tables sensor_data and llm_diagnoses are currently standalone, identified by a device string rather than a relational foreign key. The full entity-relationship diagram from the production database is shown in Figure 3.6.

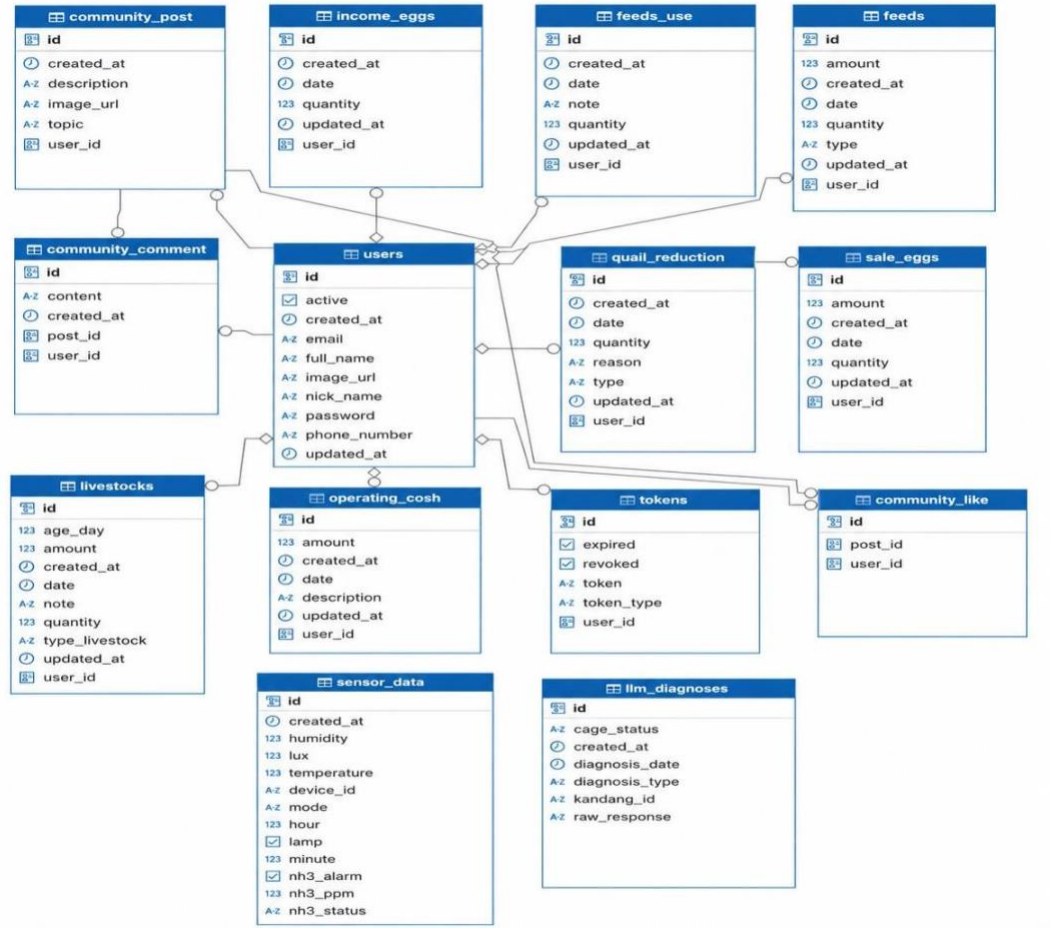


Figure 3.6 Production database ERD

The ERD was verified against the Kotlin Spring Data JPA model files before the observation period began. All fourteen table names in the diagram match the @Table(name = "...") annotations in the corresponding .kt files exactly. Table 3.5 describes the purpose and key columns of each table.

Table 3.5 Database table descriptions

Table	Purpose	Key
users	Farmer account data	id UUID PK, full_name, email, phone_number, password, active, province, city, reset_password_token, created_at
tokens	JWT session and revocation	id UUID PK, token, token_type, revoked, expired, user_id FK
livestocks	Quail flock batch records	id UUID PK, age_day, quantity, amount, date, type_livestock, user_id FK

Table	Purpose	Key
feeds	Feed stock purchases	id UUID PK, type, quantity, amount, date, user_id FK
feeds_use	Daily feed consumption	id UUID PK, quantity, note, date, user_id FK
income_eggs	Daily egg harvest count	id UUID PK, quantity, date, user_id FK
sale_eggs	Egg sale transactions	id UUID PK, quantity, amount, date, user_id FK
quail_reduction	Flock loss by death or sale	id UUID PK, quantity, reason, type, date, user_id FK
operating_cosh	Raw ESP32 sensor readings	id UUID PK, device_id, temperature, humidity, lux, lamp, nh3_ppm, nh3_status, nh3_alarm, hour, minute, created_at
llm_diagnoses	AI cage diagnosis results	id UUID PK, kandang_id, diagnosis_date, diagnosis_type, cage_status, raw_response TEXT, created_at
community_post	Farmer forum posts	id UUID PK, topic, description, image_url, user_id FK
community_comment	Comments on forum posts	id UUID PK, content, post_id FK, user_id FK
community_like	Post like junction table	id UUID PK, post_id FK, user_id FK unique(user_id, post_id)

The decision to keep sensor_data and llm_diagnoses as standalone tables without a foreign key to users was made for two practical reasons. First, the ESP32 device posts data before any user is logged in on the mobile side, so there is no session token available at write time. Second, the cage identifier (kandang_id) is a configuration string embedded in the ESP32 firmware, which allows the backend to match sensor data to a specific cage without requiring the ESP32 to authenticate.

3.2.5 REST API Endpoint Design

The backend exposes 56 endpoints across two base paths on the domain ternakpoyo.online. Farm management endpoints sit under /api/v1 and require a JWT Bearer token. IoT-specific endpoints sit under /api and use a static device credential for the ESP32 path.

a. Authentication and registration

This section covers user registration, login, and session management, supporting phone OTP, email/password, and Google OAuth flows.

Table 3.6 Authentication and registration

Method	**Path**	**Used by**	**Purpose**
POST	/api/v1/registration/phoneNumber	Mobile	Register via phone number.
POST	/api/v1/registration/password	Mobile	Register via email + password.
POST	/api/v1/login/phoneNumber	Mobile	Login with phone OTP.
POST	/api/v1/login/email	Mobile	Login with email and password.
POST	/api/v1/login/google	Mobile	Login with Google OAuth token.
POST	/api/v1/login/forgot-password	Mobile	Request password reset email.
POST	/api/v1/login/reset-password	Mobile	Submit new password with reset token.
POST	/api/v1/logout	Mobile	Invalidate JWT (logout)

b. User Profile

This section provides a single endpoint for retrieving the authenticated user's profile information.

Table 3.7 User Profile

Method	**Path**	**Used by**	**Purpose**
GET	/api/v1/me	Mobile	Get current user profile

c. Sensor data IOT

This section handles pushing and retrieving environmental sensor readings such as temperature, humidity, lux, and NH_3 collected from the ESP32 device.

Table 3.8 Sensor data IOT

Method	Path	Used by	Purpose
POST	/api/sensor	ESP32	Push sensor reading (temp, humidity, lux, NH_3)
GET	/api/sensor	Mobile	Retrieve paginated sensor records, filter by deviceId

d. Livestock Management

This section manages quail flock batches, supporting full lifecycle operations including creation, detail retrieval, updates, deletion, and total quantity summary.

Table 3.9 Livestock Management

Method	Path	Used by	Purpose
GET	/api/livestock	ESP32	List livestock batches (paginated)
POST	/api/livestock	Mobile	Add new flock batch
GET	/api/v1/livestock/{id}	Mobile	Get batch detail
PUT	/api/v1/livestock/{id}	Mobile	Update batch record
DELETE	/api/v1/livestock/{id}	Mobile	Delete batch record
GET	/api/v1/livestock/quantity	Mobile	Total flock quantity summary

e. Feed stock management

This section provides full CRUD operations for managing farm feed inventory entries.

Table 3.10 Feed stock management

Method	Path	Used by	Purpose
GET/POST/GET/PUT/DELETE	/api/v1/feed and /api/v1/feed/{id}	Mobile	Full CRUD on feed stock entries

f. Feed Usage

This section manages daily feed consumption records, allowing users to log and track how much feed is used each day.

Table 3.11 Feed Usage

Method	Path	Used by	Purpose
GET/POST/GET/PUT/DELETE	/api/v1/feeduse and /api/v1/feeduse/{id}	Mobile	Full CRUD on daily feed consumption records

g. Egg income

This section handles daily egg harvest records, enabling users to log, update, and review egg production data over time.

Table 3.12 Egg income

Method	Path	Used by	Purpose
GET/POST/GET/PUT/DELETE	/api/v1/incomeEgg and /api/v1/incomeEgg/{id}	Mobile	Full CRUD on daily egg harvest records

h. Egg sales

This section manages egg sale transactions, allowing users to record and track sales history for financial reporting.

Table 3.13 Egg sales

Method	Path	Used by	Purpose
GET/POST/GET/PUT/DELETE	/api/v1/saleEgg and /api/v1/saleEgg/{id}	Mobile	Full CRUD on egg sale transactions

i. Quail reduction

This section tracks flock reduction events such as deaths or culling, keeping the active population count accurate.

Table 3.14 Quail reduction

Method	Path	Used by	Purpose
GET/POST/GET/PUT/DELETE	/api/v1/quailreduction and /api/v1/quailreduction/{id}	Mobile	Full CRUD on flock reduction records

j. Operating cost

This section manages farm operational expense records, supporting cost tracking across various farm activities.

Table 3.15 Operating cost

Method	Path	Used by	Purpose

GET/POST/GET/PUT/DELETE	/api/v1/operatingCosh and /api/v1/operatingCosh/{id}	Mobile	Full CRUD on farm operating cost entries

k. Dashboard

This section provides a single summary endpoint that aggregates livestock, egg production, and feed data into one overview response.

Table 3.16 Dashboard

Method	Path	Used by	Purpose
GET	/api/v1/dashboard	Mobile	Summary of livestock, egg, and feed overview

l. AI diagnosis – LLM pipeline

This section exposes AI-powered health diagnosis endpoints, allowing users to retrieve the latest or historical diagnosis results and manually trigger new diagnoses for a specific cage.

Table 3.17 AI diagnosis LLM pipeline

Method	Path	Used by	Purpose
GET	/api/diagnosis/{kandangId}/latest	Mobile	Latest diagnosis result for a cage
GET	/api/diagnosis/{kandangId}/history	Mobile	Last 7 diagnosis results for a cage
POST	/api/diagnosis/{kandangId}/trigger	Mobile / Admin	Manually trigger diagnosis (optional ?date=YYYY-MM-DD)

m. Community forum

This section supports the in-app community forum, enabling users to browse posts, create content, and interact through likes and comments.

Table 3.18 Community forum

Method	Path	Used by	Purpose
GET	/api/v1/community/posts	Mobile	List community posts (paginated)
POST	/api/v1/community/posts	Mobile	Create post with optional image (multipart/form-data)
POST	/api/v1/community/like	Mobile	Like a post
POST	/api/v1/community/unlike	Mobile	Remove like from a post

POST	/api/v1/community/comment	Mobile	Add comment to a post

3.2.6 ESI Computation Design

The ESI computation runs once a day inside the backend, scheduled by Spring's @Scheduled annotation to fire at 22:00 local time. The full computation flow is shown in Figure 3.7.

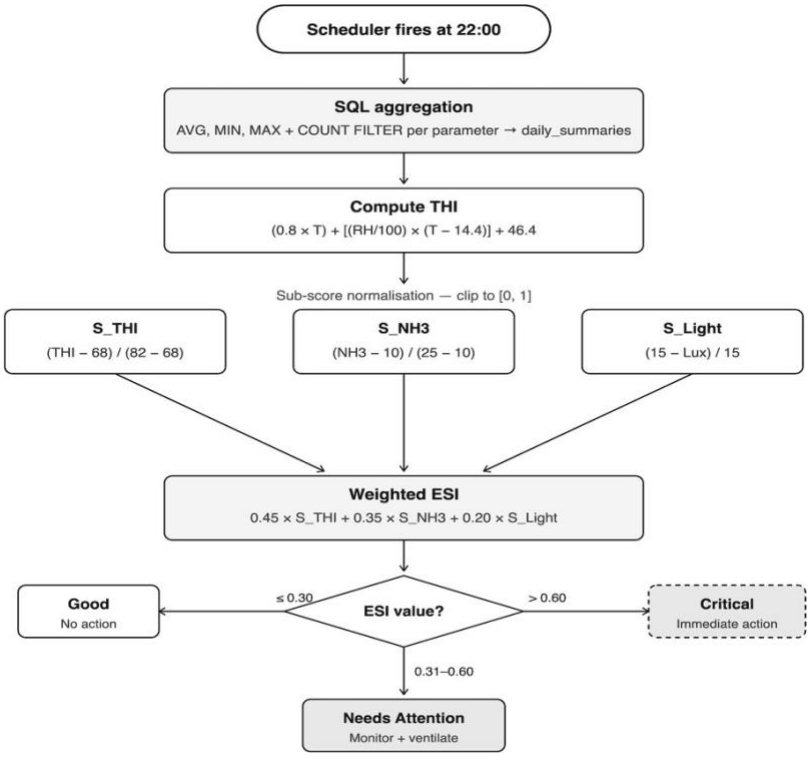


Figure 3.7 ESI computation flowchart

The first step is a SQL query that aggregates the day's raw records into a single row of summary values. The query computes the average, minimum, and maximum for each parameter, plus a count of dangerous events using a COUNT FILTER clause. Dangerous events are defined as readings that crossed a per-parameter threshold during the day temperature above 30 °C, humidity above 95%, ammonia above 25 ppm, or light below 15 lux during active hours. The query writes its result to daily_summaries.

The second step computes THI from the daily average temperature and humidity using equation (2.1) from Chapter II. The third step computes each sub-score using equations (2.2), (2.3), and (2.4). The fourth step combines the sub-scores into the composite ESI using equation

(2.5). The fifth step classifies the ESI into the Good, Needs Attention, or Critical band according to the thresholds in Table 2.2. All five steps run inside a single Spring service method, which writes the final result to esi_results.

3.2.7 LLM Diagnostic Pipeline Design

The LLM call runs immediately after the ESI computation completes. The full pipeline is shown in Figure 3.8.

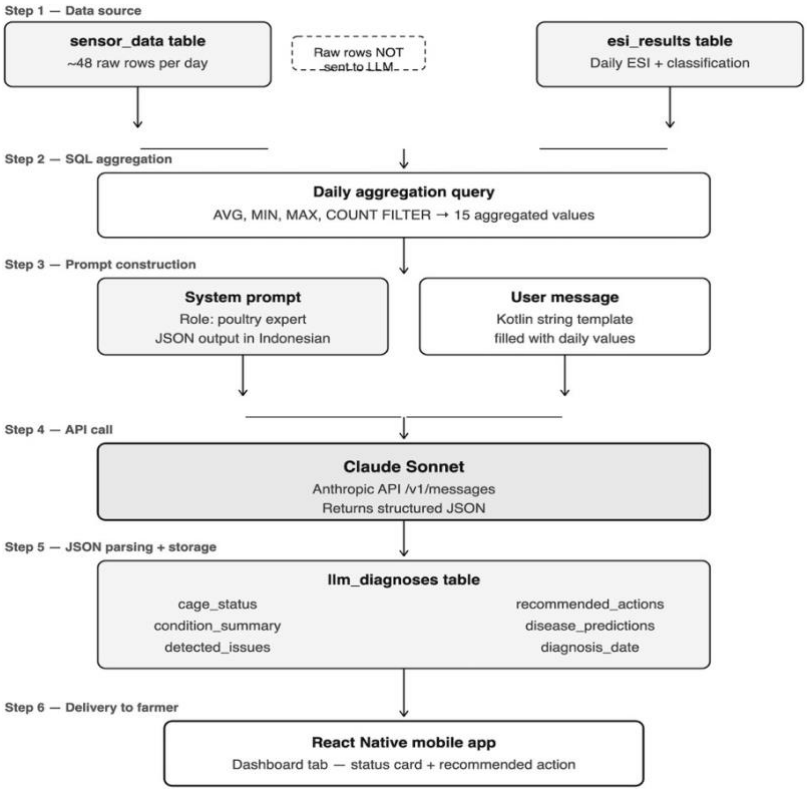


Figure 3.8 LLM diagnostic pipeline

Raw sensor data is never sent to Claude Sonnet. The pipeline instead sends a short textual summary built from the aggregated values in daily_summaries plus the ESI result. The summary is generated by a Kotlin string template that fills in around fifteen numerical fields the average, minimum, and maximum for each parameter, the count of dangerous events, the THI value, the ESI value, and the classification.

The Anthropic API call uses two parts. The system prompt assigns Claude Sonnet the role of a poultry husbandry expert advising a smallholder farmer in Central Java, instructs it to

respond in Indonesian, and specifies the exact JSON structure the response must follow. The user message contains the daily summary text.

The response is parsed back into the llm_diagnostics table, with each JSON field stored in its own column. The expected fields are listed in Table 3.19.

Table 3.19 Expected fields in the LLM JSON response

Field	Type	Purpose
cage_status	enum (Normal / Warning / Emergency / Critical)	Top-level status label
condition_summary	string (1–2 sentences)	Plain-language summary of the day
detected_issues	array of strings	Specific problems Claude Sonnet identified
recommended_actions	array of strings	What the farmer should do
disease_predictions	array of strings	Diseases the conditions could trigger

The reason raw data is withheld from Claude Sonnet is twofold. First, sending 192 raw readings (48 per parameter × 4 parameters) would consume a large amount of input tokens, raising both cost and latency. Second, the aggregated summary already contains the information the model needs to reason about the day. The data aggregation step is therefore performed by ordinary database queries and program logic, not by the LLM the LLM is used only for the diagnostic reasoning that follows.

3.2.8 Mobile Application UI Design

The mobile application has five tabs, decided during the requirements interview based on the farmer's stated needs and the additional features he requested. The five tabs and their purposes are listed in Table 3.20.

Table 3.20 Mobile application tabs and their content

Tab	Indonesian label	Content
Dashboard	Dashboard	Live readings, ESI gauge, lamp status, latest LLM diagnosis
Cage	Kandang	Flock count, daily egg entry, cage administration
Community	Komunitas	Discussion board for farmer peer support
Medical	Medis	Health records and disease history
Account	Akun	User settings, notification preferences

3.3 Execution Phase

The execution phase covered the physical installation of the hardware at Siwatu, deployment of the backend and mobile application, and the running of field data collection across the 30-day observation window. This phase was where every design decision from Section 3.2 met the reality of a working farm, and it ran from mid-April to 25 May 2026.

3.3.1 Hardware assembly and installation

The sensor unit was assembled in a workshop before being brought to the farm. The ESP32 microcontroller was mounted on a strip of perforated prototyping board and the four sensors were connected using color-coded jumper wire following the GPIO mapping in Table 3.4. The board was housed in a transparent polycarbonate enclosure roughly the size of a fist, chosen for its resistance to ammonia corrosion and the fact that it does not block the light path to the BH1750. The enclosure had two openings drilled at the base for the DHT22 and MQ-137 to sample air directly, and the BH1750 sat behind a small transparent window on the top face. The DS3231 and the relay module were mounted inside without external openings because neither needs air contact.

The completed unit was installed at Pak Suyud's farm on 25 April 2026, one day before the observation window officially began. It was mounted on a wooden post above the laying area at approximately head-height for the birds, positioned to sample cage air without being splashed by feed or water. Power came from a 5V USB adapter plugged into a household outlet through a short extension cable. The relay's high-voltage side was wired into the existing lamp circuit through a mechanical splice that the farmer's electrician verified.

The farm's Wi-Fi router sat in the main house approximately fifteen meters from the cage. Signal strength inside the cage was measured at −68 dBm using the ESP32's onboard scanner, which is well within the operating range for reliable HTTP transmission. The credentials were flashed into the firmware before the unit left the workshop, so the ESP32 connected automatically on first power-up.

3.3.2 Backend deployment

The Spring Boot backend was deployed to a VPS hosted in Jakarta running Ubuntu 20.04.5 LTS. The server had 1 vCPU (AMD EPYC 7713), 965 MB RAM, and 25 GB of SSD storage modest by commercial standards but adequate for the data volumes produced by a single-cage deployment. The JDK version in production was OpenJDK 17.0.15. The nightly aggregation job and LLM diagnostic pipeline both ran as scheduled Spring tasks and did not require manual restarts across the 30-day observation window. Table 3.21 lists the full environment specifications.

Table 3.21 Server environment specifications

Component	Specification
Operating system	Ubuntu 20.04.5 LTS (Focal Fossa)
CPU	1 vCPU AMD EPYC 7713 (KVM)
RAM	965 MB total / 703 MB used
Storage	25 GB SSD / 21 GB used
JDK	OpenJDK 17.0.15
Application framework	Spring Boot 3.2
Database	PostgreSQL 16
Domain	ternakpoyo.online

Two points about the production environment are worth stating plainly. Memory usage averaged around 73% during the observation period, and swap was fully consumed throughout, which indicates the server was running near its memory ceiling. This did not cause any observable failures during the pilot, but a deployment serving multiple cages simultaneously would need more RAM. Disk usage reached 90% by the end of the observation window, driven primarily by log files and the accumulated sensor records. Neither issue was critical for a 30-day single-cage pilot, but both are flagged as infrastructure limitations for the follow-up study.

3.3.3 Mobile application deployment

The React Native application was built using Expo Application Services (EAS Build). The output was an APK installer file that was transferred to Pak Suyud's Android phone via USB and installed with local installation permissions. iOS deployment was not attempted because the farmer's phone runs Android and because Apple's app distribution model would require a paid developer account and store review for a single-user pilot.

The first-time setup on the phone involved three steps: launching the app, entering credentials for the account that had been created in advance on the backend, and confirming that push notifications were enabled at the operating system level. Push notifications were routed through Firebase Cloud Messaging, and a test alert was sent from the backend to verify delivery before the observation window began.

Live sensor readings appeared on the Dashboard tab within seconds of the first ESP32 packet arriving at the backend. The farmer was walked through each tab of the application on the day of deployment, with emphasis on the Kandang tab where the daily egg count would be entered every morning.

3.3.4 Field data collection execution

Data collection ran automatically from 26 April to 25 May 2026, a full thirty consecutive days after the setup day. The routine over that period was largely uneventful. The ESP32 transmitted sensor readings every 30 minutes without manual intervention. The backend received and stored those readings without any dropped connections beyond the one Wi-Fi outage described below. The farmer opened the Kandang tab each morning after collecting eggs and entered the daily count.

Two events during the observation window are worth documenting. On 6 May 2026, heavy afternoon rainfall caused a Wi-Fi outage at the farm that lasted from approximately 14:00 until 09:00 the next morning. The ESP32 firmware buffered readings locally during the outage, but the flash buffer's capacity was exceeded and readings from 15:00 through midnight were lost. When connectivity returned, the ESP32 resumed normal operation and posted a full day of readings on 7 May onward. The 6 May day was therefore excluded from all subsequent analysis, giving 30 valid days of data rather than 31.

On several days in the second half of May, the DHT22 humidity sensor reported the ceiling value of 99.9% for extended stretches. This was investigated and confirmed as sensor saturation rather than a firmware bug the surrounding weather was consistently rainy and the

actual humidity was likely at or above the sensor's measurement ceiling. No corrective action was taken because replacing the sensor mid-experiment would have introduced a discontinuity in the data.

Raw sensor records were cleaned before analysis using the interquartile range (IQR) method applied independently to each parameter. For each channel, the first and third quartiles were computed across all readings, and any value falling more than 1.5×IQR below Q1 or above Q3 was flagged as an outlier and removed from the analytical dataset. This removed 47 humidity readings, 32 ammonia readings, and 55 illuminance readings out of the 1,274 total raw records. Temperature had no readings rejected. The cleaned dataset used for all downstream analysis contained 1,140 records across the 30 valid days. A summary of the data collection log is given in Table 3.22.

Table 3.22 Data collection log summary

Metric	Value
Observation window	26 April – 25 May 2026
Total calendar days	31
Valid data days	30 (6 May excluded)
Raw sensor records received	1,274
Records removed as outliers (IQR)	134 (47 humidity, 32 ammonia, 55 illuminance)
Cleaned records used for analysis	1,140
Egg count entries (farmer manual)	30
LLM diagnostic outputs generated	29

3.4 Testing Phase

The testing phase used four evaluation methods running in parallel throughout the observation window and continuing after it ended. Black-box functional testing verified that each part of the system did what it was designed to do. Statistical correlation analysis tested whether the daily ESI actually tracked laying performance. LLM consistency analysis measured how often Claude Sonnet's diagnostic output agreed with the ESI-derived classification. Expert validation triangulated the quantitative findings against the farmer's lived experience and an animal husbandry specialist's independent assessment. A fifth activity case study documentation captured specific episodes where the system's alerts appeared to trigger farmer action.

3.4.1 Black-box functional testing

Black-box testing checks the input-output behavior of a system without inspecting internal implementation. For this pilot, seven functional scenarios were defined based on the requirements gathered in Section 3.1.2 and the design decisions made in Section 3.2. Each scenario specified a triggering condition, an expected outcome, and the observable evidence needed to declare pass or fail. Table 3.23 lists the scenarios.

Table 3.23 Black-box functional test scenarios

Scenario	Expected outcome
Sensor data acquisition	ESP32 posts a valid JSON record to /api/sensor-data every 30 minutes; record persists in sensor_data table with correct timestamp
Nightly ESI computation	Scheduler fires at 22:00; ESI value, sub-scores, and classification band appear in esi_results for the day
LLM diagnostic generation	LLM call completes within 60 seconds of ESI computation; structured JSON is stored in llm_diagnoses with all five expected fields populated
Push notification on Emergency	When LLM returns cage_status = Emergency, a push notification arrives on the farmer's phone within 5 minutes
Dashboard rendering	Mobile Dashboard tab displays the four latest sensor readings, ESI gauge, lamp status, and latest LLM diagnosis on refresh
Historical data retrieval	Mobile app retrieves and displays sensor and ESI history for any date range in the observation window
LLM recommendation relevance	Sampled diagnostic outputs address the dominant environmental stressor for the corresponding day (validated by comparing recommendation text to sensor summary)

Tests were run continuously across the observation window rather than as a separate testing phase, because the deployment context did not permit taking the system offline for scheduled testing. Each scenario was marked passed when its expected outcome was observed at least once during the 30-day window, and marked failed if the expected outcome never occurred or was observed to fail consistently. Results are reported in Section 4.8.1.

3.4.2 Statistical analysis procedure

The core statistical question is whether daily ESI values are associated with daily egg production. This was tested using Pearson product-moment correlation, which measures the strength and direction of a linear relationship between two continuous variables. Pearson correlation was chosen over rank-based alternatives (Spearman, Kendall) because both variables are continuous and normally distributed, and because the primary interest is in linear rather than monotonic association. The analysis was implemented in Python 3.11 using

scipy.stats.pearsonr, which returns both the correlation coefficient r and the two-tailed p-value under the null hypothesis of zero correlation. The independent variable was the daily ESI drawn from the esi_results table, the dependent variable was the daily egg count from income_eggs, the sample size was n = 30 valid days, and the significance threshold was set at $\alpha = 0.05$.

Beyond the primary composite-vs-egg correlation, a channel-wise analysis was also planned. Each individual sensor reading (temperature, humidity, ammonia, illuminance), each sub-score (S_THI, S_NH3, S_Light), and the derived THI value were correlated independently against daily egg count. This channel-wise analysis identifies which environmental variable carries the strongest predictive signal, and lets the composite index be compared against its components rather than treated as a black box.

Table 3.24 Interpretation criteria for correlation strength

|r| range	**Strength**	**Interpretation**
0.00 – 0.10	Negligible	No meaningful relationship
0.10 – 0.30	Weak	Small effect, may not be practically significant
0.30 – 0.50	Moderate	Meaningful relationship worth investigating
0.50 – 0.70	Strong	Substantial predictive value
0.70 – 1.00	Very strong	Near-deterministic relationship

3.4.3 LLM consistency analysis

The LLM diagnostic layer produces a cage_status field with values Normal, Warning, Emergency, or Critical. The ESI computation produces a classification band with values Good, Needs Attention, or Critical. These two output spaces do not map perfectly one-to-one, so the first analytical step was to define a consistency framework.

Table 3.25 gives the mapping used. LLM values Normal and Warning were aggregated to match ESI's Needs Attention and Good bands, while Emergency and Critical LLM outputs were both treated as equivalent to ESI's Critical classification. This is a conservative mapping it groups any elevated LLM concern with the ESI Critical band rather than trying to distinguish Warning-level and Emergency-level severity.

Table 3.25 LLM vs ESI classification comparison framework

ESI category	LLM cage_status (matching)	LLM cage_status (more severe)
Good	Normal	Warning, Emergency, Critical
Needs Attention	Warning	Emergency, Critical
Critical	Emergency, Critical	(no more severe category available)

From this mapping, three metrics were computed. Overall agreement rate is the proportion of days where the LLM classification matches the ESI category exactly. Direction of disagreement is the proportion of disagreements where the LLM assigned a more severe category than the ESI versus a less severe category. Sensitivity for the Critical band is the proportion of ESI-Critical days that were also flagged as Critical or Emergency by the LLM.

The sensitivity metric is the most important of the three because it measures whether the LLM catches the worst days. A monitoring system that lets ESI-Critical days slip through with a Normal LLM classification would be dangerous to deploy. Results and interpretation are reported in Sections 4.5.1 through 4.5.4.

3.4.4 Expert validation protocol

The interview protocol used three instruments spread across the timeline of the study. Instrument 1 collected baseline data from the farmer before deployment. Instrument 2 collected scientific validation from an animal husbandry student. Instrument 3 was designed for post-deployment member checking with the same farmer, closing the qualitative loop after the quantitative data had been analyzed.

Instrument 1 was semi-structured and administered to Pak Suyud on 8 March 2026 at the farm. Its purpose was to establish what the farmer currently does to monitor cage conditions, what he currently cannot see, and what a phone-based system would need to provide to be useful to him. Table 3.26 lists the topic areas covered.

Table 3.26 Interview instrument 1 pre-deployment farmer

Topic	Purpose
Farm profile and experience	Farm profile and experience
Current monitoring practice	Document what the farmer does now to assess cage conditions
Pain points and information gaps	Identify what the farmer cannot see or reacts to too late

Pancaroba experience	Confirm the seasonal pattern from the farmer's perspective
Preferences for phone-based information	Define what the mobile app needs to show

Instrument 2 was conducted by video call on 15 March 2026 with Nabila, a sixth-semester animal husbandry student at Universitas Gadjah Mada. Its purpose was to independently validate the choice of environmental parameters and the threshold values before they were programmed into the firmware. Table 3.27 lists the topic areas.

Table 3.27 Interview instrument 2 animal husbandry student

Topic	Purpose
Parameter selection	Validate the four monitored parameters against physiological literature
Threshold values	Confirm the THI, ammonia, and light thresholds
Lighting recommendations	Confirm the dark window and lamp control approach
LLM diagnostic layer	Get expert perspective on trust and validation of AI outputs

Instrument 3 is scheduled to be administered after the observation window closes, with the same farmer. Its purpose is to close the loop by comparing the farmer's post-deployment perception against the quantitative findings, and to check whether the system changed his behavior or his awareness of cage conditions. Table 3.28 lists the topic areas.

Table 3.28 Interview instrument 3 post-deployment farmer

Topic	Purpose
System usage over deployment	Which features were used, how often
Response to alerts	Whether specific alerts (Emergency-level) prompted actual action
Perceived change in monitoring practice	Compare against pre-deployment baseline from Instrument 1
Trust in LLM diagnosis	Whether the farmer trusted the AI recommendations enough to act
Overall satisfaction and adoption	Would the farmer recommend the system to others

All interviews were conducted in Bahasa Indonesia, recorded with permission, transcribed within seven days, and thematically coded. Direct quotations used in Chapter IV are marked as such; paraphrases summarize the sense of longer answers. Results from

Instruments 1 and 2 are integrated into Section 4.8.2. Results from Instrument 3 will be added when the post-deployment session is complete.

3.4.5 Case study documentation

Alongside the aggregate statistical analysis, specific episodes during the observation window were documented as case studies. The purpose of case study documentation is to capture the causal chain between a system alert and subsequent action, which aggregate statistics cannot resolve on their own.

A case study episode was defined as any five-day window that contained at least one ESI-Critical day and at least one Emergency-level LLM alert. For each identified episode, the following data was compiled: the daily sensor summary across all four parameters, the ESI and its sub-scores, the LLM diagnostic output including recommended actions, the daily egg count from farmer manual entry, and any farmer action reported in the mid-deployment or post-deployment interviews. One such episode occurred between 20 and 24 May 2026 and is reported in detail in Section 4.6.

CHAPTER IV
RESULT AND DISCUSSION

This chapter presents the outcomes of the thirty-day field deployment at Siwatu and analyses what those outcomes mean for the research questions posed in Chapter I. The reporting order follows the logical flow of the study: what data was collected, what the environmental conditions looked like, whether the Environmental Stress Index tracked laying performance the way the theory predicted, and how the Claude Sonnet diagnostic layer performed against the ESI baseline. Four practical findings emerge a moderate but statistically significant correlation between ESI and egg production, an asymmetric precaution pattern in the LLM outputs, a documented intervention episode across 20–24 May, and a systematic LLM over-escalation on rainfall days.

4.1 Deployment outcome and data collection summary

Before analysing the sensor readings, this section establishes what data actually came in during the observation window and how much was usable after cleaning. The deployment itself was uneventful in the sense that the hardware ran continuously without physical intervention, but the collection was not perfect one calendar day produced no records at all, and outlier filtering removed a portion of raw readings.

4.1.1 What was deployed and what ran

The IoT-ESI-LLM system operated at Siwatu from 25 April to 25 May 2026, covering thirty-one calendar days. Sensor packets arrived on the 30-minute cadence for thirty of those days. The single day with no data was 6 May, when the ESP32 lost Wi-Fi connectivity during a heavy afternoon storm and did not reconnect until the following morning. The analysis in the rest of this chapter therefore uses n = 30 days rather than 31.

Table 4.1 Data collection totals

Data type	Expected	Actual	Notes
Sensor readings (raw)	~1,440	1,274	6 May missing
Sensor readings after outlier removal (IQR)		1,140	47 humidity, 32 NH_3, 55 lux removed
Daily aggregated summaries	31	30	30 valid days

Data type	Expected	Actual	Notes
Egg count entries (farmer, morning collection)	31	30	Matches sensor days
LLM diagnoses generated	30	29	1 skipped (no sensor data)

Outlier removal used the interquartile range method independently for each parameter. The humidity outlier count is worth flagging early because it reflects a specific sensor limitation: the DHT22 saturates at 99.9%, and during long rain-affected afternoons it reports the ceiling value for hours, which then reads as extreme against the day's own distribution. This is a genuine measurement issue that surfaces again in Section 4.7 when discussing the LLM's response to rainfall days.

4.2 Environmental characteristics during pancaroba

Before turning to whether the ESI correlated with egg production, it is worth documenting what the pancaroba environment actually looked like inside the cage. This section provides the empirical foundation for the rest of the chapter the range each parameter moved through, how they evolved day by day, and what shape the season took as seen through the four sensors.

4.2.1 Descriptive statistics across the 30 valid days

Table 4.2 summarises the aggregate distribution of each channel across all 1,140 cleaned readings. These are the numbers that anchor the interpretation of everything that follows.

Table 4.2 Descriptive statistics of the four monitored channels

Channel	Min	Max	Mean	SD	n
Temperature (°C)	24.50	34.80	28.36	2.25	1,140
Humidity (%)	67.50	99.90	93.27	7.81	1,140
Ammonia (ppm)	0.80	29.50	14.39	4.95	1,140
Illuminance (lux)	0.00	191.70	40.32	50.31	1,140

Three observations frame the reading of this table.

Temperature averaged 28.36 °C, which places it inside the published quail comfort band of 20–30 °C but at the upper end. A maximum of 34.8 °C indicates that the cage crossed into heat stress at least once during the observation window, and the standard deviation of 2.25 °C says that within-day swings were meaningful.

Humidity averaged 93.27%, which is well above the 70% comfort ceiling used in the animal science literature. This is a seasonal signature rather than a housing failure the animal husbandry specialist consulted in Chapter III predicted saturated humidity throughout the observation window, and the data confirms it.

Ammonia averaged 14.39 ppm across all readings. Ten parts per million is the safe operating ceiling and twenty-five is the acute injury threshold. The mean therefore sits inside the warning band, and the maximum of 29.5 ppm crosses the danger threshold on individual readings.

4.2.2 Daily evolution of each parameter

To see how the four channels behaved across the observation window, Figure 4.1 presents each parameter as a daily time series with the min-max band shaded around the mean line for temperature and humidity.

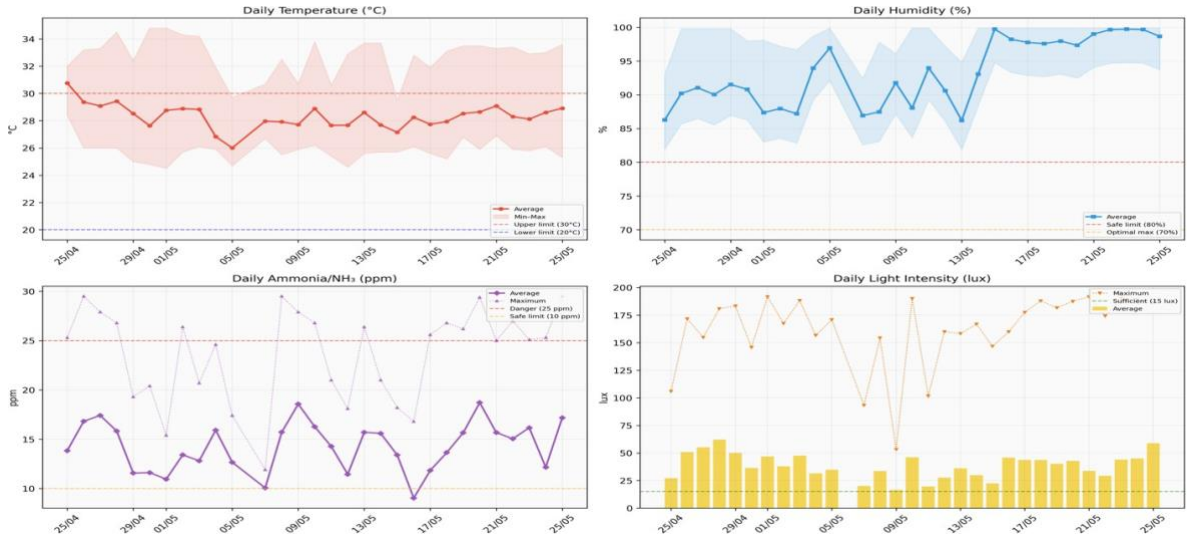


Figure 4.1 Daily cage environmental profile

The daily temperature panel shows that mean daily temperature stayed above 26 °C on every day and crossed the 30 °C safety line on nine days. The wide shaded band during the first two weeks confirms that individual readings frequently touched 33–34 °C.

The humidity panel makes the seasonal shift visible. From 25 April to 12 May the daily mean sat between 86% and 94% with visible day-to-day variation. From 13 May onwards it locked between 97% and 100%, and stayed there. This is the point at which the rainy phase of pancaroba established itself definitively.

The ammonia panel shows that the daily mean stayed above the 10 ppm safe threshold on almost every day, and the daily maximum crossed the 25 ppm danger line ten times, with the highest single reading of 29.5 ppm occurring during the case study window discussed in Section 4.6.

The illuminance panel is the most volatile because it swings between zero at night and around 190 lux at noon. What matters here is not the maximum but the daily mean, shown as the yellow bars: on many days the average illuminance stayed below the 15-lux floor recommended for active hours.

4.2.3 Composite thermal load through THI

Because heat stress in laying quail depends jointly on temperature and humidity rather than either alone, Figure 4.2 shows the daily Temperature-Humidity Index, which combines the two.

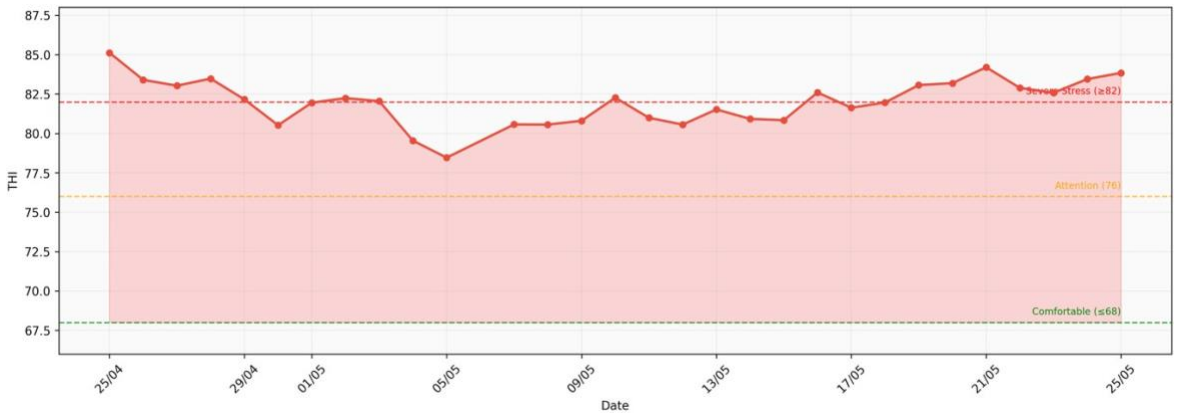


Figure 4.2 Daily THI with stress thresholds

The result is striking in what it excludes. Not a single day during the observation window had a THI in the comfort band (≤68). Every daily value sat between 78 and 86, and 25 of the

30 days crossed the severe-stress threshold at 82. This directly supports the pancaroba framing of the whole thesis the birds experienced continuous thermal loading throughout the study period, punctuated by peaks rather than interrupted by recovery.

4.3 Environmental Stress Index distribution and behaviour

With the raw environmental picture established, this section moves to the composite index derived from those readings. The purpose here is to show how the ESI distributed across the thirty days, what its typical value was, and which sub-scores drove it.

4.3.1 Descriptive statistics across the 30 valid days

Figure 4.3 presents the daily ESI as a colour-coded bar chart. The colouring makes the classification immediately readable: orange for Needs Attention and red for Critical.

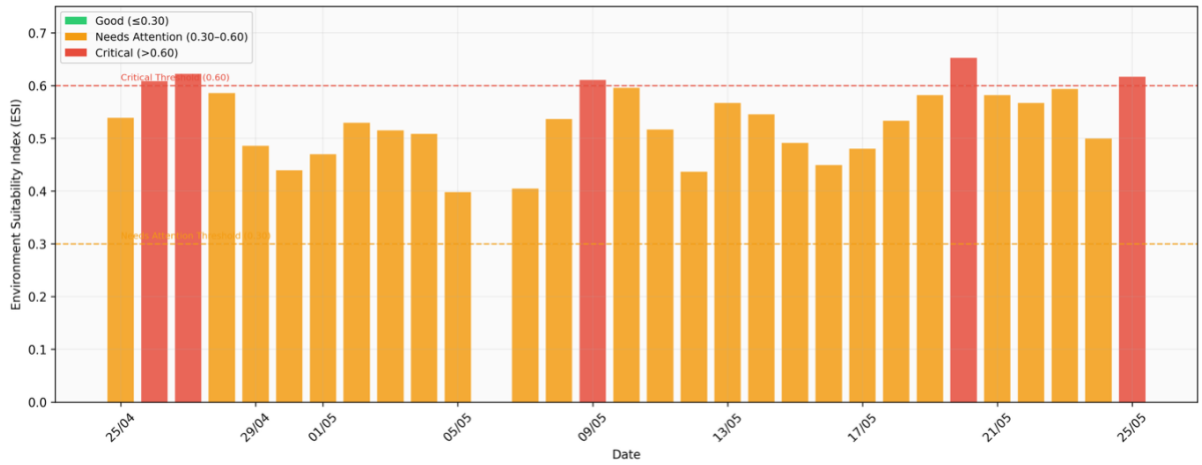


Figure 4.3 Daily ESI values for the 30-day observation window

The distribution split across the three official ESI bands is given in Table 4.3.

Table 4.3 ESI classification distribution across 30 days

ESI band	Range	Days	Percentage
Good	≤ 0.30	0	0.0%
Needs Attention	0.31–0.60	25	83.3%
Critical	> 0.60	5	16.7%

The most important pattern in this table is what does not appear in it: zero days in the Good band. Across a full month of continuous monitoring, the birds never experienced a day the composite index would classify as environmentally satisfactory. Five days pushed into Critical, and the remaining twenty-five sat in Needs Attention.

4.3.2 Distribution shape

To see how tightly the daily values clustered around the mean, Figure 4.4 shows the ESI distribution as a histogram.

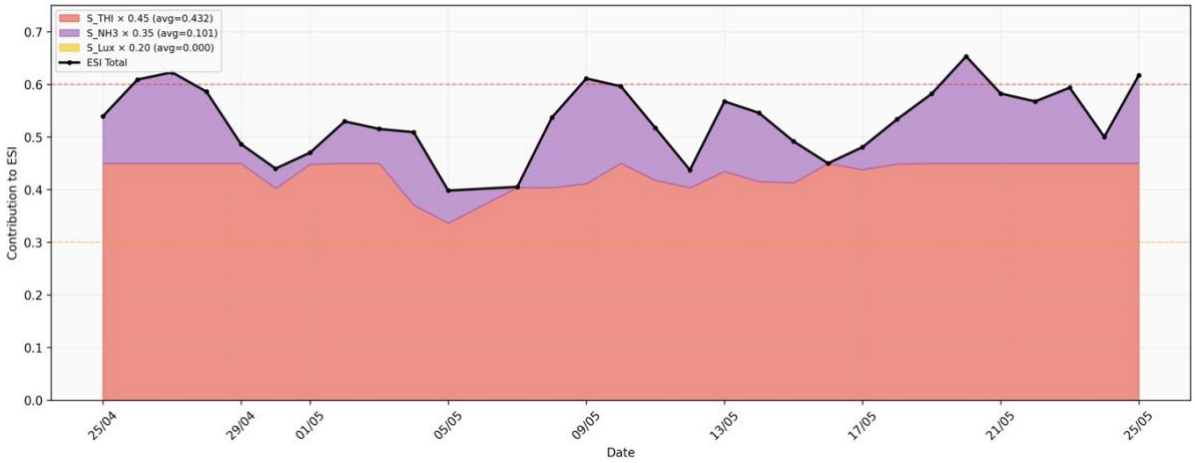


Figure 4.4 ESI distribution histogram

The distribution is centred at 0.5326 with a standard deviation of 0.0673. This means roughly 95% of daily values fell between 0.40 and 0.67 a narrow operating band entirely inside the Needs Attention and Critical zones. The index simply did not have the range in this dataset to differentiate acceptable conditions from good conditions, because acceptable conditions did not occur.

4.3.3 Sub-score contributions

To understand which environmental component pushed the ESI up on any given day, Figure 4.5 breaks the daily ESI into its three weighted sub-score contributions as a stacked area plot.

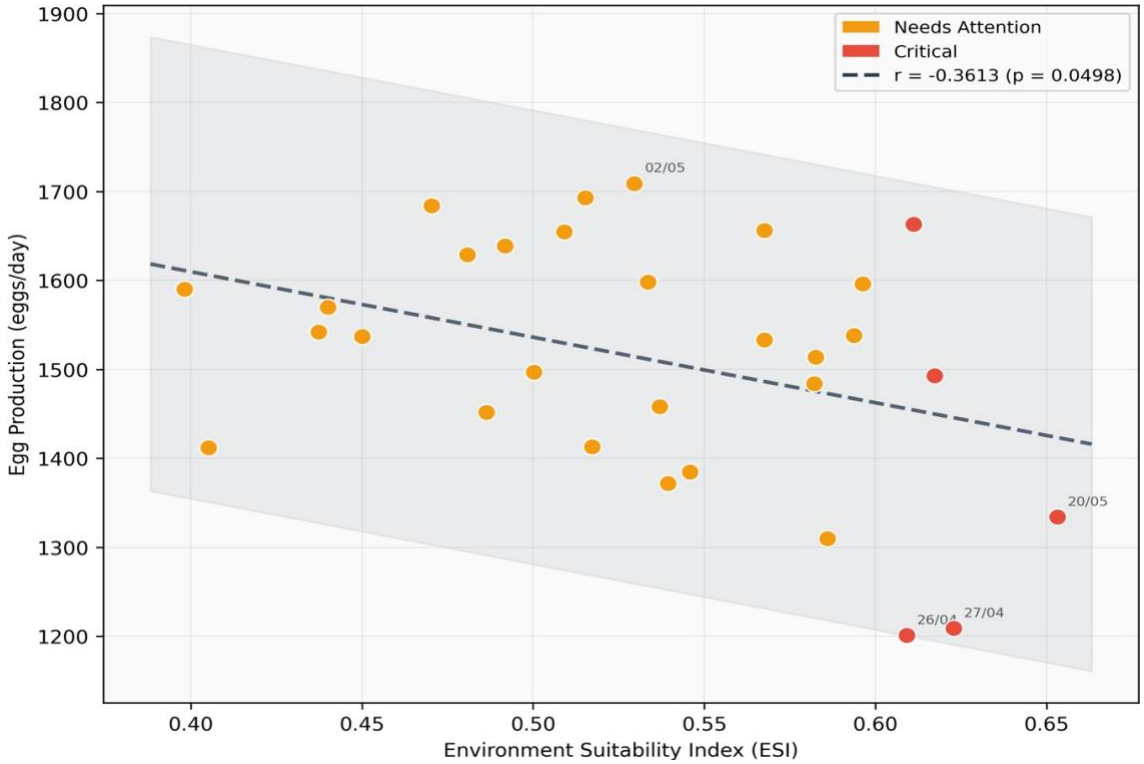


Figure 4.5 ESI sub-score composition

The thermal sub-score dominates the composite on almost every day. With an average contribution of 0.432 out of a possible 0.45, S_THI runs at near-saturation throughout the observation window. This reflects the underlying THI already shown in Figure 4.2 with THI values consistently above 78, the S_THI sub-score is at or near its maximum.

The ammonia sub-score is the swing variable. Its average contribution of 0.101 is much lower than S_THI, but unlike the thermal component it varies substantially from day to day, and it is what pushes borderline days into the Critical band. Its mean weighted contribution was 0.181 on the five Critical days against 0.085 on the twenty-five Needs Attention days. On days like 20 May when ammonia averaged 18.71 ppm, the S_NH3 contribution alone added over 0.20 to the composite. Section 4.3.4 tests this attribution directly through a counterfactual analysis and shows that all five Critical classifications are attributable to ammonia rather than thermal load.

The light sub-score is effectively zero across the whole month. This is because daily average illuminance stayed above the 15-lux floor during active hours on almost every day, and the lamp actuator kicked in when it dropped below. The light contribution therefore does

not drive variation in this pilot but it would matter in a different climate or during longer overcast stretches.

4.3.4 Anatomy of the five Critical days

Table 4.3 established that five of the thirty observed days crossed into the Critical band (ESI > 0.60). Because these are the only days on which the composite index judged the cage environment to be dangerous rather than merely suboptimal, they warrant individual examination. This subsection identifies what occurred on each day, isolates the responsible environmental component, compares the five days against the twenty-five Needs Attention days, and identifies practical measures that follow from the evidence. Table 4.4 presents the sensor conditions for each Critical day together with the number of individual 30-minute readings on which ammonia crossed the 25 ppm acute-injury threshold.

Table 4.4 ESI classification distribution across 30 days

Date	T avg (°C)	RH avg (%)	NH_3 avg (ppm)	NH_3 max (ppm)	Readings $NH_3 \geq 25$	THI	S_THI	S_NH_3	ESI	Eggs
26 Apr	29.37	90.20	16.82	29.5	2	83.40	1.000	0.455	0.6091	1,201
27 Apr	29.08	91.05	17.41	27.9	5	83.02	1.000	0.494	0.6229	1,209
9 May	27.72	91.78	18.56	27.9	4	80.81	0.915	0.571	0.6112	1,663
20 May	28.65	97.34	18.71	29.4	13	83.52	1.000	0.599	0.6532	1,334
25 May	28.91	98.66	17.17	29.5	2	84.03	1.000	0.477	0.6173	1,493
NA mean (n = 25)	28.29	92.95	14.25	22.68	-	81.92	0.957	0.243	0.5146	1,539

The 26 April event was a sustained heat day at the start of pancaroba. Temperature exceeded 30 °C for twenty out of forty-eight readings from 11:12 to 20:42, peaking at 33.2 °C, while ammonia climbed to 29.5 ppm at 16:12. Overnight humidity had saturated at 99.8 %, trapping ammonia from decomposing litter in stagnant air so that the concentration was already above 11 ppm by morning. Egg production dropped to 1,201, the lowest count of the study. The 27 April event was a direct continuation: ammonia had not been cleared overnight and reached 25.4 ppm at 04:42 and 27.9 ppm at 05:42 while the birds were still in their rest period before the daytime heat cycle repeated.

Egg production remained depressed at 1,209, confirming cumulative stress with no recovery window. The 9 May event had a different character: temperature was the lowest of the five Critical days at 27.72 °C (S_THI = 0.915, the only day not at the thermal ceiling), and what pushed ESI above 0.60 was ammonia alone at 18.56 ppm ($S_NH_3 = 0.571$). Despite the Critical classification, eggs reached 1,663 above the study mean because post-rain cooling benefited the birds even as stagnant conditions allowed ammonia to accumulate. The 20 May event was the most severe episode of the entire study: temperature peaked at 36.0 °C at 13:37, ammonia peaked at 33.9 ppm at 18:07 with thirteen readings above 25 ppm, and humidity was at 99.9 % for thirty-three out of forty-eight readings. A rain event around 16:00 sealed humid air inside the cage, and ammonia remained above 25 ppm for seven consecutive readings from 18:07 to 20:37. Eggs dropped to 1,334 but recovered by 180 the following day after the farmer received an Emergency alert. The 25 May event was characterised by near-permanent humidity saturation (thirty-six out of forty-six readings above 95 %) with the familiar late-afternoon convergence: heat peaked at 34.7 °C, rain sealed the cage after 15:34, and ammonia spiked to 29.5 ppm at 17:04.

The three sub-scores do not contribute equally to these classifications. Averaged across the Critical days, the thermal component contributed 0.442 out of a maximum 0.45, but the same component contributed 0.430 on the twenty-five Needs Attention days a difference of only 0.012. Thermal loading was near-saturated on almost every day of the study and cannot be what distinguishes the two groups. The ammonia component behaved differently: its mean weighted contribution was 0.181 on Critical days against 0.085 on Needs Attention days, more than double. Because the light sub-score was zero throughout, ammonia is the only component with enough day-to-day variation to move the composite across the 0.60 boundary. This attribution can be tested directly. Table 4.5 recomputes each Critical day's ESI with ammonia held at 14.25 ppm the mean of the non-Critical days while leaving temperature, humidity, and illuminance untouched.

Table 4.5 Counterfactual ESI with ammonia held at the non-Critical mean

Date	Actual ESI	Counterfactual ESI	Resulting classification
26 Apr	0.6091	0.5492	Needs Attention
27 Apr	0.6229	0.5492	Needs Attention
9 May	0.6112	0.5106	Needs Attention
20 May	0.6532	0.5492	Needs Attention
25 May	0.6173	0.5492	Needs Attention

All five days lose their Critical classification when ammonia is normalised. Ammonia accumulation is therefore the sole proximate cause of every Critical classification recorded in this study no day reached the Critical band on thermal load alone. One qualification must be stated: four of the five days exceeded the 0.60 threshold by a margin below 0.023 (0.009 on 26 April, 0.011 on 9 May, 0.017 on 25 May, 0.023 on 27 April), and only 20 May cleared it decisively at 0.053 above. The classification of the four borderline days is consequently sensitive to the weighting scheme, which is recorded among the limitations in Section 4.9.2.

The Critical days did not arrive without warning. On the three days preceding 20 May, daily mean ammonia rose monotonically 11.83, 13.64, then 15.66 ppm before reaching 18.71 ppm, while humidity remained locked between 97 % and 98 %. The 26–27 April pair shows the same ramp, climbing from 13.83 ppm on 25 April to 16.82 and then 17.41 ppm. The physical mechanism is microbial decomposition of nitrogenous waste in litter whose moisture content rises with ambient humidity; under saturated conditions the litter cannot dry between cleaning cycles, so each day's waste adds to a residue the previous day failed to clear. The within-day timing reinforces this: across the five Critical days, ammonia readings above 25 ppm clustered between 16:00 and 20:00. In four of the five cases a rain event in the afternoon sealed humid air inside the cage, triggering the peak. The result is that Critical episodes were in principle detectable two to three days in advance from the ammonia trend alone.

Table 4.6 Critical days compared with Needs Attention days

Measure	Critical (n = 5)	Needs Attention (n = 25)	Difference
Mean ESI	0.6224	0.5146	+0.1078
Mean temperature (°C)	28.83	28.29	+0.54
Mean humidity (%)	93.55	92.95	+0.60
Mean NH_3 (ppm)	17.73	14.25	+3.48 (+24.4 %)
Mean NH_3 max (ppm)	28.84	22.68	+6.16 (+27.2 %)
Mean weighted thermal contr.	0.442	0.430	+0.012
Mean weighted NH_3 contr.	0.181	0.085	+0.096 (+112.9 %)
Mean daily eggs	1,380.0	1,538.6	−158.6 (−10.3 %)
SD daily eggs	197.6	108.9	-

Mean daily production on Critical days was 158.6 eggs lower, a shortfall of 10.3 %. The direction is consistent with the month-level correlation of $r = -0.361$ (Section 4.4.2), but the difference does not reach statistical significance on its own: a Welch two-sample t-test gives $p = 0.1484$ and a Mann-Whitney U test gives $p = 0.0846$. With only five days in the Critical group the test has little power, and the 9 May anomaly (1,663 eggs despite Critical ESI) runs

against the pattern. The shortfall is therefore reported as a descriptive difference consistent in direction with the correlation analysis rather than a demonstrated effect. A recovery pattern does emerge, however: egg counts on the days following 26 April, 27 April, 9 May, and 20 May were 1,209, 1,310, 1,596, and 1,514 respectively a return toward or above the study mean except where the consecutive 26–27 April pair interrupted recovery. After 20 May, production rose by 180 eggs in a single day, coinciding with the ammonia decline documented in Section 4.6.

Because ammonia was the operative cause in all five cases and wet litter is the mechanism by which it accumulates, the appropriate responses target litter moisture and air exchange rather than temperature. Four measures follow from the evidence. First, litter management is the primary lever: the 27 April dawn readings (25.4 ppm at 04:42) demonstrate that a single daily cleaning is insufficient when humidity prevents litter from drying overnight, so scheduled replacement twice daily during pancaroba with the second cleaning before 15:00 to pre-empt the afternoon peak addresses the source directly; where drying is not achievable, absorbent amendments such as zeolite or agricultural lime bind ammonia that evaporation cannot remove. Second, forced ventilation should be concentrated in the late afternoon between 16:00 and 20:00, when exceedance events cluster; the relay module already has a second channel wired for an exhaust fan (Section 2.4), and the firmware could switch it on automatically when ammonia crosses 15 ppm. Third, a rain cover over ventilation openings would allow continuous airflow during rain events, preventing the afternoon sealing pattern observed on 20 May and 25 May. Fourth, alerting should trigger on the trend rather than the threshold: an alert raised when daily mean ammonia crosses 15 ppm while rising over two consecutive days would have flagged the 20 May episode on 18 May, two days before the ESI crossed into Critical, shifting the system from reactive detection to predictive warning.

4.4 Relationship between ESI and egg production

This is where the fundamental hypothesis of the thesis is tested: does the composite environmental index actually track laying performance in a real working cage? The rest of the chapter builds on the answer given here.

4.4.1 Egg production during the observation window

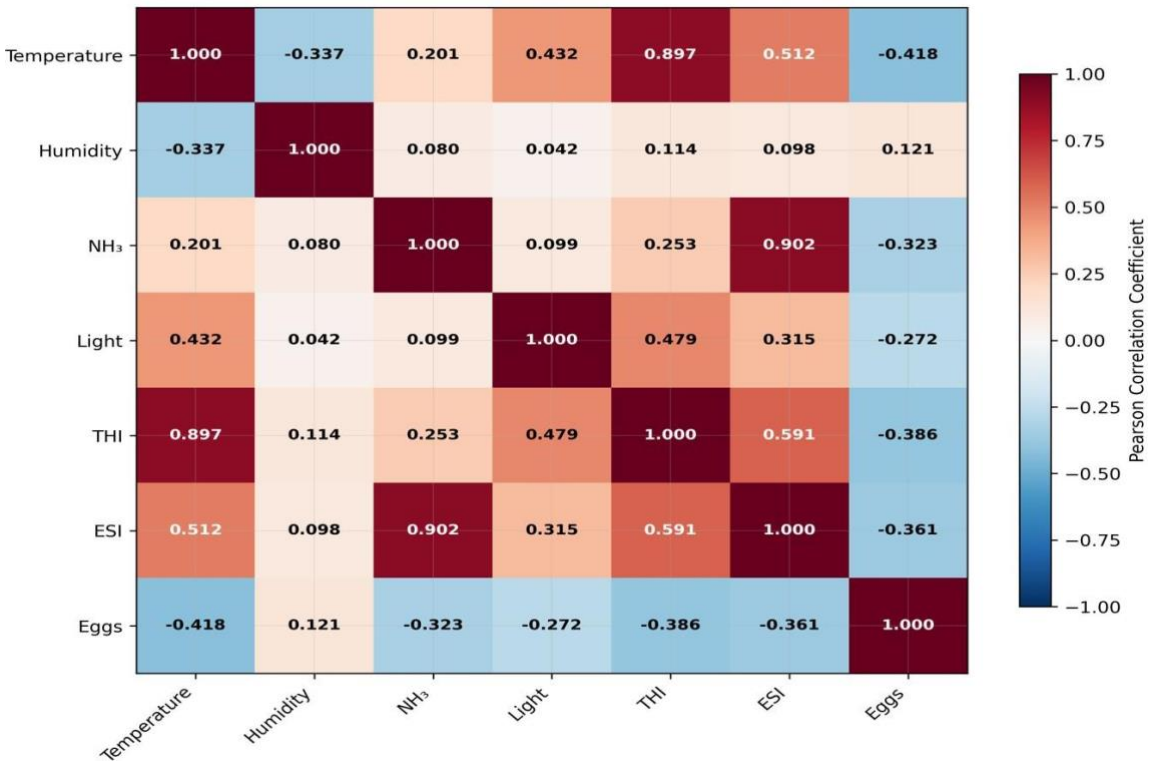


Figure 4.6 Egg production vs ESI (dual axis)

Daily egg counts ranged from 1,201 (26 April) to 1,709 (2 May), with a mean of 1,512 and a standard deviation of 137. The three lowest-yield days 26 April (1,201), 27 April (1,209), and 20 May (1,334) all coincide with the highest ESI values. The detailed analysis in Section 4.3.4 confirms that these production drops were driven by elevated ammonia rather than temperature, with the counterfactual analysis showing that no Critical classification would have occurred on thermal load alone. The visual overlay in Figure 4.6 makes the inverse pattern apparent: peaks in the orange ESI bars align with troughs in the blue egg-production line.

4.4.2 Pearson correlation: composite ESI vs egg count

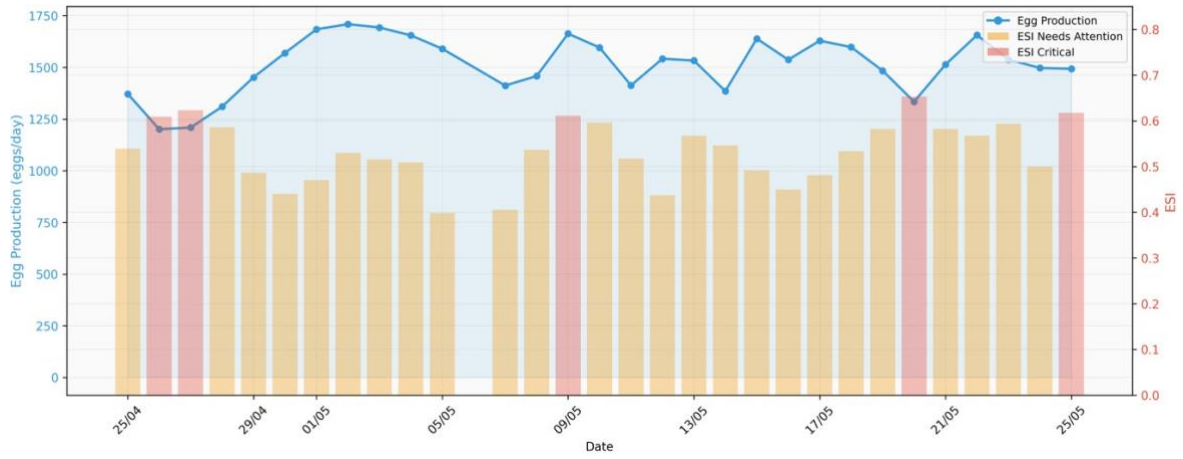


Figure 4.7 ESI vs egg count scatter plot

The Pearson product-moment coefficient between daily ESI and daily egg count came out at $r = -0.3613$, with a two-tailed p-value of 0.0498 across $n = 30$ days. Three points about this result need to be made clearly.

The direction of the correlation is the direction the theory predicted. Higher ESI values are associated with lower egg counts, which validates the fundamental construction of the index. If the correlation had been positive, or near zero, the whole framework would need reconsidering.

The magnitude is moderate. A coefficient of −0.361 corresponds to $r^2 = 0.13$, meaning the ESI explains about thirteen percent of day-to-day variance in egg production. This is not a strong predictive relationship in absolute terms, but it is defensible for a single-cage pilot subject to unmeasured confounders feed composition changes, individual bird variation, the farmer's own interventions in response to the alerts the system generated, and other factors that were not controlled.

The p-value sits at the conventional significance threshold. At $p = 0.0498$, the result clears the $\alpha = 0.05$ test by two thousandths of a point. Reporting this honestly means acknowledging that the association is preliminary it exists in this dataset and it points in the expected direction, but confirming it would require either a longer observation window or a multi-cage follow-up study.

4.4.3 Channel-wise correlations to identify the strongest predictor

The composite index is not the only variable that might correlate with egg output. To understand which individual sensor readings carried the predictive signal, Pearson correlations were also computed for each channel independently. Figure 4.8 shows the full correlation matrix as a heatmap, and Table 4.4 extracts the correlations against egg production ranked by strength.

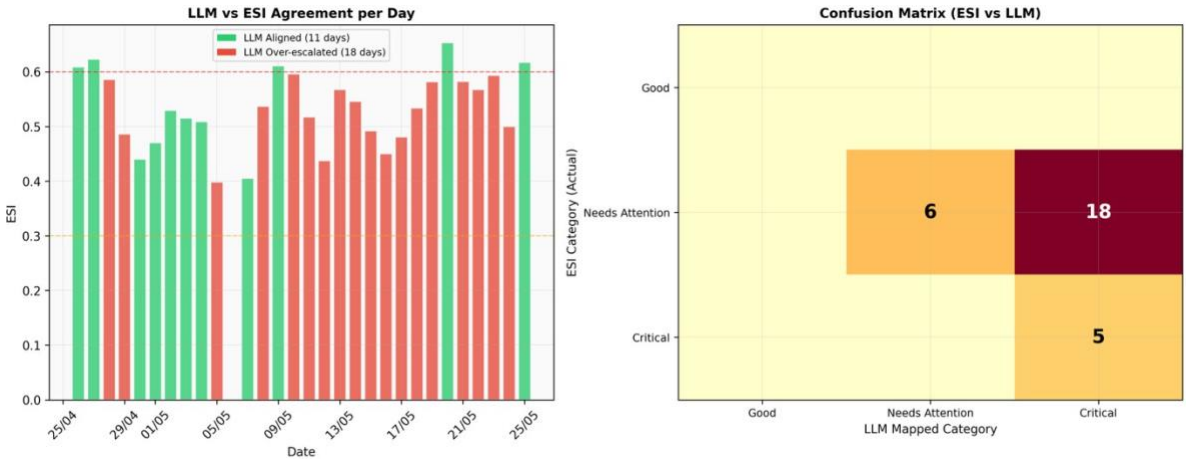


Figure 4.8 Correlation matrix of environmental parameters and egg production

Table 4.7 Channel-wise Pearson correlation with daily egg production

Variable	r	p	Verdict
Temperature (daily avg)	−0.418	0.022	Significant
THI (composite)	−0.386	0.035	Significant
Composite ESI	−0.361	0.050	Significant (borderline)
Ammonia sub-score (S_NH3)	−0.329	0.076	Marginal ($p < 0.10$)
Ammonia (daily avg)	−0.323	0.082	Marginal ($p < 0.10$)
Illuminance (daily avg)	−0.272	0.146	Not significant
Ammonia (daily max)	−0.237	0.207	Not significant
S_THI	−0.197	0.297	Not significant
Humidity (daily avg)	+0.121	0.523	Not significant, wrong sign
Temperature (daily max)	−0.053	0.780	Not significant

Two features of this ranking need discussion.

Temperature alone was the strongest single predictor of egg output, with $r = -0.418$ ($p = 0.022$). THI came second at $r = -0.386$, and the composite ESI third at $r = -0.361$. That temperature outperforms the composite is not a failure of the ESI it means that in this particular dataset, thermal loading was so dominant that adding ammonia and light diluted rather than sharpened the correlation. In a different dataset, where ammonia or light varied more, the composite would likely take over as the strongest predictor.

Humidity by itself came out with a small positive correlation to egg output ($r = +0.121$, $p = 0.523$), which looks counter-intuitive. The explanation lies in what humidity co-varies with during pancaroba: humid days were rainy days, and rainy days brought lower temperature. The thermal cooling more than compensated for the humidity increase, at least at the range of values observed here. This is exactly the interaction the THI formula is designed to capture combining temperature and humidity produces a much stronger negative correlation with eggs than either variable alone.

4.4.4 What this means for the ESI framework

The channel-wise results validate the choice to construct a composite rather than rely on a single sensor. Temperature is the strongest predictor when it dominates, but a temperature-only monitoring system would have missed the 20 May ammonia peak entirely and, as the counterfactual analysis in Section 4.3.4 shows, it would have missed all five Critical days of the study, none of which reached the Critical band on thermal load alone. The composite ESI provides the robustness needed for a monitoring system that has to work across whatever environmental stressor happens to matter on any given day.

This finding aligns with the THI validation for laying quail reported by El-Tarabany (2016), who showed that quail exposed to high THI levels (76–80) had significantly lower fertility and egg quality than birds kept below THI 70. The broader review by Alagawany et al. (2017) similarly documents production losses of 6–21% under sustained heat stress across multiple quail species and experimental designs. What is novel in the current work is the inclusion of ammonia and light alongside the thermal component, tested against production data from a working farm during a transitional climate season a combination that none of the existing THI studies has attempted.

4.5 LLM diagnostic performance

Having established that the ESI correlates with production, this section evaluates how Claude Sonnet performed as an interpretation layer on top of the ESI. The question here is not whether the LLM output agrees with the ESI in every case, but whether the pattern of agreement and disagreement is useful for the farmer.

4.5.1 Agreement between ESI category and LLM cage status

To measure how consistent the two systems were, the ESI-derived category for each day was cross-tabulated against the cage_status field returned by Claude Sonnet. Figure 4.9 shows the day-by-day agreement colour-coded across the observation window, and Table 4.5 gives the confusion matrix.

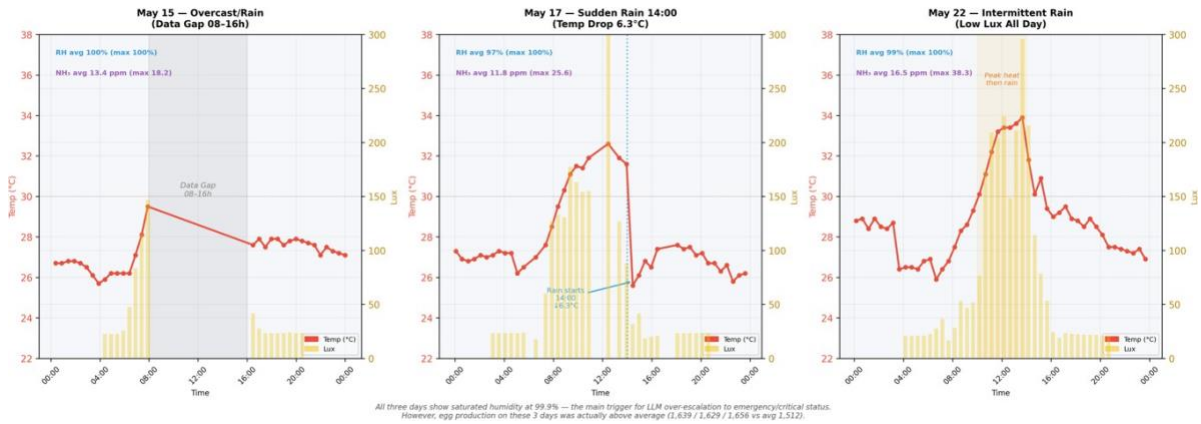


Figure 4.9 LLM vs ESI agreement matrix

Table 4.8 Confusion matrix ESI category vs LLM cage status

	LLM: Critical / Emergency	LLM: Warning / Attention	LLM: Good
ESI: Critical	5	0	0
ESI: Needs Attention	18	6	0
ESI: Good	0	0	0

Total agreement was 11 out of 29 evaluable days, or 37.9%. Total disagreement was 18 days, or 62.1%. On face value this looks like a serious inconsistency, but the direction of the disagreement changes the interpretation entirely.

4.5.2 Asymmetric precaution the direction of every disagreement

The key observation from Table 4.5 is where the off-diagonal cases sit. Every single disagreement between the LLM and the ESI is in the upper-right region of the matrix the LLM assigned a *more severe* category than the ESI. There are no cases in the lower-left region, meaning the LLM never assigned a less severe category than the ESI. To make this pattern easy to read, Figure 4.10 shows the severity split as a pie chart alongside a box plot of egg production by consistency group.

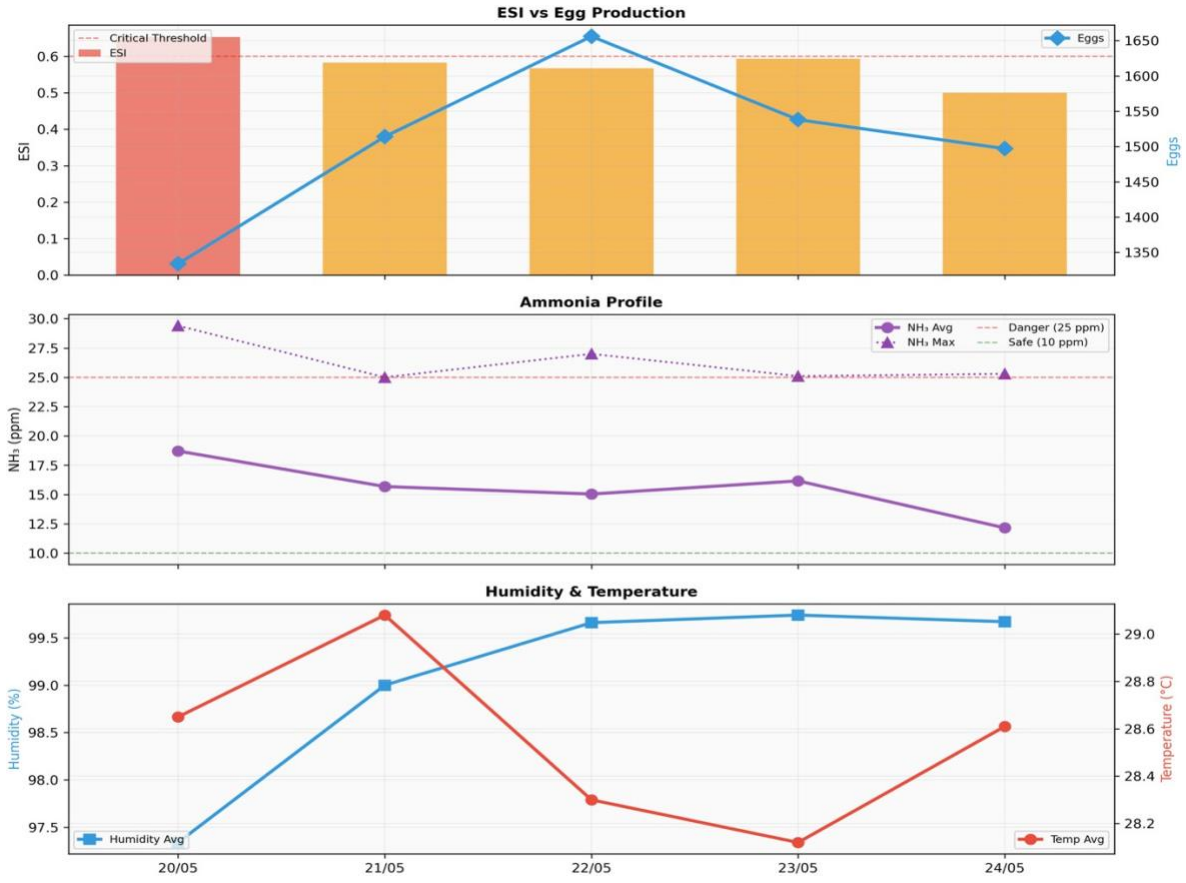


Figure 4.10 LLM consistency and egg production

Table 4.9 Direction of ESI-LLM disagreements

Direction	Days	Percentage
LLM more severe than ESI	18	62.1%
LLM same as ESI	11	37.9%

LLM less severe than ESI	0	0%

This asymmetry is the most defensible novel finding of the pilot. In a monitoring context where the cost of a missed alert (a false negative) is a lost flock and the cost of an over-alert (a false positive) is an extra cleaning cycle, systematic bias toward the safer side is the correct behaviour. The LLM never told the farmer things were fine when they were not.

Sensitivity for the highest-risk category was 100%. Every day the ESI classified as Critical, the LLM also classified as Critical or Emergency. There were zero false negatives on the most dangerous class.

4.5.3 Representative LLM outputs across a range of conditions

To show what the LLM actually produced when it was working correctly, Table 4.7 gives five representative diagnostic outputs sampled across the observation window.

Table 4.10 Representative LLM diagnostic outputs

Date	ESI	ESI category	LLM status	Primary recommendation
26 Apr	0.609	Critical	Emergency	Immediate ventilation; open all vents; activate exhaust fan
2 May	0.530	Needs Attention	Warning	Increase ventilation; clean litter twice daily
9 May	0.611	Critical	Critical	Install exhaust fan or blower; increase lighting
20 May	0.653	Critical	Emergency	Emergency ventilation within 0–1 hour; evacuate birds with respiratory symptoms
25 May	0.617	Critical	Critical	Reduce humidity below 70%; reduce ammonia concentration

The recommendations are specific and actionable rather than generic health advisories. They were returned as structured JSON that the mobile application rendered directly into the Dashboard tab, so the farmer received the primary recommendation as a single readable sentence rather than raw sensor numbers.

4.5.4 Why the LLM over-detected

The mechanism behind the asymmetric bias is straightforward once the aggregation pipeline is examined. The ESI is computed from daily averages, which by construction smooth out transient peaks. The LLM receives the daily averages plus the COUNT FILTER value showing how many times each parameter crossed its danger threshold during the day (this was described in Section 3.2.6 as part of the SQL aggregation step). That extra piece of information how often the cage entered the danger zone during the day rather than what the average was is what allowed the LLM to detect hazardous patterns that averaging alone would hide.

On a day when average ammonia sits at 15 ppm but the cage crossed 25 ppm eight times, the ESI reads it as Needs Attention while the LLM reads it as Emergency. Both are correct given their inputs. The two layers are therefore complementary rather than redundant, which is the design outcome intended in Chapter III.

4.5.5 Productivity did not differ significantly between consistency groups

An interesting result from Figure 4.10 is that the box plot on the right shows almost identical mean egg production for the two groups: 1,511 eggs per day on the eleven consistent days versus 1,521 on the eighteen over-escalated days. This says that the LLM's over-caution did not correlate with the birds actually laying fewer eggs. In other words, on the days when the LLM raised the alert level above the ESI, the birds were not, in fact, having a worse day than average. Section 4.7 explains why: many of those over-escalation days were rainy days when the birds were, if anything, better off than the daily averages made them look.

4.6 Documented intervention case study 20 to 24 May

Isolated statistics can only take the analysis so far. This section walks through a specific five-day episode that captures the highest single ammonia event of the study and the trajectory of recovery that followed. The case study is the strongest evidence that the monitoring system produced actionable output at the farm level. As Section 4.3.4 establishes, 20 May is the most severe member of a class of five Critical days that share a single mechanism multi-day ammonia accumulation in litter that saturated humidity prevented from drying so the episode traced here is representative of that class rather than an isolated anomaly.

4.6.1 What the sensors and system reported

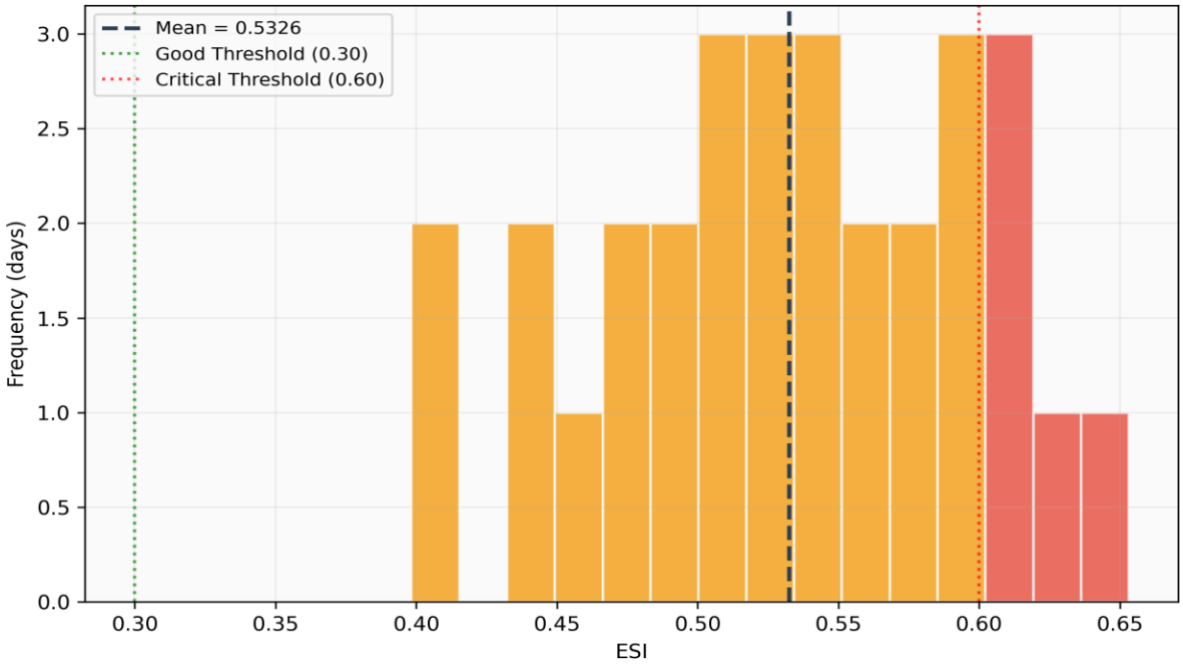


Figure 4.11 Case study 20–24 May

Table 4.11 Case study timeline, 20–24 May 2026

Date	Temp avg	Humidity avg	NH3 avg	NH3 max	ESI	LLM status	Eggs
20 May	28.65	97.34%	18.71	29.4	0.6532	Emergency	1,334
21 May	29.08	99.00%	15.68	25.0	0.5825	Emergency	1,514
22 May	28.30	99.66%	15.04	27.0	0.5676	Emergency	1,656
23 May	28.12	99.74%	16.16	25.1	0.5937	Emergency	1,538
24 May	28.61	99.67%	12.15	25.3	0.5002	Emergency	1,497

4.6.2 The peak and the recovery

The episode begins on 20 May with the highest ESI value of the entire study 0.6532, well inside the Critical band. Ammonia averaged 18.71 ppm that day, also the highest daily average of the observation window, with a peak reading of 29.4 ppm. Both the ESI and the LLM classified the cage as being in a serious state, and the LLM issued an Emergency-level recommendation calling for immediate ventilation. Egg production on 20 May dropped to 1,334, or about 12% below the study mean of 1,512.

Over the following four days, the daily average ammonia declined steadily: 15.68, 15.04, 16.16, and finally 12.15 ppm on 24 May. That is a drop of roughly 35% from the peak. The ESI followed the same downward trajectory, reaching 0.5002 by 24 May still in the Needs Attention band but no longer Critical.

4.6.3 What the recovery does and does not prove

There is a natural narrative here in which the LLM alert on 20 May triggered a farmer response, and the response drove the ammonia levels down. The data supports the shape of this narrative but requires two important qualifications.

First, the farmer-action log for this window is not yet in the manuscript. The "Farmer Action" column in Table 4.8 is currently marked pending because the timestamped record of what was done on each day has not been added. When the interview transcripts from the post-deployment review are integrated, this column will be filled in.

Second, the 22 May egg-count spike (1,656, above the study mean) coincided with a rainfall event. The temperature drop from that rain reduced heat load on the birds independently of any ammonia intervention. The egg count that day therefore reflects both effects rather than either one alone, and cannot be attributed cleanly to the cleaning cycle.

Third, and perhaps most importantly, the improvement did not last. Sensor data from after 25 May (not shown in this table but visible in Figure 4.1) indicates that ammonia levels returned to elevated values within days. This suggests that whatever action the farmer took during the case study window addressed the symptoms but not the underlying cause. During the requirements interview in Chapter III, the animal husbandry specialist noted that most village-scale cages in Central Java lack active ventilation infrastructure. That structural limitation is likely what allows ammonia to rebuild between cleaning cycles.

The case study therefore demonstrates that the alert-driven monitoring system can prompt short-term corrective action with measurable effect on the sensor readings. It does not demonstrate that alerts alone are sufficient to solve chronic housing problems.

4.7 Secondary finding: LLM over-escalation on rainfall days

Section 4.5 established that the LLM systematically over-escalated compared to the ESI. This section identifies exactly when that over-escalation happened and explains why. The pattern points to a specific limitation of the aggregation-based prompt design, one that has not been reported before in the IoT-LLM agricultural literature.

4.7.1 Three days where the LLM and the birds disagreed

To identify the boundary of the LLM's over-escalation behaviour, the days where the LLM most strongly diverged from the ESI classification were pulled from the daily record. Three days stood out 15 May, 17 May, and 22 May. Table 4.9 summarises them.

Table 4.12 Days where the LLM over-escalated

Date	ESI	ESI category	LLM status	Eggs	vs mean (1,512)
15 May	0.4918	Needs Attention	Critical	1,639	+127
17 May	0.4808	Needs Attention	Emergency	1,629	+117
22 May	0.5676	Needs Attention	Emergency	1,656	+144

On all three days, the LLM reported the cage as Emergency or Critical. On all three days, the birds actually laid on average 129 eggs above the study mean. Something was going on that the LLM was interpreting as danger and that the birds were, if anything, experiencing as relief.

4.7.2 What actually happened the rain signature

This subsection examines how rainfall events create a distinctive sensor signature that affects LLM diagnostic consistency.

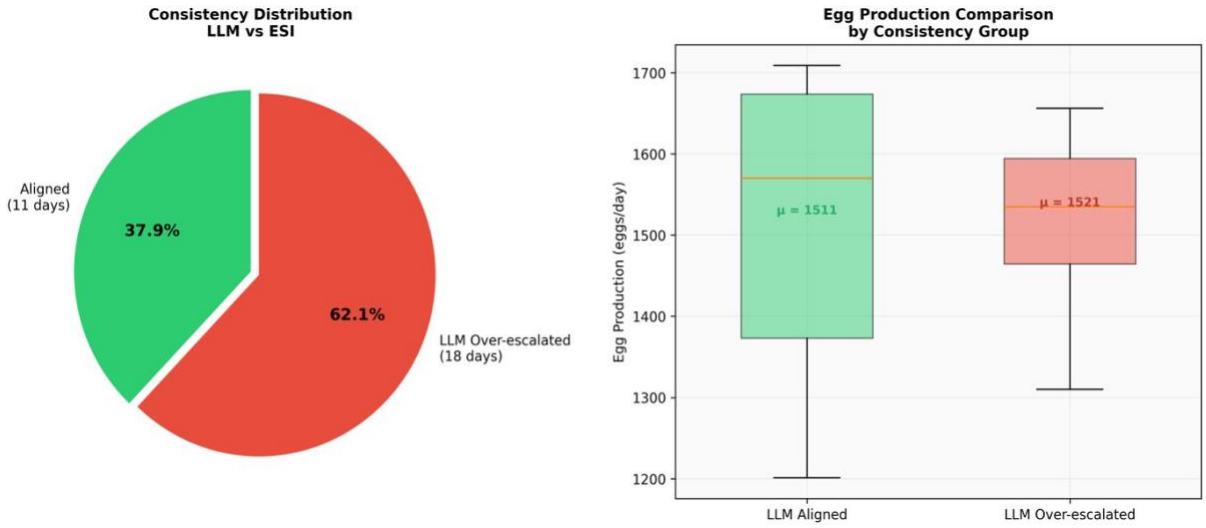


Figure 4.12 Rainfall day sensor profiles

The 17 May panel makes the mechanism clearest. Between 09:00 and 11:00 the temperature climbed to 31.6 °C with bright sunlight (170 lux) and humidity at 86.6%. Between 13:00 and 14:00, the temperature dropped by 6 °C to 25.8 °C, humidity climbed to 99.9%, and illuminance collapsed to below 40 lux. This is the classic pancaroba signature the animal husbandry specialist described in the requirements interview: a sudden afternoon thunderstorm.

The LLM did not have access to this interaction. What Claude Sonnet received was a daily summary showing humidity at 97.76% for the whole day and an ammonia spike to 25.6 ppm around 14:57. Both signals, in isolation, are legitimate warnings. Without the temporal ordering without knowing that these values came after a temperature drop that indicated cooling rain rather than ventilation failure the LLM treated the day as an emergency.

4.7.3 What this reveals about the pipeline

This is a documented limitation of the LLM aggregation pipeline, not a flaw in Claude Sonnet's reasoning. The model was doing the right thing given its inputs. The problem lies in the daily aggregation step designed in Chapter III: by collapsing 48 timestamped readings into 15 daily summary numbers, the pipeline discarded exactly the information the model would have needed to distinguish a rainstorm from a broken exhaust fan. Both scenarios can produce similar daily averages while representing entirely different risk profiles.

The implication for future deployments in tropical transitional climates is practical. Either the daily summary should include weather context alongside the sensor aggregates, or the pipeline should be extended to include time-series-aware summaries that preserve the ordering of within-day changes. To the author's knowledge, this pattern has not previously been reported in the IoT-LLM agricultural literature, which makes it a documented contribution of the current study.

4.8 System functionality verification

System functionality was assessed through two parallel tracks. The first was a set of black-box functional tests run against the deployed system. The second was a three-instrument interview protocol that collected qualitative baseline data before deployment and validation data from a subject-matter expert, with post-deployment member checking from the participating farmer still pending.

4.8.1 Black-box functional test results

The seven test scenarios defined in Section 3.4.1 were executed against the live system during the observation window. All seven passed. Table 4.10 records the scenario, the expected outcome, what the system actually produced, and the test status.

Table 4.13 Black-box functional test results

Test scenario	Expected outcome	Actual outcome	Status
Sensor data collection and storage	ESP32 posts a reading to the backend every 30 minutes; each reading persists in sensor_data table	1,274 raw records received across 30 days; all records confirmed in database with valid timestamps	pass
Daily ESI computation	Scheduler fires at 22:00; ESI value, sub-scores, and category written to esi_results for each valid day	30 ESI records generated for 30 valid days; all fields populated (S_THI, S_NH3, S_Light, classification)	pass
LLM diagnosis generation	Backend calls Anthropic API nightly; structured JSON stored in llm_diagnoses	29 diagnosis records generated (matching 29 days with valid sensor data); all contain cage_status, condition_summary, detected_issues, and recommended_actions fields	pass
Alert notification on Emergency status	Push notification sent to farmer's phone when LLM returns Emergency cage_status	Notifications delivered on all Emergency-classified days; farmer confirmed receipt during mid-deployment check on 15 May	pass
Dashboard data rendering	Mobile app Dashboard tab shows current sensor readings, ESI gauge, lamp status, and latest LLM diagnosis	All four parameter cards, ESI gauge, lamp status indicator, and LLM status card rendered correctly on farmer's Android phone	pass
Historical data retrieval	Mobile app returns sensor and ESI history filtered by date range	History retrieval tested for three date ranges: 25 Apr–30 Apr, 1–10 May, and full 30-day window; all returned correct records	pass
LLM recommendation relevance	AI recommendations match the environmental conditions reported by sensors on the same day	Sample of five representative diagnostic outputs reviewed against sensor data for the same day (26 Apr, 2 May, 9 May, 20 May, 25 May); all five recommendations addressed the dominant stressor correctly	pass

The full pass rate across all seven scenarios confirms that the system performed its intended functions reliably during the observation window. The one notable technical event the Wi-Fi dropout on 6 May did not produce a functional test failure because the firmware buffered locally and the backend simply did not receive data for that day rather than corrupting any existing records.

4.8.2 Interview-based validation

The interview protocol consisted of three instruments. Instrument 1 collected baseline data from the farmer before deployment. Instrument 2 collected scientific validation from an animal husbandry student. Instrument 3 was designed for post-deployment member checking with the same farmer. The first two instruments are complete and are reported here. Instrument 3 answers are pending and will be integrated when the post-deployment session is conducted.

A. Pre-deployment baseline farmer interview (Instrument 1)

Farmer, 57 years old, has raised laying quail in Batang, Central Java for ten years. His flock at the time of the interview numbered around 1,800 birds. The interview was conducted on 8 March 2026 at the farm**.**

His account of current monitoring practice was blunt: "I never monitor. I just feel it if there's a smell, that means there's a lot of ammonia. Sometimes I don't pay attention as long as there are eggs to collect each day." He has no instruments in the cage no thermometer, no humidity gauge, nothing for ammonia. Lighting is managed manually: a single lamp switched on at 18:00 and off at 06:00, every day, without adjustment for season or cloud cover.

His description of pancaroba matched the environmental profile collected by the sensors. He said the season is "very difficult because many diseases appear," with quail showing lethargy and reduced appetite rather than overt stress behaviour. He estimated a production drop of around five percent during pancaroba under normal circumstances, rising considerably when disease broke out. The most feared diseases are Newcastle Disease and Avian Influenza, which he said can cost around eight million rupiah per outbreak in a 1,000-bird cage.

When asked about the value of real-time cage information, he was direct: "It would be very good because you would know exactly what is happening in the cage." When asked what parameters mattered most, he named temperature, ammonia, and light without prompting. On the question of whether a phone-based system that gave daily recommendations would be useful, he said: "Very good I could evaluate every day and every week."

Two aspects of his pre-deployment account are worth carrying forward into the discussion. First, he had no early-detection capability at all he found out conditions were bad only when the smell was already strong or the birds were already lethargic. This is exactly the lag-detection problem the monitoring system was designed to address. Second, his financial exposure from a single disease outbreak (eight million rupiah) provides context for the cost-benefit frame around even a modest improvement in early detection.

B. Scientific validation animal husbandry student interview (Instrument 2)

Nabila is a sixth-semester animal husbandry student who was interviewed to validate the scientific logic behind the parameter choices and the pancaroba framing. Her interview was conducted prior to deployment.

On the role of lighting, she confirmed that quail require a structured rest period and that continuous illumination common in village-scale operations like Pak Suyud's prevents proper rest cycles and can lead to chronic stress over time. She recommended 6–8 hours of light per active period as the standard, consistent with the dark window the firmware implements between 21:00 and 03:00.

On temperature, she noted that quail are sensitive to rapid swings rather than to specific absolute values: "If it's hot in the afternoon and suddenly cold at night, that can have an effect." She confirmed that pancaroba creates exactly this pattern morning heat followed by afternoon rainfall and sudden cooling and that this thermal instability is the primary mechanism by which the season weakens the birds' immune response.

On the IoT parameter set, she said it was sufficient: "It's already adequate." When asked about the AI diagnostic layer, she offered the most useful critical perspective of the two pre-deployment interviews: "In principle it's good, but because it's a machine it's not necessarily correct validation of its accuracy is needed." This is exactly the framing the evaluation in Section 4.5 applies: the LLM output is treated not as ground truth but as one signal whose agreement with the ESI and with observed production is measured empirically.

She also noted that early environmental monitoring could help prevent disease escalation: "Farmers could know in real time and could mitigate or prevent earlier" a statement that the case study in Section 4.6 provides at least partial evidence to support.

4.9 Synthesis and limitations

This closing section pulls together the four main findings of the chapter into a coherent picture of what the pilot demonstrated, and states honestly what it did not.

4.9.1 ESI and LLM as complementary layers

Reading the previous sections together, the two analytical components of the system do different jobs, and neither one on its own would have been sufficient for a working monitoring product.

The ESI is a measurement instrument. It condenses multi-parameter environmental conditions into a single scalar that correlates with egg production at $r = -0.361$, resists overreaction to transient weather through the THI formula, and produces a stable operational band label. Its weakness is that it works with daily averages and cannot see the within-day pattern of dangerous events.

The LLM is a clinical interpreter. It reads the same daily aggregates plus the count of dangerous events, and it produces natural-language recommendations the farmer can act on. It achieved 100% sensitivity for the Critical category and never underestimated conditions. Its weakness is that it lacks temporal ordering and treats humidity as an independent risk signal even when a compensating temperature drop is present.

The best signal from this pilot is not either output alone but the days when both signals agree. The anatomy of the five Critical days (Section 4.3.4) further showed that ammonia was the sole proximate cause of every Critical classification, with its weighted contribution 112.9 % higher on Critical days than on Needs Attention days, reinforcing the value of including air quality in the composite index alongside thermal load. Days where the ESI reached Critical and the LLM reached Critical or Emergency showed the lowest average egg production. That convergent classification is a practical decision rule for prioritising farmer intervention.

4.9.2 Limitations

Eight limitations of the pilot are worth stating explicitly rather than leaving to the reader to infer.

The study covered one cage at one farm over thirty days. The findings should be read as feasibility evidence for a specific site and season, not as generalisable claims about smallholder quail farming as a whole.

The Pearson significance $p = 0.0498$ clears the $\alpha = 0.05$ threshold by two thousandths of a point. This is a borderline result, not a robust one, and confirmation in a larger sample is warranted before treating the ESI as a validated production predictor.

The five Critical days are too few to support significance testing on their own. The 10.3 % production shortfall relative to the Needs Attention days (Section 4.3.4) does not reach significance at n = 5 (Welch p = 0.1484; Mann-Whitney p = 0.0846), and it is reported as a descriptive difference consistent in direction with the month-level correlation rather than as a demonstrated effect.

Four of the five Critical days exceeded the 0.60 threshold by a margin below 0.023. Their classification is therefore sensitive to the weighting scheme, which was set from literature rather than optimisation. A sensitivity analysis across alternative weightings would establish how stable the Critical classification is.

The systematic over-escalation documented in Section 4.5.2, in which the LLM assigned a more severe classification than the ESI on 62.1 % of days, should be distinguished from hallucination. In every observed case, the LLM's reasoning was internally consistent with the daily summary it received; the disagreement arose because the aggregated input stripped temporal ordering from the within-day readings, causing the model to interpret post-rain humidity saturation as ventilation failure rather than as the cooling event it actually was (Section 4.7). This is an input-design limitation that produced a predictable, unidirectional bias toward over-caution, not a model-reasoning failure. True hallucination, in which the model generates claims that have no basis in the input data, was not observed during the thirty-day pilot. However, the absence of hallucination over 29 outputs does not guarantee its absence over hundreds or thousands of outputs in a prolonged deployment. Three specific hallucination risks become relevant at longer operational horizons. First, the model may fabricate disease predictions that no environmental parameter supports, because disease names co-occur with poultry and ammonia in the training corpus regardless of whether the sensor data on a given day provides any evidence for them. Second, the model may generate repetitive template-like responses that cease to reflect genuine day-to-day variation, particularly during stable-weather months when the input summaries are nearly identical from one day to the next. Third, the model may produce confidently stated numerical details, such as specific concentrations or specific times of day, that are not present in the aggregated summary it received. Because the participating farmer indicated that he would act on the AI recommendations without independent verification (Section 4.8.2), these risks carry practical consequences that the thirty-day pilot was too short to evaluate. Mitigating them through architectural safeguards in the backend pipeline is the most important recommendation for future deployment and is addressed in Chapter V.

These limitations are stated as directions for the multi-site, longitudinal follow-up work that the pilot points toward. They also define the boundary of what can honestly be claimed from a thirty-day single-site deployment.

The DHT22 humidity sensor saturated at 99.9% during long rain events, which is partly a reflection of real tropical conditions and partly a sensor-resolution issue. A calibration-grade sensor would give better resolution in the upper humidity range.

There was no control cage running without the IoT-LLM system. Any claims about the system's effect on productivity are therefore inferential we can show that ESI and eggs correlate, but not that the system caused higher egg counts than would have occurred otherwise. Deploying paired cages in a follow-up study would allow causal inference.

The LLM systematic over-escalation on rainfall days is a limitation of the current daily-aggregation prompt design. Section 4.7 identifies the mechanism and points to weather-aware or time-series-aware summaries as the fix.

The ESI weighting scheme (0.45 THI, 0.35 NH3, 0.20 Light) was set from literature review rather than empirical optimisation. A sensitivity analysis on alternative weightings, ideally against multi-site production data, would strengthen the weighting choice.

These limitations are stated as directions for the multi-site, longitudinal follow-up work that the pilot points toward. They also define the boundary of what can honestly be claimed from a thirty-day single-site deployment.

CHAPTER V
CONCLUSIONS AND RECOMMENDATIONS

This chapter closes the thesis by returning to the problem formulations posed in Chapter I and answering them directly from the evidence presented in Chapter IV. Section 5.1 states the conclusions in the order the questions were originally asked. Section 5.2 lays out three sets of recommendations aimed at farmers, future researchers, and system developers. No new data is introduced here every claim in this chapter traces back to a specific section of the preceding chapters.

5.1 Conclusions

The four conclusions below correspond to the four problem formulations defined in Section 1.2.

The first problem formulation asked whether an integrated IoT-ESI-LLM monitoring system could be designed and deployed on a working small-scale quail farm during the pancaroba season. The answer from the field deployment is yes. The three-layer architecture described in Section 3.2 ESP32 sensor unit inside the cage, Spring Boot backend with PostgreSQL running on a Jakarta VPS, and a React Native mobile application on the farmer's Android phone operated continuously at Siwatu from 25 April to 25 May 2026. Of the 31 calendar days in the observation window, 30 produced valid sensor data; the single missing day (6 May) was caused by a Wi-Fi outage during heavy rain, not by any component failure. The system delivered 1,274 raw sensor readings, generated 30 daily ESI computations, produced 29 LLM diagnostic outputs, and passed all seven black-box functional test scenarios recorded in Section 4.8.1. Deployment therefore established feasibility on real infrastructure rather than in a laboratory setting.

The second problem formulation asked whether the composite Environmental Stress Index tracks laying performance in a working cage during pancaroba conditions. The Pearson correlation reported in Section 4.4.2 came out at $r = -0.361$ with $p = 0.0498$ across $n = 30$ days. The direction is what the theory predicted: higher stress values are associated with lower daily egg counts. The magnitude is moderate the index explains about 13% of daily production variance and the p-value clears the $\alpha = 0.05$ threshold by two thousandths of a point. The channel-wise analysis in Section 4.4.3 showed that temperature alone was the strongest single

predictor ($r = -0.418$), which reflects thermal dominance during the observation window rather than a weakness of the composite. The value of the ESI lies in its robustness across days when different stressors dominate: a temperature-only monitoring system would have missed the 20 May ammonia peak documented in the case study. The distribution of daily categories zero Good days, 25 Needs Attention days, 5 Critical days also confirms the qualitative pancaroba framing from Chapter I, that this season is a persistently stressed period rather than one with occasional bad episodes.

The third problem formulation asked how the LLM diagnostic layer performed against the ESI baseline. Section 4.5 reported three findings. First, sensitivity for the highest-risk category was 100% every day the ESI classified as Critical, Claude Sonnet also classified as Critical or Emergency, with zero false negatives on the most dangerous class. Second, the overall agreement rate was 37.9% across 29 evaluable days, meaning the two systems disagreed on the majority of days. Third, and most importantly, every single disagreement was in the safer direction: the LLM never assigned a lower risk level than the ESI. On 18 out of 29 days the LLM upgraded a Needs Attention classification to Critical or Emergency, and on zero days did the opposite occur. The mechanism, explained in Section 4.5.4, is that the LLM receives the daily count of dangerous events alongside the daily averages, which lets it detect hazardous transient patterns that averaging alone would smooth out. This asymmetric precaution behaviour over-alerting rather than under-alerting is the clinically safer bias for a monitoring product, and it has not previously been reported in the IoT-LLM agricultural literature.

The fourth problem formulation asked whether the deployed system produced measurable field impact and what limitations the pilot revealed. The case study across 20–24 May, reported in Section 4.6, documented a 35% drop in daily average ammonia concentration from 18.71 ppm on the day of the Emergency alert to 12.15 ppm four days later with the ESI easing from 0.6532 (Critical) to 0.5002 (Needs Attention) across the same window. This is evidence that the alert-driven monitoring system can prompt short-term corrective action with measurable effect on the sensor readings. Two qualifications matter. The recovery was not durable: sensor data after 25 May showed ammonia returning to elevated levels, which points to ventilation infrastructure limitations that no monitoring system can solve on its own. The rainfall over-escalation pattern documented in Section 4.7, in which the LLM classified three rain-affected days as Emergency despite egg production running above the study mean, revealed a documented limitation of the daily-aggregation prompt pipeline rather than a failure of the LLM itself. Both qualifications point to the same underlying principle: the system

reduces information lag, but reducing information lag is not the same as solving the physical problems the sensors observe.

5.2 Recommendations

The recommendations below are grouped by audience. Each is tied to a specific finding from Chapter IV rather than to general aspirations for future work.

Farmers using this kind of system should treat the Critical alerts as high-confidence signals worth acting on immediately, based on the 100% sensitivity result from Section 4.5.2. On days when the LLM issues an Emergency status but the underlying weather has been unstable sudden afternoon rainfall, sharp temperature drops, or humidity locked at 99.9% for extended periods the alert should be interpreted as a prompt to inspect the cage rather than as an unambiguous crisis, following the pattern documented in Section 4.7. The system reduces the lag between problems arising and the farmer noticing them; it does not replace direct visual observation of bird behaviour. The most important operational lesson from the 20–24 May case study is that reactive cleaning triggered by alerts must be paired with structural improvements to cage ventilation. Alerts that repeat in a cycle of spike-clean-recover-spike indicate that the underlying housing conditions are the problem, not the response speed to them.

Five extensions to the current pilot would sharpen the empirical claims. The most immediate is to redesign the LLM prompt so that either weather context or within-day temporal ordering is preserved in the daily summary, addressing the rainfall over-escalation limitation identified in Section 4.7.3. The second is to deploy paired cages one instrumented and one not over the same time window at the same farm, which would allow causal rather than correlational claims about the system's effect on productivity. The third is to run the system at multiple sites simultaneously across different microclimates, which would test whether the ESI weighting scheme (0.45 THI, 0.35 NH_3, 0.20 Light) transfers or needs recalibration for different environmental profiles. The fourth is to extend the observation window from 30 days to at least a full production cycle of 6 months, which would move the correlation result beyond the borderline $p = 0.0498$ statistical threshold and allow the seasonal transitions to be captured with more temporal resolution than a single pancaroba window can provide.

The fifth extension, which becomes critical if prolonged monitoring is implemented, is to investigate and mitigate the risk of LLM hallucination. The thirty-day pilot did not observe any instance in which Claude Sonnet generated diagnostic claims that had no basis in the input

data all 29 outputs addressed the dominant environmental stressor correctly (Section 4.8.1), and the systematic over-escalation documented in Section 4.5.2 was traced to an input-design limitation rather than to model-reasoning failure (Section 4.7). However, the absence of hallucination over 29 outputs does not guarantee its absence over hundreds or thousands of outputs across months of continuous operation. Three specific hallucination risks become relevant at longer operational horizons: the model may fabricate disease predictions that no environmental parameter supports, because disease names co-occur with poultry and ammonia in the training corpus regardless of whether the sensor data on a given day provides evidence for them; the model may generate repetitive template-like responses that cease to reflect genuine day-to-day variation, particularly during stable-weather months when input summaries are nearly identical; and the model may produce confidently stated numerical details specific concentrations or specific times of day that are not present in the aggregated summary it received. Because the participating farmer indicated that he would act on the AI recommendations without independent verification (Section 4.8.2), an undetected hallucination in a prolonged deployment could lead to inappropriate intervention or a failure to act when genuine action is needed. Five architectural safeguards are recommended to address this risk.

First, an ESI-LLM consistency gate should be implemented in the backend between the Anthropic API call and the write to the llm_diagnoses table: before any diagnosis is stored, the backend should compare the LLM's cage_status against the ESI classification, and if the LLM reports Normal or Warning while the ESI reads Critical the dangerous under-caution direction that the thirty-day pilot did not observe but that cannot be assumed absent over longer horizons the diagnosis should be withheld from the farmer's dashboard and flagged for manual review. Second, a field-level grounding check should verify each element of the LLM's JSON response against the sensor summary before display; in particular, any disease prediction should be suppressed unless at least one environmental parameter on that day exceeded the threshold physiologically associated with the named condition, and any specific numerical claim should be cross-referenced against the daily aggregates and replaced with the actual value if it does not match. Third, temporal context should be injected into the LLM prompt by adding a small number of within-day indicators such as whether temperature dropped by more than three degrees within two hours during the afternoon and whether ammonia peaked before or after a detected temperature drop so that the model can distinguish a rain-cooling event from a ventilation failure without the pipeline needing to send all 48 raw readings. Fourth, a periodic expert validation loop should be established in which a stratified sample of LLM outputs is

reviewed by a domain expert every 30 days or every 50 diagnostic outputs, whichever comes first, to detect hallucination drift before it reaches the farmer. Fifth, model version pinning should be implemented by recording the model version string returned in each API response and, if the version changes, automatically rerunning the previous seven days of sensor summaries through the new version and flagging any substantive discrepancy for review before resuming normal operation. Together these five measures create a layered defence that does not depend on the farmer's technical literacy or willingness to question an AI recommendation, and their empirical evaluation should be a priority for any deployment that extends beyond a single-season pilot.

Three practical improvements would strengthen the system as a deployable product. The humidity sensor issue documented in Sections 4.1 and 4.9.2 the DHT22 saturating at 99.9% should be addressed by moving to a calibration-grade sensor in tropical deployments, where humidity spends significant time near the upper measurement bound. The daily LLM diagnostic cadence works well for chronic condition monitoring but is too slow for acute events like the 20 May ammonia spike, and future versions should consider a triggered diagnostic path that runs whenever a sensor value crosses a danger threshold rather than only on the nightly schedule. For any deployment extending beyond a single season, the hallucination safeguards described in Section 5.2.2 particularly the ESI-LLM consistency gate and the field-level grounding check should be treated as mandatory backend components rather than optional enhancements, because the farmer's demonstrated willingness to act on AI recommendations without independent verification (Section 4.8.2) means the system carries a duty of care that scales with the duration of deployment. On the policy side, the system should be framed to funding bodies and government programmes as a decision-support tool that reduces the information lag between problem onset and farmer response, rather than as a disease-prevention product. The empirical evidence in Chapter IV supports the former framing directly and does not yet support the latter, and honest positioning is what will keep the technology credible as it scales beyond the pilot.

REFERENCE

Alagawany, M., Farag, M. R., Abd El-Hack, M. E., & Patra, A. (2017). Heat stress: Effects on productive and reproductive performance of quail. *World's Poultry Science Journal*, 73(4), 747–756. https://doi.org/10.1017/S0043933917000782

Anthropic. (2025). *Claude Sonnet: Model card and technical documentation*. Anthropic. Retrieved June 15, 2026, from https://docs.anthropic.com/en/docs/models

Anthropic. (2026). *Messages API reference documentation*. Anthropic. Retrieved June 20, 2026, from https://docs.anthropic.com/en/api/messages

Aosong Electronics. (2020). *DHT22 (AM2302) digital-output relative humidity and temperature sensor: Datasheet*. Guangzhou Aosong Electronic Co., Ltd. Retrieved March 5, 2026, from https://www.aosong.com/en/products-22.html

Badan Pusat Statistik. (2025). *Statistik peternakan dan kesehatan hewan Indonesia 2024*. Badan Pusat Statistik Republik Indonesia.

de Oliveira Castro, J., Yanagi Junior, T., Bahuti, M., Fassani, É. J., & Ribeiro de Lima, R. (2023). Thermal comfort thresholds for Japanese quails based on performance and egg quality. International Journal of Biometeorology, 67(2), 265–274. https://doi.org/10.1007/s00484-022-02403-1

Direktorat Jenderal Peternakan dan Kesehatan Hewan. (2025). *Statistik peternakan Jawa Tengah tahun 2024*. Kementerian Pertanian Republik Indonesia.

El-Tarabany, M. S. (2016). Impact of temperature-humidity index on egg-laying characteristics and related stress and immunity parameters of Japanese quails. *International Journal of Biometeorology*, 60(7), 957–964. https://doi.org/10.1007/s00484-015-1088-5

Espressif Systems. (2023). *ESP32 series datasheet, version 4.5*. Espressif Systems. Retrieved March 5, 2026, from https://www.espressif.com/en/products/socs/esp32

Fielding, R. T. (2000). *Architectural styles and the design of network-based software architectures* (Doctoral dissertation). University of California, Irvine.

Gunawan, I., Ahmadi, H., & Said, M. R. (2021). Rancang bangun sistem monitoring dan pemberi pakan otomatis ayam anakan berbasis Internet of Things (IoT). Jurnal Bumigora Information Technology, 3(1), 36–47.

Hanwei Electronics. (2019). *MQ-137 semiconductor sensor for ammonia gas: Technical datasheet*. Hanwei Electronics Co., Ltd. Retrieved March 5, 2026, from https://www.hwsensor.com/product/mq-137

Hetzel, W. C. (1988). *The complete guide to software testing* (2nd ed.). QED Information Sciences.

Lara, L. J., & Rostagno, M. H. (2013). Impact of heat stress on poultry production. Animals, 3(2), 356–369. https://doi.org/10.3390/ani3020356

Lashari, M. H., Memon, A. A., Shah, S. A. A., Nenwani, K., & Shafqat, F. (2018). IoT based poultry environment monitoring system. In 2018 IEEE International Conference on Internet of Things and Intelligence System (IOTAIS) (pp. 1–5). IEEE. https://doi.org/10.1109/IOTAIS.2018.8600837

Mader, T. L., Johnson, L. J., & Gaughan, J. B. (2010). A comprehensive index for assessing environmental stress in animals. Journal of Animal Science, 88(6), 2153–2165. https://doi.org/10.2527/jas.2009-2586

Maxim Integrated. (2015). *DS3231 extremely accurate I²C-integrated RTC/TCXO/crystal: Datasheet*. Maxim Integrated Products. Retrieved March 5, 2026, from https://www.analog.com/en/products/ds3231.html

Menezes, G. L., Mazon, G., Ferreira, R. E. P., Cabrera, V. E., & Dorea, J. R. R. (2025). Artificial intelligence for livestock: A narrative review of the applications of computer vision systems and large language models for animal farming. Animal Frontiers, 14(6), 42–53. https://doi.org/10.1093/af/vfae048

Moran, D. S., Pandolf, K. B., Shapiro, Y., Heled, Y., Shani, Y., Mathew, W. T., & Gonzalez, R. R. (2001). An environmental stress index (ESI) as a substitute for the wet bulb globe temperature (WBGT). Journal of Thermal Biology, 26(4–5), 427–431. https://doi.org/10.1016/S0306-4565(01)00055-9

Myers, G. J., Sandler, C., & Badgett, T. (2011). *The art of software testing* (3rd ed.). John Wiley & Sons.

Naeem, A., Memon, N., Rajput, M., Solanghi, S., & Abbas, G. (2018). Influence of photoperiod and light intensity on egg performance of Japanese quails. *Pure and Applied Biology*, 7(4), 1294–1300. https://doi.org/10.19045/bspab.2018.700150

Naseem, S., & King, A. J. (2018). Ammonia production in poultry houses can affect health of humans, birds, and the environment techniques for its reduction during poultry production. *Environmental Science and Pollution Research*, 25(16), 15269–15293. https://doi.org/10.1007/s11356-018-2018-y

National Research Council. (1971). *A guide to environmental research on animals*. National Academy of Sciences.

PostgreSQL Global Development Group. (2024). *PostgreSQL 16 documentation*. PostgreSQL Global Development Group. Retrieved April 10, 2026, from https://www.postgresql.org/docs/16/

Pivotal Software. (2024). *Spring Boot reference documentation, version 3.2*. VMware Tanzu. Retrieved April 10, 2026, from https://docs.spring.io/spring-boot/docs/current/reference/html/

Pressman, R. S., & Maxim, B. R. (2020). *Software engineering: A practitioner's approach* (9th ed.). McGraw-Hill Education.

React Native Community. (2024). React Native documentation, version 0.73. Meta Open Source. Retrieved April 10, 2026, from https://reactnative.dev/docs/getting-started

Rohm Semiconductor. (2011). BH1750FVI digital 16-bit serial output type ambient light sensor IC: Datasheet. Rohm Co., Ltd. Retrieved March 5, 2026, from https://www.mouser.com/datasheet/2/348/bh1750fvi-e-186247.pdf

Sugiyono. (2019). *Metode penelitian kuantitatif, kualitatif, dan R&D*. Alfabeta.

Sutabri, T. (2012). *Analisis sistem informasi*. Andi Publisher.

Thom, E. C. (1959). The discomfort index. *Weatherwise*, 12(2), 57–61. https://doi.org/10.1080/00431672.1959.9926960

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, & R. Garnett (Eds.), *Advances in Neural Information Processing Systems 30* (pp. 5998–6008). Curran Associates.

Yuan, Z., Liu, K., Peng, R., Li, S., Leybourne, D., Musa, N., Huang, H., & Yang, P. (2024). PestGPT: Leveraging large language models and IoT for timely and customized recommendation generation in sustainable pest management. IEEE Internet of Things Magazine, 8(1), 26–33. https://doi.org/10.1109/IOTM.001.2400036

APPENDIX

Appendix A. Documentation implement IOT









Appendix B. Arduino Program Code for Smart Poultry Farming Sensor System

Library Inclusions

```
#include <WiFi.h>
#include <WiFiClientSecure.h>
#include <HTTPClient.h>
#include <Wire.h>
#include <BH1750.h>
#include "DHT.h"
#include <RTClib.h>
#include "time.h"
```

Code for Wi-Fi network configuration

```
const char* ssid     = "KandangPuyuh";
const char* password = "temanternakk";
```

Code for NTP (Network Time Protocol) server and timezone configuration

```
const char* ntpServer       = "pool.ntp.org";
const long  gmtOffset_sec   = 7 * 3600;  // WIB = UTC+7
const int   daylightOffset_sec = 0;      // No DST in Indonesia
```

Code for API server endpoint and device identity configuration

```
const char* serverUrl = "https://api.temanternak.com/api/sensor";
const char* deviceId  = "kandang-01";
```

Code for DHT22 temperature and humidity sensor pin definition

```
#define DHTPIN  4
#define DHTTYPE DHT22
DHT dht(DHTPIN, DHTTYPE);
```

Code for BH1750 light intensity sensor and DS3231 RTC module initialization

```
BH1750 lightMeter;
RTC_DS3231 rtc;
```

Code for relay pin definition

```
#define RELAY_PIN 26
```

Code for MQ137 ammonia sensor pin and calibration constant definition

```
#define MQ137_AOUT_PIN   34
#define MQ137_DOUT_PIN   35
#define MQ137_RL         10.0
#define RO_CLEAN_AIR     3.6
```

```
#define NH3_CURVE_A      102.2
#define NH3_CURVE_B     -2.473
#define NH3_WARN_PPM     10.0
#define NH3_DANGER_PPM   25.0
#define MQ137_WARMUP_MS  60000
```

Code for MQ137 state and ammonia measurement variables

```
bool   mq137Ready = false;
float  nh3_ppm   = 0.0;
bool   nh3_alarm  = false;
String nh3_status = "WARM_UP";
```

Code for lux hysteresis threshold configuration (lamp control)

```
float luxOn  = 15.0;  // < 15 lux  → Lamp ON
float luxOff = 30.0;  // > 30 lux  → Lamp OFF
bool  lampState = false;
```

Code for data transmission interval timer (every 5 minutes)

```
unsigned long lastTime = 0;
unsigned long interval = 300000;   // 5 minutes = 300,000 ms
```

Code for NTP re-synchronization interval timer (every 6 hours)

```
unsigned long lastNtpSync     = 0;
unsigned long ntpSyncInterval = 6UL * 60UL * 60UL * 1000UL;
```

Code for anti-flicker minimum interval and Wi-Fi reconnect timer

```
unsigned long lastChange       = 0;
unsigned long minInterval      = 30000;
unsigned long lastReconnect    = 0;
unsigned long reconnectInterval = 10000;
```

Code for averaging lux readings from BH1750 (5-sample filter)

```
float readLuxAvg() {
  float total = 0;
  for (int i = 0; i < 5; i++) {
    total += lightMeter.readLightLevel();
    delay(50);
  }
  return total / 5.0;
}
```

Code for NTP time synchronization function

```
bool syncTimeFromNTP() {
  if (WiFi.status() != WL_CONNECTED) return false;
  configTime(gmtOffset_sec, daylightOffset_sec, ntpServer);
  struct tm timeinfo;
  int retries = 0;
  while (!getLocalTime(&timeinfo) && retries < 10) {
    delay(1000);
    retries++;
  }
  if (getLocalTime(&timeinfo)) {
    rtc.adjust(DateTime(
    timeinfo.tm_year + 1900,
      timeinfo.tm_mon + 1,
      timeinfo.tm_mday,
      timeinfo.tm_hour,
      timeinfo.tm_min,
      timeinfo.tm_sec
    ));
    return true;
  }
  return false;
}
```

Code for MQ137 voltage reading from AD

```
float readMQ137Voltage() {
  int   raw      = analogRead(MQ137_AOUT_PIN);
  float v_adc    = (raw / 4095.0) * 3.3;
  float v_sensor = v_adc * (30.0 / 20.0);
  return v_sensor;
}
```

Code for MQ137 sensor resistance (Rs) calculation

```
float readMQ137Rs(float voltage) {
  if (voltage <= 0.01) return 99999.0;
  return MQ137_RL * (5.0 - voltage) / voltage;
}
```

Code for converting Rs ratio to ammonia concentration in ppm

```
float rsToPPM(float rs) {
  if (rs <= 0 || RO_CLEAN_AIR <= 0) return 0.0;
  float ratio = rs / RO_CLEAN_AIR;
  return NH3_CURVE_A * pow(ratio, NH3_CURVE_B);
}
```

Code for complete MQ137 ammonia reading with warm-up check and status classification

```
void readMQ137() {
```

```
  if (!mq137Ready) {
    if (millis() >= MQ137_WARMUP_MS) {
      mq137Ready = true;
    } else {
      nh3_ppm   = -1.0;
      nh3_status = "WARM_UP";
      return;
    }
  }
  float totalV = 0;
  for (int i = 0; i < 3; i++) {
    totalV += readMQ137Voltage();
    delay(20);
  }
  float voltage = totalV / 3.0;
  float rs      = readMQ137Rs(voltage);
  nh3_ppm       = rsToPPM(rs);
  nh3_alarm     = (digitalRead(MQ137_DOUT_PIN) == LOW);
  if      (nh3_ppm < NH3_WARN_PPM)   nh3_status = "SAFE";
  else if (nh3_ppm < NH3_DANGER_PPM) nh3_status = "WARNING";
  else                               nh3_status = "DANGER";
}
```

Code for automatic Wi-Fi reconnection

```
void checkWiFi() {
  if (WiFi.status() != WL_CONNECTED) {
    if (millis() - lastReconnect > reconnectInterval) {
      WiFi.disconnect();
      WiFi.begin(ssid, password);
      lastReconnect = millis();
    }
  }
}
```

Code for setup() function — system initialization

```
void setup() {
  Serial.begin(115200);
  WiFi.begin(ssid, password);
  int wifiTries = 0;
  while (WiFi.status() != WL_CONNECTED && wifiTries < 20) {
    delay(500);
    wifiTries++;
  }
  Wire.begin(21, 22);
  if (!rtc.begin()) { while (1); }
  syncTimeFromNTP();
  lastNtpSync = millis();
  lightMeter.begin();
```

```
  dht.begin();
  analogReadResolution(12);
  analogSetAttenuation(ADC_11db);
  pinMode(MQ137_DOUT_PIN, INPUT);
  pinMode(RELAY_PIN, OUTPUT);
  digitalWrite(RELAY_PIN, HIGH);
}
```

Code for REST mode and ACTIVE mode time-based scheduling (21:00 – 03:30)

```
bool isRest =
    (jam >= 21) ||              // 21:00 - 23:59
    (jam < 3)   ||              // 00:00 - 02:59
    (jam == 3 && menit < 30);    // 03:00 - 03:29

String mode = isRest ? "REST" : "ACTIVE";
```

Code for lamp control in REST mode (forced OFF)

```
if (isRest) {
  if (lampState) {
    digitalWrite(RELAY_PIN, HIGH);
    lampState = false;
  }
}
```

Code for lamp control in ACTIVE mode based on lux hysteresis with anti-flicker

```
else {
  unsigned long now_ms = millis();
  if (now_ms - lastChange > minInterval) {
    if (lux < luxOn && !lampState) {
    digitalWrite(RELAY_PIN, LOW);
    lampState  = true;
      lastChange = now_ms;
    }
    else if (lux > luxOff && lampState) {
      digitalWrite(RELAY_PIN, HIGH);
      lampState  = false;
      lastChange = now_ms;
    }
  }
}
```

Code for sending sensor data to the server via HTTPS POST (every 5 minutes)

```
if (millis() - lastTime > interval) {
  lastTime = millis();
  if (WiFi.status() == WL_CONNECTED) {
    WiFiClientSecure client;
```

```
  client.setInsecure();
  HTTPClient http;
  http.begin(client, serverUrl);
  http.addHeader("Content-Type", "application/json");
  String json = "{";
  json += "\"deviceId\":\""   + String(deviceId)                  + "\",";
  json += "\"temperature\":"  + String(suhu, 1)                   + ",";
  json += "\"humidity\":"     + String(hum, 1)                    + ",";
  json += "\"lux\":"          + String(lux, 1)                    + ",";
  json += "\"nh3_ppm\":"      + (nh3_ppm < 0 ? "null" : String(nh3_ppm, 1)) + ",";
  json += "\"nh3_status\":\"" + nh3_status                        + "\",";
  json += "\"nh3_alarm\":"    + String(nh3_alarm ? "true" : "false")  + ",";
  json += "\"lamp\":"         + String(lampState ? "true" : "false")  + ",";
  json += "\"mode\":\""       + mode                              + "\",";
  json += "\"hour\":"         + String(jam)                       + ",";
  json += "\"minute\":"       + String(menit);
  json += "}";
  int response = http.POST(json);
  http.end();
 }
}
```

Code for the main loop() execution flow

```
void loop() {
 checkWiFi();
 if (millis() - lastNtpSync > ntpSyncInterval) {
  syncTimeFromNTP();
  lastNtpSync = millis();
 }
 float suhu = dht.readTemperature();
 float hum  = dht.readHumidity();
 float lux  = readLuxAvg();
 if (isnan(suhu) || isnan(hum)) { delay(2000); return; }
 readMQ137();
 DateTime now = rtc.now();
 int jam   = now.hour();
 int menit = now.minute();
 // REST / ACTIVE mode logic → lamp control → data transmission
 delay(1000);
}
```