



DOCTORAL SCHOOL
MEDITERRANEA UNIVERSITY OF REGGIO CALABRIA
DEPARTMENT OF INFORMATION ENGINEERING, INFRASTRUCTURES
AND SUSTAINABLE ENERGY
(DIIES)

PHD IN
INFORMATION ENGINEERING

S.S.D. CHEM-07/B, ING-INF/01-Elettronica
XXXVIII CYCLE

Agri-Food Chain Data Collection for Advanced Traceability and Certification

CANDIDATE

Mohamed Riad Sebti

ADVISOR

Prof. Mariateresa Russo

CO-ADVISOR

Prof. Massimo Merenda

COORDINATOR

Prof. Tommaso Isernia

REGGIO CALABRIA, JANUARY 2026

The PhD Board of the PhD course in
INFORMATION ENGINEERING
is composed of:

ARANITI Giuseppe
BUCCAFURRI Francesco Antonio
CAMPOLO Claudia
CAROTENUTO Riccardo
COPPOLA Giuseppe
COTRONEI Mariantonia
CROCCO Lorenzo
DALLET Dominique
DE CAPUA Claudio
DELLA CORTE Francesco Giuseppe
DI DONATO Loreto
FAGGIO Giuliana
FAILLA Gioia
GIUFFRÈ Sofia
GROZA Voicu
ISERNIA Tommaso (coordinator)
LAX Gianluca
LAY-EKUAKILLE Aimè
Lohan Elena Simona
MANZONI Pietro
MESSINA Giacomo
MISTRETTA Marina
MOLINARO Antonella
MORABITO Andrea Francesco
MORABITO Francesco Carlo
MORELLO Rosario
MUNTEAN Gabriel-Miro
MUSOLINO Giuseppe
PEZZIMENTI Fortunato
PRATICO' Filippo Gianmaria
RAO Sandro
ROMANO Maria
ROSACI Domenico
RUGGERI Giuseppe
RUSSO Mariateresa
VINEL Alexey

Acknowledgements

I would like to express my deepest gratitude to my parents, my brothers, and my family for their unconditional love, constant encouragement, and unwavering support throughout my academic journey. Their trust, patience, and sacrifices have always been a source of strength and motivation for me. I would also like to sincerely thank my supervisors, Prof. Massimo Merenda and Prof. Mariateresa Russo, for their guidance, support, and valuable advice during the development of this PhD work. Their experience, encouragement, and academic insight have played a fundamental role in shaping this thesis. My sincere thanks also go to my friends, whose support, understanding, and encouragement have accompanied me during both the challenging and rewarding moments of this journey. I am grateful to the Electronics Lab, the FocusLab members, and all the lab colleagues for providing a stimulating and collaborative research environment. Sharing ideas, experiences, and daily work with them has been both valuable and inspiring. I would also like to thank the DIIES Department and all its members for their support and for offering an academic environment that contributed to my professional and personal growth. Special thanks are also extended to the Tech4You project for the opportunities and support provided during my research activities. Finally, I would like to express my gratitude to the University of Reggio Calabria for giving me the opportunity to pursue this PhD path and for supporting my academic development.

Abstract

Modern agri-food supply chains require reliable and real-time data collection mechanisms to enable advanced traceability and certification, ensuring product authenticity, safety, and regulatory compliance. The growing deployment of IoT sensors, identification technologies, and embedded devices across the supply chain has significantly increased the availability of monitoring data. However, this distributed sensing paradigm also generates large volumes of heterogeneous data that exceed the processing, storage, and communication capabilities of low-cost, resource-constrained nodes. As a result, conventional cloud-centric approaches introduce latency, bandwidth overhead, and scalability limitations, hindering continuous and trustworthy traceability. This thesis addresses these challenges by proposing a framework for real-time data collection for advanced traceability and certification, specifically designed for operation at the edge of the agri-food chain. The framework enables data acquisition, preprocessing, and management directly on embedded IoT devices, reducing dependence on centralized infrastructures. To cope with the large data volume problem, on-device data reduction and dataset distillation strategies are introduced to compact raw measurements into representative, traceability-preserving information. This approach minimizes redundancy, lowers memory usage and communication costs, and enables continuous operation on microcontroller-based platforms while maintaining the integrity and usefulness of collected data. Building upon efficient edge-level data collection, the proposed system integrates secure identification and verification mechanisms, combining RFID-based product tagging, local intelligence, cloud services, and a blockchain-inspired integrity layer to create a robust binding between physical products and their digital records. This integration strengthens both traceability and anti-counterfeiting capabilities across the supply chain. The framework is validated through two representative study cases that demonstrate its feasibility and effectiveness in realistic agri-food environments. Overall, this work shows that real-time, edge-based data

collection and processing constitute key enablers for scalable, low-cost, and trust-worthy traceability and certification systems in modern agri-food ecosystems

Contents

Acknowledgements	I
Abstract	III
Conferences	3
Journals	4
1 Introduction	5
1.1 Background and Context	5
1.2 Thesis Motivation and Contributions	7
1.3 Thesis Organization	8
2 State-of-the-Art	11
2.1 Introduction	11
2.2 Traceability and Anti-Counterfeiting in Agri-Food	12
2.2.1 Evolution of Traceability Over Years and Across Borders in Agri-Food Supply Chains	13
2.2.2 Challenges in Modern Agri-Food Traceability and Anti-Counterfeiting Solutions	16
2.3 Enabling Technologies for Next-Generation Traceability	18
2.3.1 Barcoding and QR Code-Based Systems	19
2.3.2 RFID and NFC	21
2.3.3 IoT for Agri-Food Traceability	22
2.3.4 Artificial Intelligence in Agri-Food Monitoring	23
2.3.5 Blockchain and Distributed Ledger Technologies	24
2.4 Positioning and Comparative Evaluation of Enabling Technologies ...	26
2.4.1 Positioning Across the Agri-Food Supply Chain	26
2.4.2 Comparative Analysis Based on Core Criteria	26
2.5 Gaps in the State-of-the-Art and Thesis Positioning	28
2.5.1 Limitations of Existing TRAC Technologies	29

2.5.2 Research Gap Addressed by This Thesis	29
2.5.3 Summary	31
3 Data Reduction Strategies for Resource-Constrained On-Device	
Machine Learning	33
3.1 Introduction	33
3.2 Histogram-Based Dataset Distillation	36
3.3 Average Based Dataset Distillation	37
3.4 Preliminary Tests on Driver Behavior Dataset	38
3.4.1 Purpose and Experimental Setup	38
3.4.2 Dataset Description	39
3.4.3 Model Architecture and Experimental Evaluation	39
3.5 Distillation on MNIST and Application Potential in Agri-Food	44
3.5.1 Purpose and Experimental Setup	44
3.5.2 Data Transmission and Processing	46
3.5.3 Model Architecture and Training	48
3.5.4 Test Results	50
3.5.5 Preliminary Distillation in the Agri-Food Sector	52
3.5.6 Discussion	53
3.6 Full Pipeline on Grain Dataset: From Capture to Distillation	54
3.6.1 Purpose and Experimental Setup	54
3.6.2 Dataset Description, and Processing	55
3.6.3 On-device Pipeline	57
3.6.4 Model Architecture and Training	59
3.6.5 Model Architecture	59
3.6.6 Model Training Structures	60
3.6.7 Model Training	60
3.6.8 Results and Discussion	60
3.7 Conclusion	62
4 On-Device Edge AI Solutions for the Agri-Food Sector	65
4.1 Introduction	65
4.2 AI-Enhanced IoT for Ochratoxin A Detection in Grape Production	67
4.2.1 State of the Art and Motivation for On-Device OTA Risk Assessment	67
4.2.2 Proposed System	69
4.2.3 Evaluation Function	71
4.2.4 Hardware, Software, and Dataset Description	73
4.2.5 Model Development	80

4.2.6 Experimentation and Results	85
4.2.7 Results and Discussion	87
4.2.8 Conclusion	90
4.3 On-Device AI for Fermentation Stage Classification	91
4.3.1 Problem Statement and Motivation	91
4.3.2 State of the Art in Fermentation Monitoring	92
4.3.3 Case Study Overview and Pipeline	94
4.3.4 AI Model Design	97
4.3.5 Training and Offline Evaluation	99
4.3.6 Proposed Setup and Model Deployment	101
4.3.7 Model Conversion	102
4.3.8 Model Real-Time Inference and Discussion	103
4.3.9 Discussion and Conclusion	105
4.4 Conclusion	106
5 RFID-Augmented Edge-AI Framework for End-to-End Agri-Food	
Traceability and Anti-Counterfeiting	109
5.1 Introduction	109
5.2 Problem Definition, Design Requirements, and State-of-the-Art	
Positioning	111
5.3 Proposed RFID-Augmented Edge-AI Traceability Framework	113
5.3.1 Overall System Architecture	113
5.3.2 End-to-End Traceability Workflow	116
5.3.3 Edge Intelligence and On-Device Data Processing	116
5.3.4 RFID-Based Product-Data Binding Strategy	117
5.3.5 Blockchain-Inspired Data Integrity and Provenance Layer	117
5.4 Hardware and Software Implementation	117
5.4.1 Hardware	117
5.4.2 Software	119
5.4.3 Embedded AI Module	119
5.5 Experimental Case Study: Grape Harvest Traceability	120
5.5.1 Case Study Context	120
5.5.2 Experimental Setup	122
5.5.3 Implementation Workflow	122
5.5.4 Data Encoding and Storage Strategy	124
5.5.5 Backend Integration and Verification Mechanisms	124
5.6 Results	125
5.6.1 On-Device Execution and Resource Utilization	125

5.6.2 RFID Read and Write Performance.....	126
5.6.3 Operational Efficiency and Energy Consumption.....	127
5.6.4 BC-Based Data Integrity Evaluation.....	127
5.6.5 End-to-End Traceability Verification.....	128
5.7 Discussion and Limitations.....	129
5.7.1 Security and Threat Mitigation Considerations.....	129
5.7.2 System-Level Evaluation and Architectural Limitations.....	131
5.7.3 Data Protection and Privacy Considerations.....	132
5.8 Conclusion.....	132
6 Conclusions and Future Work.....	135
References.....	139

List of Figures

1.1 Overall architecture of the proposed TRAC framework, showing on-device data collection and distillation (Chapter 3), model training and Edge AI deployment (Chapters 3-4), and the integration of RFID, cloud, BC, and web dashboard components (Chapter 5).	9
2.1 Agri-food supply chain	12
2.2 EVOO Agri-food supply chain	14
2.3 Development stages of TS	16
2.4 QR Code encoding and decoding	19
2.5 Workflow of DNA barcoding	19
2.6 Example of UHF RFID Tag Applied to an Apple for Automated Tracking and Billing	21
2.7 The uses of IoT in Agriculture	22
2.8 An example of a system architecture based on the integration of IoT with AI	23
2.9 Annual Growth of BC-related Publications (2019–2024)	25
2.10 TRAC Technologies in AFSCs Systems	27
2.11 Current limitations in the agri-food TS	30
3.1 Data Distillation technique first test	40
3.2 First Test Accuracy	41
3.3 First Test Loss	41
3.4 First test results - confusion matrix	41
3.5 Data Distillation technique second test	41
3.6 Second Test Accuracy	42
3.7 Second Test Loss	42
3.8 Second test results - confusion matrix	42
3.9 Testbed for ABDSD	45
3.10 MNIST dataset used for real-time distillation experiments	47

3.11 Example of MNIST data distillation with $\text{Dist_Value} = 3$.	47
3.12 Architecture of the fully connected NN.	48
3.13 Confusion matrix of the model trained using dataset structure A (original dataset).	51
3.14 Confusion matrix of the model trained using dataset structure B ($\text{Dist_Value} = 10$).	52
3.15 Confusion matrix of the model trained using dataset structure C ($\text{Dist_Value} = 3$).	52
3.16 Example of dataset distillation applied to pistachio images	53
3.17 MNIST dataset used for real-time distillation experiments	55
3.18 Examples of hard and soft wheat grains from the dataset.	56
3.19 Distilled soft wheat samples obtained without alignment using $\text{Dist_Value} = 3$	58
3.20 Distilled soft wheat samples obtained after alignment using $\text{Dist_Value} = 3$. Key features remain visible despite data reduction.	59
4.1 Aflatoxin	67
4.2 Ochratoxin	67
4.3 Microscopic images of fungi responsible for Mycotoxin production.	67
4.4 On-device OTA risk assessment system architecture.	70
4.5 Evaluation function relating OTA concentration to time under different temperature and humidity conditions.	72
4.6 M5Stick C Plus (M5C+) device used for environmental data acquisition and on-device OTA classification.	74
4.7 Architecture of the probability-based classification (PCD) model.	83
4.8 Distribution of OTA concentration predictions using real-world harvesting data from Reggio Calabria.	86
4.9 ADD samples distribution	87
4.10 PCD samples distribution	87
4.11 Distribution of test set samples for the ADD and PCD models.	87
4.12 ADD confusion matrix	87
4.13 PCD confusion matrix	87
4.14 Confusion matrices for the ADD and PCD models.	87
4.15 OTA classification results displayed on the M5C+ device, including the predicted class and inference time.	88
4.16 Visual representation of ADD model classification results on the M5C+.	89
4.17 Visual representation of PCD model classification results on the M5C+.	89

4.18 Examples of generated kefir images using DataDreamer showing distinct class characteristics.	96
4.19 Illustration of the cropping and cleaning process for preparing high-quality kefir fermentation images.	97
4.20 Training and validation accuracy and loss curves for the custom CNN trained on cropped images.	100
4.21 Proposed system architecture for Edge AI kefir fermentation stage monitoring, including model training, deployment, and real-time on-device inference.	101
4.22 Luxonis AI Camera PoE integrated for real-time Edge AI fermentation monitoring.	102
4.23 Model conversion workflow: trained Keras model exported to TFLite and then to .blob format for deployment on the Luxonis AI Camera PoE.	103
4.24 Examples of real-time inference outputs during live monitoring showing predicted classes and fluctuating confidence values.	104
4.25 Real-time inference examples using YOLOv10. Captions indicate original and predicted classes. Misclassification occurred primarily for over-fermented kefir samples.	105
5.1 Comparative overview of agri-food traceability architectures and the positioning of the proposed AI-IoT-BC framework with passive RFID tags.	114
5.2 Workflow of the proposed RFID-Edge-AI system during the harvest process.	115
5.3 Prototype setup showing the Arduino Nano 33 BLE, SHT40 Sensor, EM4325 RFID tag, and the battery placed behind the sensor.	118
5.4 Overall system architecture.	121
5.5 Screenshot from the mobile reader application.	123
5.6 AgriTraceX front-end displaying product registration and supply-chain tracking information linked to RFID tags.	125
5.7 Screenshot of BC information displayed after scanning the product QR code.	128

List of Abbreviations

TRAC	Traceability and Anti-Counterfeiting
OSF–OSB	One-step-forward and one-step-backward
AFSCs	Agri-Food supply chains
ICT	Information and Communication Technology
IoT	Internet of Things
AI	Artificial Intelligence
RFID	Radio Frequency Identification
NFC	Near Field Communication
BC	Blockchain
MCUs	Microcontrollers Units
EVOO	Extra Virgin Olive Oil
ML	Machine Learning
OTA	Ochratoxin A
TS	Traceability Systems
FSL	Food Safety Law
NHFPC	National Health and Family Planning Commission
EU	European Union
PDO	Protected Designation of Origin
RASFF	Rapid Alert System for Food and Feed
US	United States
COOL	Country-of-Origin Labeling
FSMA	Food Safety Modernization Act
SMEs	Small and medium-sized enterprises
CNNs	Convolutional Neural Networks
RNNs	Recurrent Neural Networks
LSTM	Long Short-Term Memory
GRU	Gated Recurrent Unit
RL	Reinforcement Learning

DRL	Deep Reinforcement Learning
FL	Federated Learning
IPFS	Interplanetary File System
ABDSD	Average Based Dataset Distillation
HBDD	Histogram-Based Dataset Distillation
M5C+	M5StickCPlus
IMU	Inertial Measurement Unit
HAR	Human Activity Recognition
NN	Neural Network
ReLU	Rectified Linear Unit
RPi	Raspberry Pi
PCA	Principal Component Analysis
PSD	Power Spectral Density
CL	Centralized Learning
HPLC	High-Performance Liquid Chromatography
LC-MS	Liquid Chromatography–Mass Spectrometry
TLC	Thin-Layer Chromatography
AOS	Artificial Olfactory Sensor
CCT	Compact Convolutional Transformer
SVM	Support Vector Machines
PBS	Phosphate-buffered saline
VICAM	Ochrates immunoaffinity column
LoD	Limits of Detection
LoQ	Limits of Quantification
TFLite	TensorFlow Lite
ADD	Anomaly Detection Dataset
PCD	Probability Classification Dataset
FLOPs	Floating Point Operations
BP	Backpropagation
e-nose	Electronic Nose
PoE	Power-over-Ethernet
EPC	Electronic Product Code
TID	Tag Identifier

List of publications

Conferences

1. M. R. Sebti, S. Carabetta, M. Russo, and M. Merenda, “*Ochratoxin A growth probability estimation in wine production using AI-powered IoT devices*”, in Proc. of the IEEE Conference on AgriFood Electronics (CAFE), Torino, Italy, 2023, pp. 152–156. doi: [10.1109/CAFE58535.2023.10292006](https://doi.org/10.1109/CAFE58535.2023.10292006).
2. A. Fotia, A. Macheda, M. R. Sebti, C. Nunnari, and M. Merenda, “*Design of a portable water pollutants detector exploiting ML techniques suitable for IoT devices integration*”, in *Applications in Electronics Pervading Industry, Environment and Society (ApplePies 2023)*, Lecture Notes in Electrical Engineering, vol. 1110, Springer, Cham, 2024. doi: [10.1007/978-3-031-48121-5_51](https://doi.org/10.1007/978-3-031-48121-5_51).
3. M. R. Sebti, A. Accettola, R. Carotenuto, and M. Merenda, “*Dataset Distillation Technique Enabling ML On-board Training: Preliminary Results*”, in Proc. of the 54th Annual Meeting of the Italian Electronics Society (SIE), 2023, pp. 379–384. doi: [10.1007/978-3-031-48711-8_46](https://doi.org/10.1007/978-3-031-48711-8_46).
4. M. R. Sebti, Z. Dakhia, A. Macheda, D. Iero, M. Russo, and M. Merenda, “*On-Device Dataset Distillation: The MNIST Use Case and Initial Experiments on Agri-Food Product Classification*”, in Proc. of the 10th International Conference on Smart and Sustainable Technologies (SpliTech), Bol and Split, Croatia, 2025, pp. 1–6. doi: [10.23919/SpliTech65624.2025.11091749](https://doi.org/10.23919/SpliTech65624.2025.11091749).
5. Z. Dakhia, A. Lazzaro, M. R. Sebti, D. Iero, M. Russo, and M. Merenda, “*Preliminary Analysis of Federated Learning in Agrifood: From MNIST to Real-World IoT Applications in a Heterogeneous Environment*”, in Proc. of the 10th International Conference on Smart and Sustainable Technologies (SpliTech), Bol and Split, Croatia, 2025, pp. 1–7. doi: [10.23919/SpliTech65624.2025.11091733](https://doi.org/10.23919/SpliTech65624.2025.11091733).
6. M. R. Sebti, K. Laabassi, M. Russo, and M. Merenda, “*Lightweight On-Device Pre-processing for Agri-food Classification*”, in Proc. of the IEEE Conference on Agri-Food Electronics (CAFE), 2025. (accepted).

Journals

1. M. R. Sebti, A. Arciello, M. Russo, and M. Merenda, “An RFID-augmented information retrieval technique for AI-enabled IoT devices ensuring agri-food traceability and anti-counterfeiting”, IEEE Journal of Radio Frequency Identification, vol. 10, pp. 122–135, 2026. doi: [10.1109/JRFID.2026.3663762](https://doi.org/10.1109/JRFID.2026.3663762).
2. M. R. Sebti, Z. Dakhia, S. Carabetta, R. Di Sanzo, M. Russo, and M. Merenda, “Real-Time Classification of Ochratoxin A Contamination in Grapes Using AI-Enhanced IoT”, Sensors, vol. 25, no. 3, p. 784, 2025. doi: [10.3390/s25030784](https://doi.org/10.3390/s25030784).
3. Z. Dakhia, A. Lazzaro, M. R. Sebti, M. Russo, and M. Merenda, “On-Device Federated Learning for Energy-Efficient Smart Irrigation”, Electronics, vol. 14, no. 21, p. 4311, 2025. doi: [10.3390/electronics14214311](https://doi.org/10.3390/electronics14214311).

Introduction

1.1 Background and Context

One of the most significant economic sectors worldwide is the agri-food industry. Establishing an effective agri-food supply chain requires substantial resources and must simultaneously ensure sustainability, adequate quantity, and affordability. Among the most crucial services that a supply chain must provide is traceability which is the ability to know where food products come from and to guarantee their quality. Reinforcing traceability strengthens the entire agri-food chain by promoting transparency, safety, and consumer trust. Historically, the industrialization of the food sector began with innovations such as the steam engine and mechanization systems in the nineteenth century, which laid the foundation for large-scale food production. Over time, these solutions evolved into digitization [1], transforming how agri-food products are produced, monitored, and distributed. In recent years, globalization has further reshaped our food habits and trading systems, dramatically expanding global import and export channels. These evolving trends, accentuated by events such as the COVID-19 pandemic, have underscored the vulnerability and importance of global agri-food production and distribution systems. Within these complex systems, services such as Traceability and Anti-Counterfeiting (TRAC) have become indispensable due to the valuable information they provide to both consumers and stakeholders throughout the supply chain. According to Regulation (EC) No 178/2002, traceability is defined as the ability to follow the movement of a food product one-step-forward and one-step-backward (OSF–OSB) [2], creating a chain-wide record of the product’s history [3] and providing details about its origin, safety, and quality across various stages of production, processing, and distribution [4]. Anti-counterfeiting, on the other hand, refers to the illegal use of registered trademarks or the substitution of genuine products with inferior imitations to deceive consumers [5]. Historically, the lack of reliable traceability systems resulted in numerous food safety issues, including the spread of foodborne diseases, fraudulent

substitutions, and the absence of data regarding transportation or storage conditions. The growing consumer demand for year-round, globally available products irrespective of seasonality has further emphasized the need for real-time traceability and authentication [6, 7]. The inability to access accurate origin or quality information negatively affects consumer health, market stability, and environmental sustainability [8]. Indeed, the absence of robust traceability mechanisms has been linked to increased food waste, with approximately 30–50% of agri-food products perishing before reaching consumers [9, 10]. This inefficiency drives up operational costs, undermines consumer confidence, and complicates compliance with regulatory standards [11]. Despite significant technological progress, ensuring global traceability, food safety, quality assurance, and timely delivery remains a challenge. The complexity of Agri-Food supply chains (AFSCs) which typically involve multi-variable stages from farming, processing, packaging, and distribution to retail—creates difficulties in maintaining end-to-end visibility [12]. Each stage must be monitored, recorded, and verified to preserve product integrity and enable accurate backtracking when necessary. In recent decades, Information and Communication Technology (ICT) developments have profoundly transformed industrial systems, enabling seamless communication between diverse components. Among these, the Internet of Things (IoT) has emerged as a key technology connecting devices, sensors, and systems to collect and exchange data across physical and digital environments. The adoption of IoT, supported by Artificial Intelligence (AI), has facilitated numerous innovations in areas such as environmental monitoring, infrastructure management, industrial automation, and healthcare. In the context of agri-food, IoT provides a foundation for intelligent traceability solutions, enabling continuous monitoring and analysis of environmental and operational parameters. However, IoT devices are typically constrained by limited processing power, memory capacity, and energy resources, while simultaneously generating large volumes of data. This poses scalability and efficiency challenges as the number of connected devices continues to rise [13]. To address these limitations, recent advances have focused on deploying AI models at the edge, where computation occurs locally on the device or near the data source. This approach reduces latency, enhances privacy, and enables advanced functions such as data distillation, feature extraction, and real-time decision-making directly on low-cost embedded systems. To ensure advanced and reliable TRAC, the integration of food databases with next-generation electronic devices has become essential. Such integration allows the fusion of sensing, communication, and computation at multiple layers of the supply chain. Emerging technologies including IoT, Radio Frequency Identification (RFID), Near Field Communication (NFC), barcoding, Blockchain (BC), and AI enable the continuous tracking of products from farm

to fork, improving transparency, responsiveness, and efficiency. These technologies collectively contribute to establishing robust, interoperable, and scalable traceability frameworks capable of supporting both industrial and small-scale agri-food actors.

1.2 Thesis Motivation and Contributions

Building upon the limitations identified in existing TRAC systems, this thesis addresses one of the main practical challenges of modern agri-food supply chains: the large volume of heterogeneous data continuously generated by distributed IoT devices and the difficulty of processing this data efficiently on resource-constrained embedded platforms. To overcome these constraints, the thesis investigates intelligent, resource-efficient mechanisms that enable local data processing, learning, and trustworthy traceability directly at the edge. The first research objective focuses on mitigating the data volume problem and enabling on-board learning through efficient data reduction. To this end, a data distillation strategy is proposed to compress and summarize sensor and monitoring data while preserving the most representative and informative characteristics. By reducing redundancy, communication overhead, and memory requirements, the proposed approach makes training and inference feasible on low-power Microcontrollers Units (MCUs) and embedded IoT nodes. This contribution, detailed in Chapter 3, demonstrates that edge learning can be achieved even under strict constraints in computation, storage, and bandwidth, thereby improving scalability and energy efficiency across the supply chain. The second objective investigates the practical deployment of embedded AI-IoT solutions across different stages of the agri-food chain. Several real-world systems based on low-cost sensors, MCUs, and Edge AI models are developed to perform local monitoring and real-time decision-making. These include environmental risk assessment for crop safety and computer-vision-based quality evaluation during processing. By executing inference directly on-device, these solutions reduce latency, enhance autonomy, and limit dependence on cloud connectivity while improving operational reliability and product quality control. The design, implementation, and validation of these on-device intelligence systems are presented in Chapter 4. Building upon these edge capabilities, the final objective proposes a complete end-to-end TRAC architecture that integrates sensing, local intelligence, and secure identification technologies. The resulting framework combines Edge AI, RFID-based product identification, cloud storage, and a BC-inspired integrity layer to securely bind physical products to their digital records and ensure data authenticity and transparency throughout the supply chain. This integrated system demonstrates how distributed IoT and AI nodes can collaboratively support trustworthy, real-time trace-

ability from farm to fork. The full framework and its experimental validation are described in Chapter 5. Collectively, these contributions establish a unified approach to agri-food traceability that couples efficient edge data processing, embedded AI, and secure system integration, enabling scalable, low-cost, and reliable certification and monitoring solutions for modern agri-food ecosystems.

1.3 Thesis Organization

The overall structure of the thesis and the relationship between its main contributions are illustrated in Fig 1.1. The remainder of this thesis is organized as follows:

- Chapter 2 presents a comprehensive review of the evolution of TRAC systems in agri-food supply chains. It discusses enabling technologies, including barcodes, IoT, AI, RFID, and BC, and analyzes the current state-of-the-art solutions that motivate the proposed framework.
- Chapter 3 focuses on data collection at the edge and introduces the proposed on-device data reduction and distillation strategy. This chapter addresses efficient sensing, local preprocessing, and data structuring on resource-constrained IoT devices, aiming to reduce memory usage and communication bandwidth.
- Chapters 4 describe the model training and on-device AI components of the framework. These chapters cover training paradigms, as well as the deployment of trained models on edge devices for real-time inference in agricultural environments.
- Chapter 5 introduces the complete TRAC framework. As shown in Fig 1.1, this chapter integrates Edge AI with RFID-based identification, cloud storage, and a BC-inspired layer to ensure data integrity, transparency, and end-to-end traceability, supported by a web-based dashboard.

Finally, Chapter 6 concludes the thesis by summarizing the main contributions and outlining potential directions for future research.




Fig. 1.1: Overall architecture of the proposed TRAC framework, showing on-device data collection and distillation (Chapter 3), model training and Edge AI deployment (Chapters 3-4), and the integration of RFID, cloud, BC, and web dashboard components (Chapter 5).

State-of-the-Art

2.1 Introduction

Agri-food supply chains are among the most complex and heterogeneous production ecosystems, involving numerous independent actors, variable product characteristics, and multiple transformations from the point of origin to the final consumer. Unlike industrial goods, agricultural products are highly perishable and sensitive to environmental conditions, requiring continuous monitoring during harvesting, processing, storage, and transportation. Because data is generated and modified at every stage, ensuring accurate, end-to-end traceability remains a major challenge, especially when heterogeneous systems, fragmented documentation practices, and limited interoperability are common across the sector. An example of a typical agri-food supply chain is illustrated in Fig. 2.1. A further complication arises from the growing prevalence of food fraud, which disproportionately affects high-value products with strong economic and cultural importance. Items such as grape-based products, Extra Virgin Olive Oil (EVOO), honey, and saffron are historically among the most targeted commodities. Wine may be diluted, mislabeled by region, or falsely assigned higher quality categories. Olive oil is often adulterated with cheaper vegetable oils or fraudulently marketed as extra virgin. Honey fraud including sugar syrup dilution, misrepresentation of floral origin, and blending with low-grade imports is recognized as one of the most widespread forms of global food adulteration. Saffron, as one of the most expensive spices, is particularly vulnerable to substitution with dyed plant material or synthetic additives. These practices are difficult to detect without robust analytical or digital verification tools, and they create significant financial and reputational damage across the sector. Such vulnerabilities highlight the limitations of traditional identification and documentation systems. Labels, certificates, manual logs, and basic barcodes are easily replicated or manipulated, and once a product is repackaged, mixed, or altered, downstream actors have few reliable ways to verify its authenticity. Effective anti-counterfeiting therefore

requires secure, tamper-resistant mechanisms capable of capturing each transformation and detecting inconsistencies in product identity, composition, or origin. This necessity has driven the integration of advanced digital technologies, including IoT-based sensing, automated identification, molecular authentication, Machine Learning (ML), and BC-backed data integrity. Over the past two decades, the state of the art has shifted toward interconnected, intelligence-driven traceability infrastructures. Modern TRAC systems combine sensing, automation, secure data exchange, and advanced analytical capabilities to provide improved transparency, reduce opportunities for fraud, and ensure that product information remains consistent and trustworthy throughout the supply chain. This chapter reviews the main developments that define the current technological landscape of agri-food TRAC. It examines the evolution of TRAC mechanisms, discusses the most influential technologies employed today, and highlights contemporary research directions aimed at improving transparency, authenticity, and trust across agri-food supply chains.




Fig. 2.1: Agri-food supply chain

2.2 Traceability and Anti-Counterfeiting in Agri-Food

Traceability in the agri-food sector refers to the structured process of recording, linking, and retrieving information that describes a product's life cycle. Unlike general definitions of traceability as simple batch identification, traceability in food systems requires a detailed reconstruction of the transformations a product undergoes from primary production to consumption. This includes the identification of all inputs, processing steps, handling conditions, and movements between supply chain actors. A traceability system must therefore provide continuity of information across heterogeneous processes, ensuring that each transformation can be associated with the specific lot or item in question. By continuously tracking product movements and collecting key data at each stage of the supply chain, it becomes possible to ensure both TRAC, as shown in Fig. 2.2. Importantly, traceability is not represented by a universal or predefined data, instead, it is shaped by the unique characteristics, risks,

and quality attributes of each commodity. For biological and high-variability products, the relevant data points may differ significantly. In the grape-based products sector, parameters such as grape variety, fermentation conditions, geographic origin, and the risk of Ochratoxin A (OTA) formation carry essential traceability value because they influence both safety and regulatory classification. In the case of honey, traceability relies on information such as floral origin, pollen composition, moisture content, and indicators of dilution or heat treatment. For olive oil, key traceability attributes include cultivar type, pressing method, storage conditions, free acidity, and chemical markers that distinguish genuine EVOO from adulterated blends. Saffron traceability often requires documentation of harvesting dates, drying techniques, geographic provenance, and biochemical profiles that verify authenticity. These examples demonstrate that traceability is intrinsically product-specific and must be designed around the parameters that define identity, safety, and quality for each commodity. While traceability provides visibility and accountability, anti-counterfeiting focuses on protecting products against intentional fraud. Counterfeiting in agri-food systems encompasses a wide range of illicit practices such as adulteration, dilution, origin misrepresentation, substitution with lower-quality materials, and the use of forged labels or certificates. High-value commodities like: honey and saffron are recurrent targets due to their economic value and the difficulty of visually detecting fraud. Anti-counterfeiting systems therefore aim to detect tampering, verify authenticity, and prevent unauthorized replication of product identifiers. Achieving these goals requires the integration of secure identification elements, tamper-evident packaging, chemical or molecular authenticity markers, and digital mechanisms capable of validating both product and data integrity. Taken together, TRAC address complementary needs within agri-food supply chains: the former establishes a transparent record of the product journey, while the latter safeguards the authenticity and reliability of that information. The following subsections expand on these foundations by examining the mechanisms, classifications, and technological approaches that support modern traceability and fraud-prevention strategies in the agri-food domain.

2.2.1 Evolution of Traceability Over Years and Across Borders in Agri-Food Supply Chains

Traceability Systems (TS) have undergone a profound transformation over the past two decades, driven by technological progress, stricter regulatory environments, and growing expectations for transparency in global food supply chains. As illustrated in Fig. 2.3 from [14], the evolution of these systems can be broadly categorized into three generations, each reflecting a significant shift in data management, technologi-




Fig. 2.2: EVOO Agri-food supply chain

cal capability, and supply chain integration. The first generation, commonly referred to as TS 1.0 [14], was characterized by fundamental record-keeping processes. Information was documented using paper-based methods or simple electronic files, typically limited to batch identifiers, basic supplier information, and delivery documentation. These systems enabled rudimentary OSF–OSB tracking, meeting the minimum regulatory requirements for locating product origin and destination. Although limited in scope, TS 1.0 established the conceptual foundation for linking product identity with movement across supply chain nodes. Around 2008, the expansion of IoT technologies initiated the transition toward TS 2.0 [14]. This generation marked the beginning of interconnected traceability, where sensors, RFID tags, and automated identification systems facilitated real-time data acquisition. IoT-enabled monitoring of parameters such as temperature, humidity, and location became especially important for perishable commodities, improving cold-chain integrity and reducing spoilage risks. By integrating sensor data with electronic information systems, TS 2.0 enabled more responsive and reliable surveillance of food conditions throughout distribution and storage. The most recent development, TS 3.0, reflects the integration of advanced data analytics, artificial intelligence, and BC technologies [14]. These systems move beyond real-time monitoring toward predictive, secure, and intelligence-driven traceability. AI models assist in anomaly detection, quality forecasting, and automated decision-making, while BC provides immutable data storage and decentralized verification, reducing opportunities for manipulation and fraud. TS 3.0 aligns with the increasing need for transparent, data-rich supply chains capable of supporting sustainability metrics, provenance claims, and rapid response to safety incidents. While technological evolution has been global, regula-

tory approaches differ significantly across major regions. Each jurisdiction has developed policies tailored to local challenges, risk perceptions, and governance structures. In China, the 2015 Food Safety Law (FSL) represents a major milestone in strengthening national traceability practices [15]. The Chinese Food and Drug Administration (CFDA) oversees implementation, while the National Health and Family Planning Commission (NHFPC) conducts risk assessments to guide safety standards [16]. With more than 50 specific rules regulating food production and distribution [17], China has established a comprehensive system designed to monitor products from primary production to retail [18]. Technological adoption has also accelerated, with increased use of IoT, QR codes, and BC in government-supported pilot programs. In the European Union (EU), traceability is governed primarily by Regulation EC 178/2002 [2], which mandates OSF–OSB documentation for all food and feed businesses. Operators must maintain detailed records of suppliers, recipients, and product flows, ensuring that items can be rapidly identified and removed from the market when necessary [19]. For products with specific designations such as Protected Designation of Origin (PDO) items or genetically modified organisms (GMOs), additional documentation is required. The EU also operates the Rapid Alert System for Food and Feed (RASFF), which enables real-time communication among member states for food safety events [20]. This system is a cornerstone of the EU's coordinated and proactive food safety infrastructure. The United States (US) has traditionally adopted a more decentralized approach. Traceability requirements vary by sector, and no nationwide mandatory system exists for all food categories. However, regulatory initiatives are increasingly shaped by consumer expectations, export requirements, and public health concerns [21]. Country-of-Origin Labeling (COOL) legislation plays a significant role in providing transparency, particularly in meat and livestock sectors [22]. Although some industries, such as beef and poultry, still lack unified traceability frameworks [23], recent efforts driven by the Food Safety Modernization Act (FSMA) have expanded traceability obligations and promoted digital record-keeping across high-risk food categories. Overall, the evolution of traceability demonstrates a clear trajectory: from basic documentation toward interconnected, data-driven, and intelligence-enabled systems. Regional regulatory differences illustrate how local contexts influence traceability priorities, yet all share the same underlying goal creating transparent, reliable, and secure agri-food supply chains that protect consumers and support fair market practices.




Fig. 2.3: Development stages of TS

2.2.2 Challenges in Modern Agri-Food Traceability and Anti-Counterfeiting Solutions

Despite significant technological progress, modern traceability systems continue to face a number of unresolved challenges. The most relevant issues are summarized below:

- **Fragmented and heterogeneous supply chain structures:** agri-food chains involve farmers, cooperatives, processors, logistics providers, distributors, and retailers, each operating with different technological capacities and data management tools. This heterogeneity leads to incompatible data formats, missing information links, and inconsistent documentation practices. As a result, traceability is often discontinuous, with key transitions breaking the continuity of product identity.
- **Data quality, integrity, and verification difficulties:** traceability systems rely heavily on accurate data, yet much information is still recorded manually or through semi-automated processes. Human errors, incorrect timestamps, mislabeled batches, duplicate entries, and calibration issues in sensors frequently compromise data integrity. Ensuring the authenticity of submitted data is also challenging, as traditional systems lack mechanisms to verify whether entries have been modified, forged, or selectively omitted.
- **Product-specific variability in traceability requirements:** there is no universal traceability data applicable across all commodities. Each product category requires unique information to represent its identity and quality. For example, OTA development, fermentation conditions, and grape provenance are essential for wine, pollen spectrum, moisture content, and heat treatment indicators are

relevant for honey, acidity, peroxide values, and cultivar type are crucial for olive oil; biochemical fingerprint and drying methods determine saffron authenticity. Designing flexible systems that can accommodate such diversity without inflating complexity remains a major challenge.

- **Low digital adoption among small and medium-sized enterprises (SMEs):** SMEs represent the majority of producers in global agri-food chains, yet they often lack the financial resources, digital literacy, or infrastructure required to adopt advanced traceability solutions. High system costs, subscription fees, and technical maintenance discourage participation, resulting in partial traceability that is limited to larger or export-oriented actors. This uneven adoption weakens chain-wide visibility and introduces structural vulnerabilities.
- **Limited real-time monitoring and event detection:** while IoT and sensor networks enable real-time data collection, their deployment remains sporadic due to connectivity constraints, energy limitations, and device maintenance issues. Consequently, many traceability systems still operate in delayed reporting mode, hindering timely detection of temperature excursions, contamination events, cold-chain abuses, or fraud attempts. Real-time traceability would allow proactive intervention, but its implementation remains technically and economically demanding.
- **Increasing sophistication of food fraud mechanisms:** fraudsters increasingly exploit gaps between physical and digital systems. Techniques such as replicating QR codes, reusing legitimate packaging, altering certificates, or blending high-value products with cheaper substitutes often bypass conventional traceability systems. Digital identifiers alone cannot prevent substitution or relabeling unless supported by secure physical markers, chemical/molecular assays, or tamper-evident packaging. Counterfeiting is particularly problematic for high-value products such as EVOO, wine, honey, and saffron.
- **Privacy, confidentiality, and data-sharing constraints:** traceability requires sharing information across all actors, yet many stakeholders are reluctant to disclose internal operational data, supplier lists, or quality parameters due to competitive concerns. Furthermore, regulations impose strict limitations on data sharing and retention. Balancing transparency with data privacy, while maintaining compliance with legal frameworks, remains a complex challenge for large-scale implementations.
- **Lack of harmonized international standards:** traceability requirements vary significantly between regulatory ecosystems such as the EU, the US, and China. Each region imposes different documentation procedures, labeling norms, and compliance expectations. This lack of harmonization creates barriers for cross-

border trade, especially for exporters who must adapt to multiple systems simultaneously. Interoperable traceability remains difficult without globally aligned data standards and certification models.

- **Scalability and integration challenges with modern technologies:** integrating IoT networks, edge devices, cloud platforms, BC ledgers, and AI-based analytics requires robust digital infrastructure. Rural and developing regions often lack stable internet connectivity, affordable sensor networks, or reliable power sources. Even when technology is available, integrating diverse platforms into unified traceability workflows demands technical expertise that many supply chain actors do not possess.
- **High implementation and lifecycle costs:** technologies such as RFID, BC, computer vision systems, and distributed sensor networks impose significant initial and ongoing costs. These include hardware acquisition, software licenses, data storage, cybersecurity, system updates, and personnel training. For commodities with low profit margins, such costs can outweigh perceived benefits, slowing widespread adoption.
- **Weak linkage between the physical product and its digital identity:** one of the most fundamental challenges is ensuring that the physical product corresponds to its digital record. Labels can be copied, RFID tags can be transferred, QR codes can be cloned, and packaging can be reused. This “digital-physical gap” enables substitution, dilution, and mislabeling. Without secure binding mechanisms chemical markers, molecular identifiers, tamper-proof seals, or embedded sensor signatures traceability data may not truly reflect the product’s real identity.

These challenges illustrate the need for more robust, integrated, and adaptable traceability solutions capable of operating across diverse agri-food contexts.

2.3 Enabling Technologies for Next-Generation Traceability

The state-of-the-art analysis shows that a wide variety of technologies are currently used to support TRAC in the agri-food sector. These solutions are applied in different ways depending on the system design: some are deployed at a single supply-chain node, others operate across multiple interconnected nodes, and more advanced implementations span the entire AFSCs, enabling end-to-end visibility. Because of this diversity, the technologies found in the literature differ substantially in terms of purpose, functionality, and deployment complexity. While many evaluation criteria exist such as accuracy, environmental resistance, communication range, interoperability, or data capacity, our analysis focuses on the criteria that most strongly influence

real-world adoption in agri-food contexts. Specifically, we consider cost, scalability, reuse potential, traceability performance, and the anti-counterfeiting strength of the technology itself as the most relevant factors for comparing and understanding the enabling technologies used in current research. Based on these criteria, this section concentrates on the technologies most frequently employed in the reviewed state-of-the-art works: barcoding and QR code-based systems, RFID and NFC, IoT, AI, and BC and Distributed Ledger Technologies. These technologies represent the pillars upon which next-generation traceability architectures are being developed.

2.3.1 Barcoding and QR Code-Based Systems




Fig. 2.4: QR Code encoding and decoding

Flow Diagram for DNA Barcoding




Fig. 2.5: Workflow of DNA barcoding

Traditional barcodes and labels have supported food traceability for decades, but their limited data capacity, low durability, and weak anti-counterfeiting protection have motivated the transition toward more advanced identification mechanisms. In this context, QR codes, shown in Fig. 2.4, have been increasingly adopted due to their ability to encode richer information and support rapid mobile-based access. Several works used QR codes to provide BC-linked production information [24, 25], while others applied them to enable product tracing in dairy and vegetable supply chains via consumer-facing mobile applications [26, 27]. Additional studies implemented QR-based anti-counterfeiting schemes by assigning unique digital identifiers verifiable online [28]. Other contributions connected QR scans to cloud and database infrastructures to support traceability for olive oil and various agricultural products [29, 30, 31]. Overall, these uses highlight QR codes as a versatile, low-cost solution for digital traceability and consumer transparency. In parallel, DNA barcoding, with its flow illustrated in Fig. 2.5 from [32], has been used as a molecular-level authentication technique capable of identifying species through standardized genetic markers. Works applying this method verified product identity in tea [33], characterized botanical origins and detected toxic species in honey [34], and revealed adulteration across seafood, spices, botanicals, and probiotics [35]. Other studies employed mini- and medium-length barcodes to authenticate fruit juices, detecting target species at low concentrations and identifying significant mislabeling [36]. These applications show that DNA barcoding provides strong scientific authentication for both TRAC across fresh, processed, and heat-treated products. Together, QR codes and DNA barcoding form complementary and widely used tools that enhance transparency, product identification, and fraud detection in agri-food supply chains.

2.3.2 RFID and NFC



Fig. 2.6: Example of UHF RFID Tag Applied to an Apple for Automated Tracking and Billing

RFID and NFC are widely adopted wireless identification technologies used across multiple domains [37, 38]. In the agri-food sector, both have been deployed at various supply chain stages to improve identification, monitoring, and traceability. Several works applied RFID tags for product tracking during production and distribution [39, 40, 41], including item level implementations such as the UHF fruit labeling example shown in Fig. 2.6 from [42]. Other studies employed RFID for warehouse management by tagging bins and storage units to maintain internal traceability [43, 44], while additional contributions monitored transport conditions, such as temperature in the pineapple chain through RFID-based systems [45]. RFID has also been integrated with gas sensors to monitor fruit quality during transport [46] and used with sensor modules to evaluate freshness indicators in vegetables [47]. Other applications combined RFID with GPS to achieve tree level traceability in orchards [6]. RFID has further been used to reinforce data reliability and product authentication. Some works combined RFID with BC to secure traceability records for fruits, vegetables, and olive oil [41, 7], while others integrated RFID with cryptographic methods to protect critical supply chain data [48]. RFID tags have also supported anti-counterfeiting by enabling verification of product specific identifiers across multiple stages [49]. NFC, as a short-range subset of RFID, has been employed primarily for consumer oriented applications. Several studies integrated

NFC tags with smartphones to allow farmers to update records throughout the supply chain and consumers to access handling and production information [50, 51]. NFC has also been applied to olive oil traceability, enabling bottle level transparency and mobile-based retrieval of production details [52]. Overall, RFID supports automated, real-time traceability across production, storage, logistics, and quality monitoring [39, 40, 41, 43, 45, 44, 48], whereas NFC enhances accessibility and consumer interaction by linking products to digital information [50, 51, 52]. Solutions combining RFID with BC or cryptographic mechanisms further strengthen data reliability and authentication [41, 48, 7]. Together, RFID and NFC play a significant role in improving transparency, product verification, and anti-counterfeiting in AFSCs.

2.3.3 IoT for Agri-Food Traceability




Fig. 2.7: The uses of IoT in Agriculture

The adoption of IoT technologies has grown rapidly across the agri-food sector, enabling continuous monitoring and rapid data exchange between sensors, devices, and servers [53, 54]. As illustrated in Fig. 2.7 from [55], IoT solutions support diverse applications in smart agriculture, playing a central role in modern traceability and transparency frameworks. Several works introduced multi-layer IoT architectures to collect and share production data throughout the supply chain. Some systems employed IoT sensors combined with workflow engines and user applications to enhance real-time traceability and operational integration [56]. Other studies integrated IoT with BC, using sensors, RFID tags, and GPS devices to ensure tamper-resistant traceability records and strengthen anti-counterfeiting measures [57]. Additional IoT-BC frameworks were used to monitor soil and crop conditions in the field, with BC securing recorded data for improved transparency [58]. Similar approaches linked IoT sensors with BC to store product metadata and enable QR-based

verification for consumers [59]. IoT-based smart farming systems have also been applied to capture environmental parameters during planting and production. Several works used IoT sensors with cloud platforms to monitor variables such as temperature, pH, conductivity, and humidity, while enabling consumers to access cultivation histories through QR code scans [60]. Other IoT–server–BC combinations supported real-time monitoring to improve quality control and reduce production inefficiencies [61]. Further studies explored specialized IoT solutions using solar-powered sensors, BLE tags, or multi-sensor nodes to monitor soil, storage conditions, and freshness indicators during transport [62]. Additional low-cost systems integrated RFID modules, Arduino-based sensors, and cloud services to support environmental monitoring and traceability for small-scale producers [63]. Overall, IoT technologies contribute significantly to agri-food traceability by enabling continuous data acquisition, real-time environmental monitoring, and transparent record sharing [56, 58, 60]. When combined with BC, RFID, or QR-based systems, IoT enhances anti-counterfeiting capabilities and strengthens the reliability of supply-chain information [57, 58, 59]. These IoT-enabled approaches support improved visibility, authenticity, and trust across all stages of agri-food production and distribution.

2.3.4 Artificial Intelligence in Agri-Food Monitoring




Fig. 2.8: An example of a system architecture based on the integration of IoT with AI.

AI has become a key component in modern agri-food traceability systems, enabling real-time data analysis, anomaly detection, and product verification across the supply chain [64, 65, 66, 67]. As illustrated in Fig. 2.8 from [64], AI is often combined with IoT infrastructures to support continuous monitoring and decision-making in

agricultural processes. Several works employed AI models for authentication and fraud detection. Convolutional neural networks (CNNs) have been applied to spectral data from IoT enabled devices to authenticate products such as EVOO [68]. Other studies used ML models such as XGBoost to analyze sensor data and predict freshness in fermented foods [69]. For meat traceability, Recurrent Neural Network (RNNs) models including Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) were integrated with IoT and BC to support predictive monitoring across supply chain stages [70]. AI has also been used for supply chain optimization. Reinforcement Learning (RL) and Deep Reinforcement Learning (DRL) frameworks were applied together with BC to optimize production, storage, inventory, and delivery processes while maintaining secure traceability records [71, 72, 73]. Additional studies leveraged drones and CNN architectures to monitor crop health, environmental conditions, and field operations for smart farming applications [74]. Embedded neural network models have also been deployed on IoT connected devices to support contamination risk prediction and environmental monitoring in agricultural settings. In livestock and dairy management, AI was combined with IoT, edge computing, and BC to improve monitoring and traceability within farm operations [75]. Other works proposed multi technology frameworks integrating IoT, AI, BC, big data, or DNA markers to strengthen authenticity verification and end-to-end transparency [76]. Further contributions introduced AI-enabled interfaces such as natural-language querying of BC traceability records [77]. More comprehensive systems used multimodal AI such as Swin Transformers for defect detection and temporal convolutional networks for environmental prediction combined with BC and Federated Learning (FL) to enhance data security and model performance [78]. Overall, AI technologies significantly advance agri-food traceability by enabling real-time monitoring, predictive analytics, and data driven decision-making across production, processing, and distribution stages [68, 69, 70, 74, 75]. When integrated with IoT devices, BC, and advanced learning techniques including Deep Learning (DL), RL, and Large Language Models (LLMs) AI enhances transparency, prevents fraud, and strengthens digital product records [71, 72, 73, 76, 77, 78].

2.3.5 Blockchain and Distributed Ledger Technologies

BC, illustrated through its recent publication growth in Fig. 2.9, has become one of the most widely adopted technologies for secure, transparent, and tamper-evident data management in agri-food traceability [79]. Its decentralized structure and immutability have led many works to integrate BC into supply-chain applications, particularly when trustworthy and auditable information sharing is required. Several studies employed BC together with IoT devices to collect, verify, and store sensor




Fig. 2.9: Annual Growth of BC-related Publications (2019–2024)

data for agri-food traceability [80, 81]. Other works introduced multi-layer systems in which BC records harvesting, processing, distribution, and environmental data for fruits, vegetables, and other agricultural products [82, 83]. BC-based frameworks were also used to store olive oil production information through layered architectures integrating IoT, routing modules, and smart contracts [84]. Additional systems combined BC with QR codes, allowing consumers to retrieve traceability information directly from product scans [85, 86, 87]. Some studies extended these designs to include the entire production process, such as bread and restaurant food traceability, enabling users to access full product histories via mobile applications [88, 89]. BC has also been integrated with distributed storage solutions. Several works used Interplanetary File System (IPFS) to store product data, environmental metrics, or multimedia content, while BC stored the corresponding hash values for verification [90, 86, 87, 91, 92, 93]. Other contributions combined BC with RFID and IoT sensors to link physical items to their digital records, strengthening authenticity and improving product verification [94]. Overall, BC has emerged as a central technology for agri-food traceability, enabling secure, transparent, and immutable data exchange [79, 95, 80, 82, 83, 85]. Integrated with IoT, smart contracts, QR codes, RFID, and IPFS, BC-based systems capture detailed product and environmental information across all supply-chain stages [81, 84, 86, 87, 92, 93]. These approaches significantly enhance traceability, support anti-counterfeiting, and provide consumers with verifiable, tamper-proof product histories [88, 90, 91, 89, 94].

2.4 Positioning and Comparative Evaluation of Enabling Technologies

2.4.1 Positioning Across the Agri-Food Supply Chain

Fig 2.10 illustrates how the main enabling technologies are positioned across the agri-food supply chain. Each technology fulfills a distinct operational role and is adopted at different stages depending on its capabilities. QR codes and NFC are primarily used at the retail and consumer levels, enabling quick verification and product information retrieval through smartphones and other compatible devices. These technologies also support intermediate stakeholders such as farmers, transporters, and retailers by allowing easy access to product records during distribution and storage. RFID is positioned mainly across production, processing, and distribution, where automated identification and real-time tracking are most beneficial. Its ability to capture data without line-of-sight makes it suitable for warehouse management, logistics, and in-transit monitoring, though consumer-side verification is better supported by QR and NFC technologies. DNA barcoding is positioned at verification points, particularly when scientific authentication of species identity or ingredient origin is required. It plays a crucial role in validating authenticity claims, especially for high-value products, although its reliance on laboratory processes limits its use to specific checkpoints rather than continuous monitoring. IoT devices operate throughout the entire supply chain, from on-farm production to post-harvest handling and storage. Sensors capture environmental, logistical, and operational parameters in real time, enabling continuous data collection. AI systems are typically used during data processing and decision-support stages, where collected data is analyzed to detect anomalies, classify products, or predict quality and safety conditions. AI often functions in combination with IoT, BC, or other data systems to provide enhanced traceability intelligence. BC is positioned as a cross-chain logging and verification layer, allowing all stakeholders to record and validate traceability data securely. By storing tamper-evident records, BC supports end-to-end transparency and facilitates shared access to trustworthy product histories. Together, these technologies form a multilayered ecosystem in which sensing, identification, verification, and data logging complement one another across different supply-chain nodes.

2.4.2 Comparative Analysis Based on Core Criteria

Table 2.1 provides a qualitative comparison of the enabling technologies based on five key evaluation criteria: cost, scalability, reuse potential, traceability effectiveness, and counterfeit protection. This comparison reflects how each technology per-




Fig. 2.10: TRAC Technologies in AFSCs Systems

forms relatively within typical agri-food systems, rather than assigning fixed or absolute values. In particular, the assessment of cost is interpreted in relation to the economic capacity of the supply chain: for low income or small scale systems, even moderately priced technologies may be considered high cost, whereas larger or high-value supply chains may more easily absorb the same investment. QR codes are characterized by very low cost, high scalability, and high reuse potential, making them accessible across almost all supply-chain contexts. Their traceability effectiveness is moderate to high, but counterfeit protection remains limited unless paired with secure backend verification. RFID and IoT systems offer a practical balance between cost and scalability, especially where shared infrastructure or low cost sensors are available. These technologies demonstrate high reuse potential and strong traceability effectiveness, though their anti-counterfeit protection is moderate unless additional authentication mechanisms are used. NFC systems, while generally medium in cost and offering high reuse potential, face scalability limitations due to short-range communication and infrastructure needs. Their traceability effectiveness is moderate, and counterfeit protection ranges from low to medium depending on the implementation. DNA barcoding and BC provide very high traceability effectiveness and strong counterfeit protection, but they incur very high implementation costs and present scalability challenges. Their reuse potential is also lower: DNA barcoding is single use, and BC infrastructures require specialized deployment and maintenance. For low income or resource-constrained supply chains, these technologies may be impractical despite their strong technical performance. AI systems, although associated with higher initial costs, offer high scalability, high reuse potential, and substantial traceability effectiveness. When integrated with other technologies, AI improves anomaly detection and verification, thereby enhancing counterfeit protection. However, feasibility again depends on the economic capacity of the system deploying it. By comparing technologies qualitatively across these five criteria and by considering the economic realities of different agri-food supply chains stakeholders can determine which solutions are both technically suitable and financially viable for their specific operational context.

2.5 Gaps in the State-of-the-Art and Thesis Positioning

Modern TRAC systems in the agri-food sector have evolved significantly, integrating identification technologies, IoT sensing, AI-driven analytics, and, in more recent works, BC-based integrity mechanisms. However, as highlighted in the comparative analysis of existing TRAC layers across supply-chain nodes, current solutions still suffer from critical limitations related to fragmented data acquisition, limited on-

Table 2.1: Comparison of TRAC technologies in agri-food systems based on cost, scalability, reuse potential, traceability effectiveness, and counterfeit protection

Technology	Cost Rate	Scalability	Reuse Potential	Traceability Effectiveness	Counterfeit Protection
QR Codes	Very Low	High	High	Medium–High	Low
RFID	Medium	High	High	High	Medium
NFC	Medium	Low	High	Medium	Medium
IoT	Low–High	High	Medium	High	Medium
AI	High–Very High	High	High	High	High
DNA Barcoding	Very High	Low	Low	Very High	Very High
BC	Very High	Medium	High	Very High	Very High

device intelligence, weak linkage between physical products and their digital identities, and strong dependence on constant connectivity for data processing and verification. The diagram of existing TRAC technology layers shows that even the most advanced IoT–AI–BC systems struggle with data availability gaps, high data volume, and poor continuity between the product and the data that represent it. These persistent shortcomings motivate the need for lightweight, real-time, and product-centric mechanisms capable of supporting secure and continuous traceability in resource-constrained agricultural environments.

2.5.1 Limitations of Existing TRAC Technologies

As illustrated in Fig 2.11 the layered TRAC systems overview, existing technologies still depend heavily on cloud connectivity, centralized storage, or passive identification mechanisms that do not travel with the product itself. RFID-based solutions provide only basic identification and remain vulnerable to data loss or manipulation during transitions. IoT-based systems generate large data streams but lack strong binding between sensed data and the physical item, while AI-enabled IoT systems improve monitoring but still rely on remote processing and remain sensitive to connectivity interruptions. Even hybrid AI–IoT–BC solutions address data integrity but do not fully solve the physical–digital identity gap or the challenge of processing large volumes of heterogeneous data on resource-limited devices. These gaps collectively constrain real-time traceability, weaken product authentication, and limit adoption among small-scale producers, as emphasized throughout our analysis.

2.5.2 Research Gap Addressed by This Thesis

A central gap that emerges from the state-of-the-art is the difficulty of managing the large volume of heterogeneous data generated across agri-food supply chains. IoT devices continuously collect environmental, operational, and product-specific




Fig. 2.11: Current limitations in the agri-food TS

information, but transmitting and processing this data using conventional cloud-based pipelines is costly, bandwidth-intensive, and often impractical in agricultural settings. Moreover, standard learning architectures assume the availability of powerful processors and stable connectivity, making them unsuitable for real-world agri-food systems that must rely on low-cost sensing hardware and embedded platforms. These constraints hinder the deployment of traditional ML models for local decision making. This exposes the need for on-device intelligence, where data is pre-processed, distilled, and partially learned directly on the device. Such an approach reduces communication overhead, lowers energy consumption, and improves responsiveness in remote or resource-limited environments. However, despite these advantages, current architectures still suffer from key weaknesses identified in earlier TRAC systems namely, incomplete data continuity, weak linkage between the physical product and its digital identity, and dependency on external infrastructure for tracking and authentication. Addressing these combined limitations requires a new, lightweight learning framework tailored to the agri-food sector, capable of operating efficiently on low-cost devices while strengthening end-to-end tracking and tracing capabilities.

2.5.3 Summary

Building on the limitations identified in this Chapter, the thesis positions itself to address these gaps through four main directions. First, it explores a data reduction technique applicable to both row sensor data and image-based datasets, enabling compact yet informative representations suitable for efficient learning. Second, it develops Edge-AI-based solutions that support on-device processing, real-time inference, and reduced dependence on continuous connectivity. Finally, the thesis introduces a comprehensive framework for end-to-end traceability and anti-counterfeiting, enabling the product to be tracked, authenticated, and verified from its earliest production stages to the final supply-chain node. Together, these contributions provide a cohesive response to the gaps highlighted in the state-of-the-art and help establish a more efficient, intelligent, and trustworthy agri-food traceability ecosystem.

Data Reduction Strategies for Resource-Constrained On-Device Machine Learning

3.1 Introduction

The increasing convergence of AI and the IoT has profoundly transformed data collection and processing across multiple domains, including the agri-food sector [13]. IoT devices are now capable of generating continuous data streams that describe environmental conditions, product quality, and process efficiency [96]. These advances have enabled intelligent and interconnected systems for traceability, monitoring, and control throughout the supply chain. However, such benefits come with a major challenge which is: the massive volume of data produced in real time. Managing and processing these data within the limited computational, memory, and energy capacities of embedded devices remains a key obstacle to the deployment of edge intelligence in agriculture [13, 97]. Modern embedded technologies, such as low-power MCUs and edge boards, have evolved to execute AI algorithms at the device level, operating at milliwatt power ranges [13]. Yet, deploying DL models in such resource-constrained environments requires significant optimization. To achieve this, several model compression techniques have been proposed to reduce model size, computation, and memory usage while maintaining acceptable accuracy. Among the most notable are quantization, which converts high-precision weights and activations into lower-bit representations [98, 99], model pruning [100, 101], which removes redundant neurons and connections to decrease complexity [13, 99], and knowledge distillation, which transfers learned representations from a large teacher model to a smaller student model [13, 102]. These methods collectively make it feasible to deploy trained AI models directly on MCUs and embedded IoT devices, significantly reducing inference latency and energy consumption. While model-level optimization improves efficiency, it does not address another crucial limitation: the handling of large datasets used to train or continuously update these models. In the agri-food domain, IoT sensors, cameras, and smart instruments produce massive quantities of heterogeneous data, including temperature, humidity, and high-resolution

images [96, 97]. Processing all raw data at the edge is often infeasible, both due to hardware restrictions and communication constraints. Therefore, complementary strategies focusing on the data itself rather than on model structure have become increasingly important. Among them, techniques such as data selection [103] and data pruning [104] have been introduced to identify and retain only the most informative samples for training, thereby reducing computational demand without substantially compromising performance. A particularly promising direction in this regard is Average Based Dataset Distillation (ABDSD) [105, 106, 107, 108]. Unlike model compression, which focuses on reducing model parameters, ABDSD aims to condense large datasets into smaller, information-rich subsets that retain the key characteristics of the original data [101]. Through summarization, averaging, or aggregation, it is possible to produce compact data representations that reduce both memory usage and communication overhead. In embedded systems, such a process allows devices to store and transmit only representative data samples, facilitating real-time learning and adaptive decision-making under severe hardware constraints [107]. In addition to ABDSD, another relevant non-intelligent distillation strategy explored in this work is Histogram-Based Dataset Distillation (HBDD). Unlike ABDSD, which compresses data by averaging groups of similar samples, histogram-based distillation focuses on the global statistical distribution of each feature within a class. Each feature can be represented by a histogram, where the bin frequencies indicate how often specific value ranges occur [109]. By selecting the most dominant (high-frequency) bins and extracting representative values typically using bin centers or weighted averages, the method produces a highly compact dataset that preserves the key distributional characteristics of the original data. This approach is particularly effective for structured numerical datasets and was employed in our first experiment [106] to evaluate the feasibility of distribution-preserving distillation under controlled PC-based conditions. In this context, it is important to clarify that the term 'distillation' can refer to two distinct but related concepts. The first is AI-based distillation, in which high-level knowledge such as predicted labels, class probabilities, or feature embeddings is transmitted from a model running on the device to another model or central server [110]. This process relies on intelligent inference and learned representations. The second, which is the focus of this chapter, is a non-intelligent, data-level distillation process that operates directly on raw data [106, 107, 108]. Instead of transferring knowledge, this method mathematically condenses the data using statistical techniques such as averaging or sample aggregation [106, 107, 108]. The objective is to maintain the essence of the data, its overall distribution and variability, while reducing redundancy. This approach is particularly suitable for low-power IoT devices deployed in agricultural environments, where simplicity, speed, and en-

ergy efficiency are crucial [106]. The contribution of this chapter is to address the challenge of managing large-scale agri-food data through the implementation of an on-board dataset distillation framework. By executing the distillation process directly within the MCU, the system minimizes the need for external computation and enables real-time data reduction in the field. The investigation follows a progressive structure comprising three experimental stages, each representing a step toward a fully integrated on-device pipeline:

1. **Experiment 1 – Baseline Evaluation in a PC Environment:** This experiment [106] validates the feasibility of dataset distillation under controlled conditions using a standard computing environment. In this stage, a HBDD technique was applied to reduce the dataset by extracting representative values from the most dominant histogram bins of each feature. The experiment evaluates how this distribution-preserving reduction impacts model training and accuracy, providing a baseline reference for subsequent on-device implementations.
2. **Experiment 2 – On-Board Distillation and Technique Exploration on-device:** The second experiment [107] transitions from simulation to a physical embedded setup. The distillation process is implemented directly on-board the M5StickCPlus (M5C+) MCU. Several data enhancement techniques—such as normalization, image smoothing, and pixel-level filtering—were evaluated in combination with the distillation methods. The results revealed that natural (raw) data distillation, performed without any intelligent or learned operation, achieved the most effective trade-off between performance and computational efficiency.
3. **Experiment 3 – Full On-device Pipeline for Agricultural Image Data:** The third experiment [108] integrates the developed methods into a complete on-board image processing pipeline for agricultural data. The pipeline combines lightweight pre-processing steps such as rotation correction and ABDSD, both executed entirely on the M5C+. The device performs real-time data reduction and transmits the distilled results to an external system via the MQTT protocol for model training. When tested on a custom grain dataset, the pipeline achieved up to 90% data reduction with minimal accuracy loss, demonstrating its feasibility for real-world agri-food monitoring.

These experiments collectively demonstrate that on-board dataset distillation offers a robust and scalable solution for managing large agri-food datasets under resource limitations. The results show that meaningful data compression can be achieved without compromising essential information, supporting real-time analytics in distributed agricultural systems. Moreover, the success of natural data distilla-

tion highlights the potential of low-complexity, non-intelligent methods to complement traditional AI-based approaches in embedded environments. In summary, this chapter contributes to the advancement of embedded AI for agriculture by presenting a real-time, on-board dataset distillation framework designed to handle the increasing volume of agricultural data efficiently. Unlike AI-based distillation, which depends on model inference, the proposed method directly operates on raw sensor and image data, ensuring fast, lightweight, and energy-efficient processing. This approach provides a practical pathway toward achieving intelligent, autonomous, and traceable agri-food systems capable of continuous operation under field conditions.

3.2 Histogram-Based Dataset Distillation

Histogram-Based Dataset Distillation is a non-intelligent data reduction technique that compresses a dataset by exploiting the statistical distribution of its feature values. For a given class in a structured dataset, each feature can be represented as a histogram in which bins correspond to specific value ranges and the bin height indicates the frequency of samples falling within that range. This histogram serves as a compact summary of how the feature values are distributed within the class [109]. The distillation process begins by constructing a histogram for every feature in each class. Once the histogram is obtained, the method identifies the bins with the highest frequencies, which correspond to the most representative or dominant value regions. Instead of retaining all the samples that contributed to these regions, the algorithm selects a small set of distilled values. These representative values can be extracted directly from the bin centers or computed as weighted averages of the samples within each dominant bin. In this way, the distilled dataset preserves the essential characteristics of the original distribution while discarding redundant or infrequent values. For example, consider a feature with the following values for a specific class:

$$X = [10, 12, 11, 50, 52].$$

When constructing a histogram with three bins, the first bin (10–20) contains the values 10, 11, and 12, making it the dominant region. The last bin (40–60) contains 50 and 52, representing the second dominant region. The distilled representative values are therefore computed as weighted averages of the values within these bins:

$$\text{Bin 1 representative: } 11, \quad \text{Bin 3 representative: } 51.$$

The resulting distilled dataset becomes:

$$X_{\text{hist}} = [11, 51].$$

Although the number of distilled samples is significantly smaller than the original, the essential statistical structure of the class is preserved. This makes histogram-based distillation particularly suitable for embedded and IoT systems, where computational and memory resources are limited. The technique efficiently summarizes large datasets using simple operations such as counting and averaging, enabling on-board data reduction while maintaining the key information required for subsequent model training [106, 107, 108].

3.3 Average Based Dataset Distillation

The principle of ABDSD is to create compact yet informative datasets by summarizing several similar samples into a single representative one. The aim is to retain the essential statistical or visual characteristics of the data while significantly reducing the number of samples to be processed, transmitted, or stored. This approach aligns with the broader concept of dataset distillation, in which a large collection of samples is condensed without substantial loss of information content [109, 99]. In traditional dataset representations, each class of a ML dataset can be characterized by the distribution of its feature values. For structured data, this distribution can be visualized using a histogram, where each bin represents a specific range of values, and the bin height corresponds to the number of samples falling within that range [109]. By analyzing such histograms, it becomes possible to identify the most frequent or representative values. In the initial implementation of our approach, these values were selected to represent the key information of each class, thus reducing dataset size while preserving its overall structure. The concept was later extended to numerical and image-based datasets, where the distillation process involves computing the average of multiple consecutive or similar samples to form a new distilled sample. Each group of N original samples is replaced by one representative sample whose values correspond to the mean of the group, effectively summarizing their shared characteristics. The parameter N , referred to as the *Dist_Value*, controls the degree of reduction: larger values of N yield stronger compression but may risk loss of fine details. For example, setting $N = 3$ results in a 66.67% data reduction, while $N = 10$ achieves up to 90% reduction [99]. For image data, ABDSD operates at the pixel level. Each input image I_i is represented as a two-dimensional matrix of pixel intensities of size $H \times W$. The distilled image I_d is computed as the average of N images aligned within the same class according to:

$$I_d(x, y) = \frac{1}{N} \sum_{i=1}^N I_i(x, y), \quad \text{for all } x \in [1, H], y \in [1, W].$$

This process maintains the main visual features of the group, such as shape and texture, while reducing redundancy and noise. Proper alignment of input images is essential: misalignment may blur the product's contours or obscure relevant patterns, particularly in agricultural datasets containing irregularly shaped objects. In practice, the ABDSD procedure was implemented across different platforms. In PC-based experiments, the algorithm was applied offline on structured datasets to validate its reduction capability and impact on classification accuracy. On embedded devices such as the M5C+, ABDSD was executed on-board in real time, where groups of incoming data samples were averaged before transmission to a central server. This implementation enabled lightweight local processing, reducing communication overhead and energy consumption while maintaining the quality of information required for accurate learning. Overall, ABDSD provides a simple yet powerful mechanism for data reduction. It requires no model inference or feature extraction, relying purely on statistical averaging to compress datasets. This property makes it particularly suitable for edge AI applications in agriculture, where devices operate under strict resource constraints but must still preserve critical information for traceability, classification, and decision support.

3.4 Preliminary Tests on Driver Behavior Dataset

3.4.1 Purpose and Experimental Setup

This first study was conceived as an initial validation of the proposed data reduction framework, focusing specifically on the use of HBDD. The main objective was to investigate whether a large dataset could be substantially reduced without causing a critical drop in model performance. This experiment served as the conceptual starting point for later on-device implementations, providing a controlled framework to study the relationship between data reduction and model accuracy under standard computing conditions. The work was conducted entirely in a personal computer environment to focus exclusively on the behavior of the distillation algorithm. This setup eliminated the limitations of embedded hardware and allowed for faster iteration during the early design stage. In this experiment, the central idea was that the statistical structure of a dataset could be captured through the frequency of its feature values. By analyzing these distributions, it becomes possible to identify and preserve the most representative samples, reducing the dataset size while maintaining its overall variability. To achieve this, the HBDD method used a histogram-based selection mechanism. Each feature's distribution for every class was divided into bins, and only the most frequent values were retained as representative points. The

number of selected values, denoted by the parameter N , controlled the compression level and was systematically varied across experiments to assess how different reduction ratios influenced training outcomes. In this way, HBDD functioned as a non-intelligent yet statistically driven form of data distillation. The experiment was implemented in a Python environment using a standard workstation with sufficient computational resources to train and test neural models. Two main test configurations were considered: one applying a strong compression with very few retained values per class, and another with a moderate reduction to preserve more detail. Both were used to examine how the trade-off between dataset size and classification accuracy evolved as the distillation strength increased. This preliminary setup established the foundation for the subsequent stages of this research. It provided key insight into how representative-value selection through HBDD could be used as a viable strategy for dataset reduction and paved the way for its adaptation to embedded systems, where real-time and energy-efficient processing are essential.

3.4.2 Dataset Description

The dataset used in this study [109] was collected from an Inertial Measurement Unit (IMU) designed for Human Activity Recognition (HAR) applications. Each record in the dataset contains eight features: six correspond to the three-axis readings of the gyroscope and accelerometer, while the remaining two represent the vehicle's speed components. The complete dataset includes approximately 32,000 observations distributed across four activity classes. Since the class distribution was initially unbalanced, a balanced subset was created by randomly selecting 6,000 samples per class, resulting in a final dataset of 24,000 observations. This balanced version was used for all subsequent training and evaluation processes.

3.4.3 Model Architecture and Experimental Evaluation

Model Architecture

For the classification task, a fully connected neural network (NN) was employed. The model architecture is composed of five dense layers: the first four layers each contain 100 neurons with Rectified Linear Unit (ReLU) activation functions, while the final layer includes four neurons with a softmax activation function for multi-class classification. The network receives as input the eight features that describe the driver's behavior, as detailed in Section 3.4.2. As a baseline, the same model was trained without applying the HBDD technique, using the complete balanced dataset of 24,000 samples over 25 epochs. Under these conditions, the network achieved a testing accuracy of 97.49%, confirming the suitability of the chosen architecture for

the task. Two separate experiments were carried out to assess the performance of the proposed HBDD approach. In both cases, the previously described fully connected NN was used for classification. The objective was to evaluate how the reduction in training samples affected the learning behavior of the model in terms of accuracy and loss evolution.

First Test: High Compression Scenario




Fig. 3.1: Data Distillation technique first test

In the first experiment, the aim was to simulate the scenario of two IoT devices performing on-board data distillation and then sharing the reduced dataset with a central server for model training. As illustrated in Fig. 3.1 from [106], each device performed local histogram-based distillation by selecting the most frequent feature values, effectively computing a weighted average of these dominant bins. Each device generated eight distilled observations (two per feature), resulting in a total of 16 samples available for training the central model. The network was trained for 15 epochs using this highly compressed dataset. The resulting accuracy and loss curves are shown in Fig. 3.2 and Fig. 3.3 from [106], respectively. Despite the extremely limited number of training samples, the model achieved a validation accuracy of 64.95% (Table 3.1 and Confusion matrix in Fig 3.4 from [106]), confirming that even a minimal amount of statistically representative data can capture basic patterns of the original dataset. The loss curve demonstrates stable convergence across epochs, indicating that the distilled data retained sufficient variability for the model to generalize, albeit with reduced precision.




Fig. 3.2: First Test Accuracy




Fig. 3.3: First Test Loss




Fig. 3.4: First test results - confusion matrix

Table 3.1: First test results

Test Accuracy	Used technique	Number of train samples	Samples per class
64,95%	HBDD	16	4




Fig. 3.5: Data Distillation technique second test

Second Test: Moderate Compression Scenario

In the second experiment shown in Fig 3.5 from [106], the same methodology was repeated using a less aggressive distillation ratio. Each class was represented by 60 samples, for a total of 240 training samples. The model was again trained for 15 epochs under identical conditions. The accuracy and loss trends for this configuration are shown in Fig. 3.6 and Fig. 3.7 from [106]. Compared to the previous test, the model exhibited faster convergence and higher overall accuracy, reaching 82.35% as shown in Fig 3.8 from [106] and summarized in Table 3.2. These results highlight the direct relationship between dataset size and classification accuracy: while smaller datasets preserve core information, moderately larger distilled datasets enable better generalization and higher model confidence.




Fig. 3.6: Second Test Accuracy




Fig. 3.7: Second Test Loss




Fig. 3.8: Second test results - confusion matrix

Table 3.2: Second test results

Test Accuracy	Used technique	Number of train samples	Samples per class
82,35 %	HBDD	240	60

Discussion and Results

The two experiments demonstrate that the HBDD approach can effectively compress datasets while maintaining meaningful predictive performance. In both cases, the model was able to learn the key characteristics of the data despite significant reductions in the number of training samples. The first test, which simulated an extreme compression scenario with only 16 training samples, still achieved a validation accuracy of 64.95%. The second test, which used a moderately larger distilled dataset of 240 samples, reached an accuracy of 82.35%. These results confirm that the proposed method can preserve essential statistical information even when the dataset is reduced by several orders of magnitude. Table 3.3 compares the performance of the ordinary training process with the two distillation-based configurations. The baseline model, trained on the complete dataset of 24,000 samples, achieved an accuracy of 97.49%. When applying data distillation, the dataset size was reduced first to 16 samples and then to 240 samples representing compression ratios of 0.06% and 1% of the original data, respectively, while still maintaining acceptable accuracy levels for both cases.

Table 3.3: Final Results

Test	Test Accuracy	Technique Used	Number of Training Samples	Samples per Class
Ordinary	97.49%	None	24,000	6,000
First Test	64.95%	HBDD	16	4
Second Test	82.35%	HBDD	240	60

These findings highlight the strong potential of HBDD for reducing both the storage and computational requirements associated with ML training. The reduced datasets not only minimized memory usage but also shortened training duration, since fewer epochs and steps per epoch were required to reach convergence. The approach therefore offers practical advantages for environments where local storage and processing resources are constrained. Overall, the preliminary results are encouraging. They demonstrate that data distillation enables substantial dataset compression down to 1% of the original size while preserving much of the classification

performance. This confirms that only a limited subset of representative samples is sufficient for effective training, reducing the need to retain or transmit large datasets. These outcomes validate the feasibility of the proposed method and provide a solid foundation for the next stage of this research: the implementation of real-time, on-board training MCUs. Future experiments will focus on evaluating the execution time and energy requirements of on-device distillation.

3.5 Distillation on MNIST and Application Potential in Agri-Food

3.5.1 Purpose and Experimental Setup

The second experiment aimed to extend the data distillation concept from an offline environment to a real-time, on-device implementation. The purpose was to evaluate whether dataset reduction could be achieved directly on an embedded device while maintaining efficient communication, memory management, and acceptable model performance. This experiment represents a practical step toward enabling embedded AI systems capable of learning and updating locally, with minimal data transfer and processing overhead. The experimental setup consisted of two main hardware components: a Raspberry Pi (RPi) 4 Model B and an M5C+. The RPi acted as the data provider and receiver, responsible for sending the dataset to the MCU, collecting the distilled data, and storing it for subsequent model training. The M5C+, was selected as the IoT node for performing the distillation process due to its low power consumption, compact form factor, and integrated connectivity features. This configuration provided an ideal environment to test real-time data processing under constrained computational and memory resources. As shown in Fig. 3.9 from [107], the system architecture was built around the MQTT communication protocol, which managed data flow between the RPi and the M5C+. The process began with the RPi transmitting the MNIST dataset row by row over MQTT. Each row was divided into smaller chunks to optimize network efficiency and ensure reliable transmission. The M5C+ subscribed to the data topic, reconstructed the received rows, and performed on-device distillation by averaging groups of images according to a parameter called *Dist_Value*. After each distillation cycle, the M5C+ published the resulting distilled row to a separate MQTT topic, where the RPi collected and stored it as a CSV file. This iterative exchange continued in real time, simulating a continuous data collection and reduction process as would occur in an IoT network. To ensure effective memory utilization on the M5C+, an overwrite strategy was implemented in which previously processed data were replaced with newly received rows. This design allowed the system to operate continuously without exceeding the

device's storage limits. The distillation and transmission code on the M5C+ was written in C and deployed via PlatformIO in Visual Studio Code, while Python scripts running on the RPi handled data streaming, reception, and subsequent model training. This configuration effectively replicated a real-world IoT setting where data are continuously collected, distilled, and transmitted for centralized analysis. By validating the feasibility of performing dataset distillation directly on a MCU, this experiment established a crucial link between theoretical data reduction techniques and their practical deployment on embedded systems, laying the foundation for scalable, energy-efficient, and privacy-preserving AI applications in agri-food and other sensor-driven domains. During this phase, several additional techniques were also explored to determine whether the performance of on-device ABDSD could be further improved. In particular, we evaluated the integration of dimensionality-reduction and signal-processing methods such as Principal Component Analysis (PCA) [111], Power Spectral Density (PSD) [112], Histogram-Based (CLAHE) [113] and other lightweight feature-transformation strategies. The intention was to enhance the representativeness of distilled samples or reduce their noise before transmission. However, experimental results showed that combining these techniques with ABDSD either increased computational overhead on the MCU or produced distilled samples that were less informative for downstream training. Among all tested configurations, the simplest version raw ABDSD without any additional transformation consistently provided the most stable and accurate results while preserving real-time performance. This finding confirmed that, under resource constraints, maintaining a minimal and non-intelligent distillation pipeline yields the best trade-off between efficiency and learning quality.




Fig. 3.9: Testbed for ABDSD

3.5.2 Data Transmission and Processing

This section outlines the communication protocol, data handling mechanisms, and on-device processing strategy employed to enable real-time dataset distillation between the RPi and the M5C+. The goal was to evaluate the efficiency of data transmission, segmentation, and on-device computation in a resource-constrained environment.

Communication Protocol

The data exchange between the RPi and the M5C+ was implemented using the MQTT protocol, chosen for its lightweight, publish–subscribe architecture well-suited to IoT systems. Three MQTT topics were used to structure the data flow: *CMD*, *MNIST*, and *DIST*. The M5C+ initially publishes its readiness status to the *CMD* topic, to which the RPi subscribes. Once the connection is established, the RPi publishes the MNIST dataset to the *MNIST* topic, where the M5C+ acts as a subscriber. After performing data distillation, the M5C+ publishes the processed data to the *DIST* topic, from which the RPi retrieves and stores the distilled elements. This structured topic hierarchy ensures robust, asynchronous communication and allows both devices to manage tasks independently while maintaining data synchronization.

Dataset Processing and Transmission Strategy

The MNIST dataset [114, 115], illustrated in Fig. 3.10, consists of 28×28 grayscale images of handwritten digits from 0 to 9. To enable transmission through MQTT under buffer constraints, each image (represented as a single CSV row of 784 pixel values) was split into multiple chunks of approximately 100 characters. These chunks were transmitted sequentially at a rate of one chunk per second, with the M5C+ reconstructing each row upon complete reception. Typically, 18–20 chunks were required per row, ensuring efficient memory use and minimizing packet loss.

On-device Dataset Distillation

Once the full dataset row was received, the M5C+ performed real-time data distillation directly on the device. The distillation process involved averaging pixel values across multiple aligned samples to produce a single representative image, reducing redundancy while preserving key visual features. The number of images averaged in each operation was determined by the parameter *Dist_Value*. In this study, two configurations were tested: *Dist_Value* = 3 and *Dist_Value* = 10. An example of distillation applied to Class 5 images is shown in Fig. 3.11 from [107], demonstrating how




Fig. 3.10: MNIST dataset used for real-time distillation experiments.

the averaging operation maintains digit contours while suppressing noise. To optimize on-device processing, the images were flattened into one-dimensional vectors of 784 elements and stored as `uint8_t` integers, reducing memory usage by a factor of four compared to floating-point representations. The M5C+ replaced older data after each distillation cycle to prevent memory overflow and maintain continuous operation. Table 3.4 summarizes the resource utilization and timing measurements recorded during execution. The system occupied 18.0% of the available RAM and 57.8% of flash memory, with each distillation cycle completed in 58 milliseconds after receiving 10 full rows ($Dist_Value = 10$). These results confirm the device's ability to handle real-time distillation efficiently within the constraints of an embedded IoT platform.




Fig. 3.11: Example of MNIST data distillation with $Dist_Value = 3$.

Table 3.4: Resource Utilization and Execution Time on the M5C+.

Resource	Usage	Task	Execution Time
Flash Memory	57.8%	One distillation process	58 ms
RAM	18.0%	One data chunk sent and received	1 s


Overall, the system achieved stable data transmission and efficient on-device processing. The combination of chunk-based transmission, lightweight integer representation, and real-time distillation proved effective in enabling continuous operation without exceeding memory or communication limits. This implementation demonstrates the feasibility of performing dataset reduction directly on IoT hardware, providing a foundation for scalable, real-time learning systems in agri-food and similar domains.

3.5.3 Model Architecture and Training

This section presents the NN model employed for classification, along with the training methodology used to evaluate the impact of different data distillation configurations. The objective was to analyze how varying degrees of data reduction (determined by the *Dist_Value*) influence model convergence and overall accuracy.

Model Architecture

A NN was used for all experiments. As illustrated in Fig. 3.12, the architecture consists of three main layers: an input layer, a hidden layer, and an output layer. The input layer receives 784 features, corresponding to the flattened 28×28 grayscale images of the MNIST dataset. The hidden layer contains 32 neurons with ReLU activation functions, while the output layer includes 10 neurons activated by the softmax function, each representing one of the digits from 0 to 9. This minimalistic structure was intentionally chosen to ensure lightweight computation and fast convergence, reflecting the type of model suitable for embedded AI applications. The model was optimized using the Adam algorithm and trained with the categorical cross-entropy loss function, which provided stable convergence and effective gradient updates during training.

**Fig. 3.12:** Architecture of the fully connected NN.

Training Dataset Configurations

To evaluate the influence of dataset size and compression, three dataset configurations were used, denoted as (A), (B), and (C), as summarized in Table 3.5.

- Structure (A) corresponds to the original dataset, comprising 50,000 training samples (5,000 per class).
- Structure (B) applies a *Dist_Value* of 10, meaning that every 10 consecutive samples were distilled into one representative sample, reducing the dataset from 50,000 to 5,000 samples (500 per class).
- Structure (C) applies a *Dist_Value* of 3, starting from 15,000 samples and distilling them into 5,000 representative samples (500 per class).

These configurations were selected to maintain balanced class distributions while allowing a controlled comparison between different degrees of data compression. Each configuration was carefully initialized to ensure that the number of distilled samples remained identical across experiments, thereby isolating the effect of distillation intensity from other variables.

Table 3.5: Dataset Configurations for Model Training.

Structure	Dist_Value	Original Samples	Distilled Samples
Original (A)	–	50,000	–
Distilled (B)	10	50,000	5,000
Distilled (C)	3	15,000	5,000

Training Procedure

The same model architecture was trained on each dataset configuration for 50 epochs. Across all three scenarios, the learning rate was fixed at 0.001, and the Adam optimizer was used for parameter updates. The training process demonstrated consistent and stable performance, with accuracy exceeding 99% and loss values remaining below 0.01 in every configuration. As expected, the training convergence was fastest for the distilled datasets (Structures B and C), as fewer samples reduced the number of steps per epoch. This observation highlights one of the key benefits of data distillation: faster model training with negligible loss in accuracy when representative information is preserved. The main training parameters are summarized in Table 3.6.

Table 3.6: Training Parameters and Performance Summary.

Parameter	Value / Observation
Optimization Algorithm	Adam
Learning Rate	0.001
Training Accuracy	> 99%
Training Loss	< 0.01

Overall, the results confirm that dataset distillation effectively reduces the data volume while preserving the statistical integrity required for accurate classification. The consistent performance across different configurations demonstrates the robustness of the proposed approach and supports its suitability for on-device AI applications where memory and energy constraints are critical.

3.5.4 Test Results

To evaluate the performance of the proposed system and assess the impact of ABDSD, three training scenarios were compared corresponding to the dataset structures (A), (B), and (C) described previously. Each trained model was tested on the MNIST test set comprising 10,000 images. Performance evaluation was conducted using confusion matrices, which provide a detailed visualization of classification results by comparing the predicted and actual classes. This allows an analysis not only of overall accuracy but also of how well the model distinguishes between similar digits. Figure 3.13 presents the confusion matrix for the model trained on the original dataset (Structure A). The model correctly classified 9,142 out of 10,000 test samples, achieving a test accuracy of 91.42%. As expected, this configuration served as the reference for subsequent comparisons. Fig 3.14 shows the results for the model trained on the distilled dataset with a *Dist_Value* of 10 (Structure B). Despite a ten-fold reduction in training data, the model maintained strong performance, correctly classifying 8,537 samples with an accuracy of 85.37%. Fig 3.15 displays the confusion matrix for the model trained with *Dist_Value* = 3 (Structure C). In this configuration, the model correctly classified 8,743 samples, reaching an accuracy of 87.43%. Notably, although the dataset was distilled more frequently, the accuracy remained close to that of the *Dist_Value* = 10 scenario, indicating that moderate compression can actually preserve class-level distinctions effectively. As summarized in Table 3.7 the model trained on distilled data maintained a high level of performance despite substantial reductions in dataset size. The test accuracy decreased by less than 7% compared to the original dataset, even when the number of samples was reduced by up to 90%. These results demonstrate that dataset distillation effectively preserves

the essential discriminative features of the data, allowing for efficient training under resource-limited conditions. The results also highlight that smaller *Dist_Value* settings can lead to slightly better generalization than larger ones, possibly because averaging over fewer samples maintains sharper class boundaries. This observation supports the hypothesis that properly tuned distillation ratios can balance information retention with computational efficiency. Overall, these findings confirm the feasibility of performing on-device dataset distillation in real time using lightweight hardware, achieving strong classification performance with a significantly reduced training load. The next subsection extends this concept to preliminary experiments in the agri-food domain, demonstrating the adaptability of the approach to visual classification tasks involving real agricultural products.




Fig. 3.13: Confusion matrix of the model trained using dataset structure A (original dataset).

Table 3.7: Summary of Test Results.

Structure	Dist_Value	Test Accuracy
A	–	91.42%
B	10	85.37%
C	3	87.43%




Fig. 3.14: Confusion matrix of the model trained using dataset structure B
($Dist_Value = 10$).




Fig. 3.15: Confusion matrix of the model trained using dataset structure C
($Dist_Value = 3$).

3.5.5 Preliminary Distillation in the Agri-Food Sector

Although dataset distillation has been widely examined using benchmark datasets such as MNIST, its application to agricultural image data remains largely unexplored. To extend our previous findings, we conducted a preliminary proof-of-concept experiment using a publicly available pistachio image dataset [116, 117, 118]. The goal was to assess whether the same data compression and feature-preserving principles could be applied to visual data of agri-food products. The

procedure involved converting the original RGB pistachio images into grayscale to simplify processing and reduce computational load. The dataset distillation algorithm was then applied to groups of images belonging to the same class to generate compact, representative samples. Fig 3.16 shows an example of this process for the Siirt pistachio class, where three images were combined into one distilled representation. The resulting image retained essential morphological traits such as contour, texture, and shading, while eliminating redundant details across samples. These results suggest that the proposed technique can summarize core visual features relevant to classification even when applied to complex, real-world imagery.




Fig. 3.16: Example of dataset distillation applied to pistachio images

3.5.6 Discussion

The experiments on both the MNIST and pistachio datasets collectively demonstrate that dataset distillation can significantly reduce data volume while preserving critical information for classification. In the MNIST case, a 90% reduction in dataset size led to only a 6% drop in test accuracy (Table 3.7), confirming that much of the discriminative structure of the data was retained. Moreover, the on-device implementation executed each distillation process in just 58 ms, validating the feasibility of real-time operation on resource-constrained microcontrollers such as the M5C+. When extended to agricultural imagery, the same principle proved effective at compressing visual data while maintaining identifiable product features. This finding indicates that dataset distillation can be generalized beyond benchmark datasets and adapted for real agricultural scenarios where bandwidth, storage, and computational resources are limited. The approach is especially relevant to IoT-based trace-

ability and quality-monitoring systems, where devices must process and transmit large amounts of image or sensor data efficiently. However, some limitations were also observed. The technique remains sensitive to spatial misalignment and inconsistent lighting conditions, which can blur distinctive features in the averaged images. This limitation is particularly relevant for agricultural datasets, where variations in object orientation and illumination are common. Next experiment will therefore focus on integrating lightweight alignment, normalization, and background-removal steps directly on the MCU to improve the robustness of the distillation process under variable conditions. Overall, these findings confirm that on-device dataset distillation offers a promising pathway toward compact, energy-efficient AI for the agri-food sector. By reducing data redundancy while maintaining accuracy, this technique can enhance both FL and centralized learning (CL) frameworks—lowering communication overhead, improving training efficiency, and enabling low-cost, real-time analysis on edge devices.

3.6 Full Pipeline on Grain Dataset: From Capture to Distillation

3.6.1 Purpose and Experimental Setup

Building upon the outcomes of the previous experiments, this stage implements the complete on-device dataset-distillation pipeline, addressing the earlier limitations related to image misalignment and inconsistent background conditions. The objective is to validate a fully functional workflow that performs lightweight preprocessing and dataset distillation directly on a low-power MCU for real agricultural data. The same hardware platform employed in the earlier tests in Section 3.5 of this chapter, was used here to ensure experimental continuity and enable a direct comparison with the MNIST and pistachio experiments. The pipeline, illustrated in Fig 3.17 from [108], combines two on-device operations, rotation correction and ADD, while the background removal stage was executed externally due to current hardware limitations. This modular design makes it possible to later integrate all stages within the same embedded device once lightweight segmentation models are developed. In this setup, a RPi 4 Model B acted as the data provider, transmitting preprocessed images via the MQTT protocol to the M5C+. Upon reception, the M5C+ reconstructed and aligned the images using the first received image as a reference, then performed real-time distillation based on a configurable parameter ($\text{Dist_Value} = 2, 3, 5, \text{ or } 10$). After each distillation cycle, the representative image was transmitted back to the RPi for storage and subsequent model training. These configurations corresponded to data-reduction ratios ranging from 50 % to 90 %, allowing the analysis of trade-offs between compression efficiency and classification accuracy. All timing and memory

statistics were measured on-board to evaluate computational feasibility within MCU constraints. The experiment demonstrates a unified and efficient pipeline capable of compacting visual data at the point of capture, bridging the gap between controlled testing and a deployable, IoT-oriented solution for agri-food monitoring.




Fig. 3.17: MNIST dataset used for real-time distillation experiments

3.6.2 Dataset Description, and Processing

Dataset Description

The dataset used in this experiment consists of RGB images of wheat grains, as illustrated in Fig 3.18 from [119]. It includes samples from two main wheat species which are hard and soft wheat, as reported in [119]. The images were acquired using four flatbed scanners under varying acquisition conditions, thereby reflecting natural variability in shape, size, orientation, and surface texture. Each image represents a dorsoventral view of a grain and exhibits both random and regular orientations, making the dataset suitable for evaluating alignment and feature-preservation techniques. For classification purposes, all sub-varieties were grouped into two broad classes corresponding to hard and soft wheat. The complete dataset contains 18,057 training and 4,073 testing images, with each sample properly labeled. Table 3.8 summarizes the number of samples per class.

Table 3.8: Number of samples used from each class in the grain dataset

Dataset	Classes	Training Samples per Class	Testing Samples per Class
Grain Dataset	2	4200	1800




Fig. 3.18: Examples of hard and soft wheat grains from the dataset.

Dataset Pre-processing

The preprocessing phase was designed as a hybrid pipeline distributed between the external data-provider device and the embedded M5C+ MCU to meet memory and computation constraints. The core stages included background removal, rotation correction, and ABDSD. To simplify on-device processing, image resizing to 128×128 pixels and grayscale conversion were performed externally. Although not part of the embedded workflow, these steps ensured input uniformity. The background removal stage was also executed externally since existing segmentation models exceeded the memory capacity of the M5C+. Nonetheless, this stage remains crucial for obtaining consistent and isolated object features. Future work will focus on developing a lightweight, task-specific segmentation model optimized for the M5C+ to achieve complete on-device processing. Following preprocessing, each image was transmitted to the MCU, where rotation correction and ABDSD were performed. Each grayscale image was flattened into a 16,384-element uint8_t vector for efficient handling. Averaging operations were performed in floating-point format and then rounded to integers to preserve compatibility.

Dataset Transmission and Row Assembly

To enable efficient and reliable data transfer between the external node and the M5C+, the MQTT protocol was adopted. Because of the limited payload size supported by MQTT, each image row was segmented into fixed-size text chunks of up to 100 characters, including comma separators. This segmentation strategy prevented buffer overflows and ensured stable communication within the M5C+ memory constraints. For 128×128 grayscale images, the complete data of one image corresponds to approximately 650 chunks. The M5C+ sequentially received and reassembled these chunks to reconstruct each image in memory before executing rotation correction and distillation. This design enables continuous and stable data transmis-

sion even under the tight resource limitations of embedded IoT devices, supporting scalable deployment for agricultural monitoring applications.

3.6.3 On-device Pipeline

The on-device pipeline integrates lightweight pre-processing and dataset distillation into a unified embedded workflow designed for real-time operation under hardware constraints. As illustrated in Fig 3.17, the pipeline consists of three main stages: background removal, rotation correction, and ABDSD. These stages work sequentially to isolate the object of interest, align all samples, and condense the data into representative images suitable for efficient on-device learning.

Background Removal

Before applying rotation correction and distillation on-device, the raw grain images undergo background removal to isolate the foreground objects and eliminate irrelevant visual noise. This step enhances consistency and quality across samples, improving the robustness of the subsequent operations. Commonly used deep-learning segmentation models such as YOLOv8 [120], DeepLabV3 [121], MODNet [122], and U²-Net [121] demonstrate excellent performance on large-scale datasets but are computationally demanding. Due to the limited memory and processing capabilities of the M5C+, deploying such models directly on-device is impractical. Among the tested alternatives, a lightweight implementation of DeepLabV3 could be deployed but yielded suboptimal segmentation results, often retaining background fragments and shadows that degraded subsequent alignment performance. As a result, this operation was performed externally on a RPi 4 or camera module before image transmission to the M5C+. Nevertheless, background removal remains a fundamental part of the proposed pipeline, as shown in Fig 3.17

Rotation Correction and Alignment

After background removal, the M5C+ performs rotation correction to ensure consistent image orientation, which is crucial for meaningful average-based distillation. Instead of relying on a predefined reference, the system dynamically uses the first received image as a baseline, adapting to real-time, unlabeled input streams. The MCU searches over rotation angles from 0° to 345° in 15° steps and tests both flipped and non-flipped versions. The best alignment is identified by minimizing pixel-wise absolute differences between the reference and candidate images using the following similarity score:

$$\text{MatchScore}(\theta, f) = 100 - \frac{1}{N \cdot 255} \sum_{i=1}^N |l_{\text{ref}}(i) - l_{\text{test}}^{\theta, f}(i)| \times 100$$

where l_{ref} denotes the reference image, $l_{\text{test}}^{\theta, f}$ represents the rotated and optionally flipped test image, $f \in \{0, 1\}$, and N is the number of overlapping pixels.

Rotation is applied about the image center as:

$$\begin{pmatrix} x' \\ y' \end{pmatrix} = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} x - x_c \\ y - y_c \end{pmatrix} + \begin{pmatrix} x_c \\ y_c \end{pmatrix}$$

If flipping is required, $x' = W - 1 - x'$, where W is the image width. Bilinear interpolation refines pixel values, minimizing distortion. The configuration that yields the highest match score is applied before distillation, ensuring robust alignment under uncontrolled imaging conditions.

Distillation on the Grain Dataset

Following alignment, the system performs ABDSD, which follows the same procedure described in Section 3.3. In this stage, groups of aligned images are averaged according to a configurable parameter, *Dist_Value*, to generate compact and representative samples. For this experiment, *Dist_Value* was set to 2, 3, 5, and 10, corresponding to data-reduction ratios of 50%, 66.7%, 80%, and 90%, respectively. The process was executed directly on the M5C+ after rotation correction, producing distilled images that maintain the essential visual characteristics of the original dataset while significantly reducing redundancy and transmission overhead. When alignment is accurate, the distilled results retain the key morphological features of the grains, whereas misalignment or inconsistent lighting can blur edges and distort contours, as illustrated in Figs 3.19 and 3.20 from [108].




Fig. 3.19: Distilled soft wheat samples obtained without alignment using $\text{Dist_Value} = 3$




Fig. 3.20: Distilled soft wheat samples obtained after alignment using `Dist_Value = 3`. Key features remain visible despite data reduction.

This step completes the on-device processing workflow, enabling real-time and memory-efficient image condensation directly on the embedded hardware for subsequent model training.

3.6.4 Model Architecture and Training

This section describes the design of the classification model, detailing its architecture, training configurations, and dataset structures used for evaluating the impact of dataset distillation on performance.

3.6.5 Model Architecture

A NN was designed for the binary classification of wheat grains into hard and soft classes. The architecture is composed of an input layer, three convolutional layers, two max-pooling layers, and two fully connected (dense) layers. The first convolutional layer applies 16 filters of size 3×3 , followed by a ReLU activation function. The second convolutional layer increases the number of filters to 32 and is followed by a max-pooling layer to reduce the spatial dimensions. The third convolutional layer further increases the number of filters to 64, followed by another max-pooling layer to extract deeper spatial features. The resulting feature maps are flattened into a one-dimensional vector and passed to a dense layer with 64 neurons and a ReLU activation function. The final dense layer outputs two neurons corresponding to the two target classes. A softmax activation function is used in the output layer to produce the final classification probabilities.

3.6.6 Model Training Structures

To assess the influence of data compression on model accuracy, the CNN was trained on datasets generated with different distillation parameters. The *Dist_Value* parameter defines the number of images averaged into a single representative sample. Increasing this value reduces dataset size while enhancing compactness. Table 3.9 summarizes the different dataset structures derived from the original dataset. When *Dist_Value* = 2 (Structure A), 4,200 samples were distilled from the initial 8,400 images. Increasing *Dist_Value* to 3 (Structure B) yielded 2,800 samples, while setting it to 5 (Structure C) produced 1,680 samples. The most compact version (Structure D), with *Dist_Value* = 10, contained only 840 distilled images. These variations allowed controlled analysis of the trade-off between dataset reduction and classification performance.

Table 3.9: Dataset structures generated using different *Dist_Value* settings

Structure	Dist_Value	Number of Distilled Samples
A	2	4,200
B	3	2,800
C	5	1,680
D	10	840

—

3.6.7 Model Training

The same CNN architecture was trained separately on each dataset structure for comparison. Each experiment was run for 50 epochs using the Adam optimizer with a learning rate of 0.001. The categorical cross-entropy loss function was used to guide optimization. In all configurations, the model achieved a training accuracy exceeding 90% and maintained a loss value below 0.1, confirming stable convergence even on highly reduced datasets. The original, non-distilled dataset served as a baseline, while subsequent experiments employed distillation with *Dist_Value* = 2, 3, 5, and 10.

3.6.8 Results and Discussion

To evaluate the performance of the proposed pipeline, a separate test set of 3,600 grain images was used. These samples were preprocessed in the same way as the training data but were not included in the distillation stage to ensure an unbiased evaluation of generalization capability. The results summarized in Table 3.10

demonstrate that the model maintained a high level of accuracy even with significant dataset reduction. When trained on the full, non-distilled dataset, the model achieved a test accuracy of 94.50%. With 50% and 66.67 reduction ratios achieved via ADD, the model retained strong accuracy at 93.75% and 92.19%, respectively. However, further reductions to 80% and 90% led to noticeable performance degradation, with accuracies of 82.94% and 78.47%. These results confirm that moderate levels of dataset distillation can substantially reduce data size while maintaining comparable accuracy. In contrast, excessive reduction leads to the loss of discriminative details necessary for reliable classification. Hence, the optimal trade-off between data efficiency and predictive performance lies within the 50–66% reduction range.

Table 3.10: Comparison of test accuracies on the Grain Dataset under different distillation ratios

Dataset	Data Reduction Ratio (%)	Test Accuracy (%)
Original Dataset (No Distillation)	0	94.50
Average Grain Data Distilled (A)	50.00	93.75
Average Grain Data Distilled (B)	66.67	92.19
Average Grain Data Distilled (C)	80.00	82.94
Average Grain Data Distilled (D)	90.00	78.47

The experimental findings highlight that ABDSD provides an efficient means of balancing model performance with resource constraints. By averaging spatially aligned images, the process successfully removes redundancy while retaining representative grain features such as contour, texture, and orientation. The resulting distilled datasets enable faster training and lower storage and communication requirements—an essential advantage for embedded learning systems deployed in field conditions. Beyond this experiment, the pipeline comprising background removal, rotation correction, and data distillation can be generalized to a variety of agricultural imaging applications. Potential use cases include:

- Ripeness estimation in fruits and vegetables, where color and texture evolution can be tracked efficiently through distilled data.

- Fungal or mold detection, enabling low-cost early warning systems for post-harvest storage.
- Variety classification in crops such as grapes, olives, or rice, where high intra-class variability can be managed through alignment and averaging.

Furthermore, distilled images preserve key morphological and color features relevant to quality assessment and traceability. This makes the proposed approach well suited for real-time decision-making, quality monitoring, and certification tasks on embedded IoT devices operating within agri-food supply chains. The adaptability of this method opens opportunities for fully autonomous systems that perform lightweight preprocessing and inference directly at the edge, minimizing reliance on cloud computing and improving responsiveness in field conditions.

3.7 Conclusion

This chapter presented a progressive exploration of HBDD and ABDSD as a lightweight data reduction strategies to enable ML directly on embedded devices. Through three successive experiments, the concept evolved from offline validation on structured data to a fully integrated on-device implementation applied to agricultural imagery. The first stage established the feasibility of dataset distillation in a controlled computing environment using the an IMU dataset. In this phase, a HBDD approach was employed to investigate whether representative values extracted from dominant histogram bins could effectively summarize the dataset. These initial tests demonstrated that even extreme data compression-down to less than 1% of the original dataset can preserve meaningful learning performance when statistically representative samples are distilled. The second stage implemented real-time, on-device ABDSD using the MNIST dataset on the M5C+. This experiment validated the capability of low-power MCUs to execute ABDSD autonomously, achieving up to 90% data reduction while maintaining test accuracy above 85%. During this phase, several additional techniques including PCA, PSD, and other lightweight feature-transformation strategies were evaluated in combination with ABDSD to improve the representativeness of distilled samples. However, these methods either increased computational overhead on the MCU or degraded the quality of the distilled data. The simplest configuration, relying solely on ordinary ABDSD without auxiliary transformations, consistently produced the best balance between accuracy, efficiency, and real-time feasibility. The communication and memory management design, based on the MQTT protocol and chunk-wise data handling, proved effective for continuous embedded operation under limited hardware resources. The third and final stage extended the approach to the agri-food domain through a complete

on-device pipeline applied to a custom wheat grain dataset. The system integrated three main stages background removal, rotation correction, and average dataset distillation executed either externally or directly on the MCU. Experimental results confirmed that moderate data reduction levels (50–66%) preserved classification accuracy above 92%, while higher compression ratios led to performance degradation. These findings underline the trade-off between data compactness and predictive reliability in embedded image classification. Collectively, the results demonstrate that ABDSD and HBDD provide simple, hardware-efficient, and generalizable solutions for on-device learning and communication-efficient data handling. By condensing redundant visual or numerical data at the source, the proposed methods enable scalable and energy-conscious AI applications across distributed IoT networks. In the context of the agri-food sector, these approaches offer a practical path toward autonomous traceability, quality assessment, and contamination detection using low-cost, resource-constrained devices. Future work will focus on integrating the entire pipeline including lightweight background segmentation within a single embedded platform. Achieving this goal would result in a fully self-contained, real-time system capable of performing both preprocessing and inference directly at the edge, paving the way for next-generation intelligent and sustainable agricultural IoT systems.

On-Device Edge AI Solutions for the Agri-Food Sector

4.1 Introduction

The agri-food sector is increasingly challenged by the need to ensure food safety, quality consistency, and regulatory compliance across complex and distributed production chains. Factors such as environmental variability, biological processes, and climate change introduce significant uncertainty in both primary production and food processing stages. In this context, timely access to reliable information is essential to support traceability, certification, and risk mitigation strategies. Conventional monitoring and quality assessment methods often rely on laboratory-based analyses or centralized data processing, which, although accurate, are typically costly, time-consuming, and unsuitable for real-time or large-scale deployment [123]. These limitations have motivated the adoption of intelligent monitoring systems that combine sensing technologies with AI and the IoT to enable continuous, automated decision-making directly within the agri-food chain [124, 13, 125, 126, 127, 128, 129, 130, 131]. Recent advances in embedded systems and ML have accelerated the development of on-device AI, also referred to as Edge AI, where data processing and inference are performed locally on resource-constrained devices such as MCUs, embedded cameras, or low-power single-board computers. Compared to cloud-centric approaches, on-device AI reduces communication latency, limits data transmission requirements, enhances data privacy, and improves system robustness in environments with limited or intermittent connectivity. These characteristics make Edge AI particularly suitable for agri-food applications, where monitoring often takes place in farms, or production facilities operating under strict energy, cost, and infrastructure constraints. Within the agri-food domain, on-device AI enables the real-time assessment of heterogeneous phenomena, ranging from environmental risk factors in crop production to process-state estimation in food transformation. One prominent example concerns mycotoxin contamination, which remains a major food safety issue worldwide. Mycotoxins are toxic

secondary metabolites produced by filamentous fungi that contaminate a wide range of food products [132, 133, 134, 135, 136]. Among them, aflatoxins and ochratoxins are considered the most hazardous due to their cumulative toxicity and carcinogenic potential [123, 137, 138, 139, 140, 141, 142]. OTA, mainly produced by *Aspergillus* and *Penicillium* species [143, 144, 145, 146, 147, 148], is of particular concern in viticulture, where its development is strongly influenced by environmental conditions such as temperature and humidity [141, 142, 149, 150, 151, 152]. While laboratory-based detection techniques such as High-Performance Liquid Chromatography or ELISA provide high analytical accuracy, they are unsuitable for continuous, in-field monitoring [123], highlighting the need for predictive, sensor-driven AI solutions deployable directly at the edge. Beyond contamination risk assessment, on-device AI also plays a crucial role in monitoring and controlling agri-food transformation processes, such as fermentation. Fermentation is a fundamental bioprocess used in the production of beverages, dairy products, and plant-based foods, yet its control remains challenging due to biological variability and dynamic environmental conditions. Traditional monitoring approaches often depend on manual sampling and offline analysis, limiting their effectiveness for real-time quality assurance. By integrating sensors, computer vision, and lightweight neural network models, Edge AI systems enable continuous classification of fermentation stages directly on embedded platforms, supporting autonomous process supervision without reliance on cloud infrastructure. Accordingly, this chapter presents two representative on-device AI solutions for the agri-food sector: (i) an embedded AI framework for real-time OTA risk assessment in viticulture, progressing from conceptual modeling to validated on-site deployment, and (ii) a lightweight Edge AI vision-based system for fermentation stage classification, designed for execution under strict hardware and energy constraints. Together, these solutions illustrate how common design principles such as low-cost sensing, lightweight model architectures, and real-time inference on embedded devices can be adapted to different agri-food scenarios. The presented approaches demonstrate the potential of on-device AI to enhance traceability, certification, and decision-making in modern agri-food systems, effectively bridging the gap between laboratory-level analysis and practical, field-deployable implementations.




Fig. 4.1: Aflatoxin




Fig. 4.2: Ochratoxin

Fig. 4.3: Microscopic images of fungi responsible for Mycotoxin production.

4.2 AI-Enhanced IoT for Ochratoxin A Detection in Grape Production

4.2.1 State of the Art and Motivation for On-Device OTA Risk Assessment

Mycotoxins, and particularly OTA, represent a major challenge to food safety due to their high toxicity and frequent occurrence in agricultural commodities such as cereals, coffee, and grapes. OTA formation is strongly influenced by environmental conditions, especially temperature, humidity, and harvesting time, which create favorable circumstances for fungal proliferation and toxin accumulation. Consequently, reliable and continuous monitoring strategies are essential for the early detection and prevention of OTA contamination throughout the agri-food chain. Conventional OTA detection methods are primarily based on chemical and instrumental analyses conducted in laboratory settings. High-Performance Liquid Chromatography (HPLC) coupled with fluorescence detection remains the reference technique for OTA quantification. As demonstrated by Kong *et al.* [153], HPLC provides excellent sensitivity and reproducibility when analyzing OTA in complex food matrices. Similar conclusions were reported in [154], which confirmed the analytical accuracy of HPLC while highlighting its limited suitability for large-scale or real-time monitoring due to the need for specialized equipment and trained personnel. Liquid Chromatography–Mass Spectrometry (LC–MS) has also been investigated as a complementary approach, offering enhanced specificity and lower detection limits. In [155], LC–MS combined with electrospray ionization enabled trace-level OTA detection, proving valuable for confirmatory analyses. However, the high cost and infrastructural requirements of LC–MS systems restrict their applicability to centralized laboratories. Alternative low-cost techniques such as Thin-Layer Chromatography (TLC) have been proposed for preliminary screening purposes. The study in [156] demonstrated the practicality of TLC in resource-constrained contexts, although its limited quantitative accuracy prevents its use for regulatory compliance and precise risk assessment. To overcome the constraints of laboratory-based methods, recent

research has explored AI- and IoT-driven approaches for real-time mycotoxin monitoring. Qu *et al.* [157] proposed a low-cost intelligent system combining an Artificial Olfactory Sensor (AOS) with a MobileNetV3 deep-learning model, achieving over 98.5% accuracy in detecting contaminated maize with sub-second inference times. Their results illustrate the potential of compact NN to replace complex sensor arrays while maintaining high classification performance. Zhang *et al.* [158] presented a hybrid framework integrating AI algorithms with biosensor technologies for food quality assessment, highlighting the advantages of multimodal sensing but also noting challenges related to sensor calibration and robustness under varying environmental conditions. Noor *et al.* [159] developed a portable electrochemical biosensor coupled with an IoT-based cloud platform for multi-mycotoxin detection, including OTA and Aflatoxin B1. Although high analytical precision was reported, the system's reliance on continuous network connectivity limits its deployment in rural or low-infrastructure environments. Akhtar *et al.* [160] adopted a different perspective by focusing on the monitoring of environmental indicators, such as CO₂ concentration and precipitation, to identify conditions favorable to fungal growth. Their Arduino-based IoT system enabled real-time surveillance but faced challenges related to energy efficiency and data transmission scalability. More recently, Kim *et al.* [161] introduced a Compact Convolutional Transformer (CCT) model for wheat contamination classification based on respiration and mold development indicators, achieving an accuracy of 83.33%. While demonstrating the feasibility of deep-learning-based risk detection, their results revealed limitations in distinguishing early contamination stages, underscoring the need for more sensitive and application-specific models. Despite these advances, most existing AI-based solutions either rely on laboratory-grade analytical measurements or depend on cloud-based processing and stable connectivity, which limits their applicability for autonomous, in-field deployment. Moreover, few studies explicitly address the constraints imposed by low-power embedded platforms when designing predictive models for OTA risk assessment. In this context, the present work develops an on-device AI framework for OTA risk prediction through two complementary studies. The first study [162] serves as a proof of concept, investigating the feasibility of predicting OTA contamination risk using a synthetically generated dataset derived from temperature, humidity, and harvesting duration parameters. This preliminary phase aimed to validate whether environmental variables alone could capture the nonlinear dynamics associated with OTA formation, without relying on chemical measurements. The encouraging outcomes confirmed the suitability of neural networks for modeling OTA-related risk patterns and provided essential insights for subsequent system design. Building on this methodological validation, the second study [96] constitutes the core contribu-

tion of this work and focuses on a chemically validated, real-world deployment. In this phase, experimental data obtained from grape samples analyzed through chemical methods were used to train and evaluate multiple ML models, including Random Forests, Support Vector Machines (SVM), k-Nearest Neighbors, Gradient Boosting Machines, and CNN. Although ensemble-based methods achieved competitive accuracy, their memory footprint and computational requirements rendered them unsuitable for MCU-based execution, in agreement with observations reported in [163]. CNN-based models, when optimized through transfer learning, offered a favorable compromise between accuracy, compactness, and inference latency. Consequently, the CNN architecture was selected and integrated into a real-time IoT monitoring system, where environmental data acquired from low-cost sensors were processed locally on an (M5C+ Technology Co., Ltd., Shenzhen, China) MCU. This implementation validated the feasibility of on-device inference for OTA risk prediction under real agricultural conditions and demonstrated that lightweight AI models can support scalable, autonomous, and energy-efficient food safety monitoring directly at the edge. Building on the methodological insights obtained from the proof-of-concept study, the same overall design philosophy was retained while introducing several key refinements to enable a practical, deployable solution. In this phase, the system architecture was redesigned to better align with the constraints and requirements of real-world agricultural deployment. In particular, a more suitable MCU was selected to support reliable on-device inference under limited computational and energy resources. Furthermore, the classification strategy was revised to conform to official EU regulations governing maximum allowable OTA concentrations in grape juice and wine production. This ensured that the proposed risk assessment framework is not only technically sound but also standardized and compliant with regulatory thresholds relevant to food safety and certification.

4.2.2 Proposed System

Based on the proof-of-concept framework previously developed in [162] to demonstrate the feasibility of OTA risk prediction using environmental data, this study proposes an enhanced on-device AI architecture designed for real-time operation and validated against chemically analyzed samples. Unlike the earlier exploratory phase, where an NFC-enabled system was employed to locally store and retrieve OTA growth predictions derived from synthetic data, the present architecture integrates a more capable embedded platform and adopts an IoT-oriented communication strategy. As illustrated in Fig 4.4 from [96], the system workflow begins with the generation of a training dataset grounded in chemical analysis of grape samples, conducted in accordance with EU regulations. This dataset provides a val-

idated reference for associating environmental conditions with OTA contamination risk. The resulting data were used to train a NN model offline, after which the trained model was optimized and deployed on an M5C+ MCU. The embedded platform is equipped with an SHT30 temperature and humidity sensor, enabling continuous acquisition of environmental parameters directly within the harvesting context. During operation, the M5C+ performs on-device inference to classify the likelihood of OTA contamination based on real-time sensor readings. In contrast to the proof-of-concept implementation, where inference results were stored locally via NFC, the proposed system transmits both raw sensor data and classification outcomes to a remote server using the lightweight MQTT protocol. The server records the received information in CSV format, facilitating data visualization, long-term storage, and further analysis. This transition from a localized NFC-based interaction model to an IoT-based communication framework significantly enhances system scalability and interoperability, enabling seamless integration into broader traceability and certification infrastructures. The following subsections present the OTA risk assessment framework in a structured manner. First, the analytical evaluation and labeling function used to relate environmental conditions to OTA contamination risk is introduced (Section 4.2.3). Then, the hardware platform, software workflow, and dataset construction adopted for on-device deployment are described (Section 4.2.4). Subsequently, the development and optimization of the lightweight AI models are detailed (Section 4.2.5). The experimental protocol and performance evaluation are presented next (Section 4.2.6), followed by a comprehensive analysis and discussion of the obtained results (Section 4.2.7). Finally, the section concludes with a summary of the main findings and implications of the proposed OTA solution (Section 4.2.8).




Fig. 4.4: On-device OTA risk assessment system architecture.

4.2.3 Evaluation Function

Evaluation Function for OTA Contamination in Grapes

To establish a chemically validated and application-oriented basis for the predictive model developed in this work, an analytical feasibility study was conducted on white grape samples. The objective of this investigation was to extract, quantify, and correlate OTA concentrations with environmental parameters, thereby defining an evaluation and labeling function suitable for training and validating an on-device AI model. Unlike purely data-driven or synthetic approaches, this step ensures that the predictive framework is grounded in experimentally measured contamination levels and aligned with regulatory assessment practices. OTA extraction was performed according to the VICAM immunoaffinity protocol [164]. Briefly, 5 mL of grape juice or pulp were mixed with 10 mL of an H_3PO_4 –NaCl solution and homogenized for 1 minute. Subsequently, 5 mL of chloroform were added, and the mixture was centrifuged at $2500\times g$ for 15 minutes. The organic phase was collected, and the extraction procedure was repeated to improve analyte recovery. The combined organic extracts were dried and reconstituted in phosphate-buffered saline (PBS) containing 15% methanol before loading onto an Ochratest immunoaffinity column (VICAM). After washing and drying, OTA was eluted with methanol, evaporated under nitrogen, and finally reconstituted in 50 L of mobile phase prior to chromatographic analysis.

HPLC Determination

OTA quantification was carried out using a Shimadzu Nexera UHPLC-RF system (Shimadzu, Kyoto, Japan) equipped with an autosampler, column oven, and fluorimetric detector. Chromatographic separation was achieved using a Phenomenex Luna C18 column (5 m, 150×4.60 mm). The mobile phase consisted of acetonitrile:water:acetic acid (99:99:2, v/v/v), delivered at a flow rate of 0.8 mL/min. Fluorescence detection was performed at an excitation wavelength of 333 nm and an emission wavelength of 477 nm. OTA concentrations were determined using the external standard calibration method, with chromatographic data processed using LC Solution software (version 5.71, Shimadzu).

Method Optimization

The calibration curve was constructed using the least-squares regression method, and linearity was verified through Mandel's test over the concentration range of 0.25–10 g/L. The limits of detection (LoD) and quantification (LoQ) were estimated

using the standard deviation approach, resulting in values of 0.08 g/L and 0.12 g/L, respectively. An overall instrumental recovery rate of 94% was obtained, confirming the robustness and suitability of the analytical method for generating reliable labels to support the proposed AI-based risk assessment framework.

Results and Definition of the Evaluation Function

The analytical feasibility study enabled the formulation of an evaluation function that relates OTA concentration to specific combinations of environmental parameters and harvesting duration. Fig 4.5 from [96] illustrates the temporal evolution of OTA levels under varying temperature and relative humidity conditions. For instance, environmental conditions of 28 °C and 70% relative humidity frequently observed during nighttime harvesting were associated with a rapid increase in OTA concentration over a six-hour interval. Conversely, harvesting conditions characterized by higher temperature (35 °C) and lower humidity (35%) resulted in substantially reduced OTA accumulation, indicating a more favorable scenario for minimizing contamination risk [162]. These experimentally derived relationships provide a chemically grounded basis for labeling environmental data according to OTA risk levels. As a result, temperature, humidity, and harvesting duration were selected as the primary input features for the on-device predictive model. By relying on analytical measurements rather than synthetic assumptions, the proposed evaluation function strengthens the reliability of the AI-based classification system and ensures consistency with regulatory and food safety requirements [96].




Fig. 4.5: Evaluation function relating OTA concentration to time under different temperature and humidity conditions.

4.2.4 Hardware, Software, and Dataset Description

In the validated OTA risk assessment study, an MCU was selected to meet the computational, energy, and communication requirements of real-time on-device inference in agricultural environments. The dataset used for model training and validation was generated according to the evaluation function described in Section 4.2.3, which is grounded in chemical analysis of grape samples and aligned with European Union regulatory thresholds for OTA contamination. This section describes the hardware platform, software workflow, and dataset structure adopted to support the proposed on-device AI framework.

Hardware

The M5C+, shown in Fig 4.6 from [96], was selected as the core IoT device for this study due to its compact form factor, portability, and suitability for rapid prototyping and embedded AI deployment. The device integrates an ESP32 MCU, a built-in display, and multiple sensors within a single unit, providing sufficient computational resources for real-time environmental data acquisition and on-device model inference [165]. In addition, the M5C+ supports Wi-Fi connectivity and is compatible with lightweight communication protocols such as MQTT, enabling efficient transmission of data and inference results to a remote server [166]. The platform benefits from the M5C+ ecosystem, which offers extensive documentation, community support, and compatibility with widely used development environments, including VS Code, Python, and Arduino [167]. The main technical characteristics of the device are summarized in Table 4.1. Low-cost and easy-to-integrate sensors bundled with the M5C+ were exploited to enable practical deployment in harvesting conditions. Specifically, the built-in 6-axis IMU was used as a trigger mechanism to initiate data acquisition, while the ENV-III Hat integrates an SHT30 digital temperature and humidity sensor (Sensirion AG, Staefa ZH, Switzerland) to capture environmental parameters relevant to OTA development. In operational use, the M5C+ is placed inside or attached to a grape harvesting box. When the box is tilted upside down for more than five seconds, as detected by the accelerometer, the system automatically initiates the recording of temperature and humidity values at 60-second intervals. This sampling rate represents a compromise between capturing sufficient temporal resolution for reliable modeling and minimizing memory usage and energy consumption. The onboard memory capacity and processing speed of the M5C+ allow the embedded classification model to operate efficiently, supporting real-time inference and wireless data transmission under field conditions.



Fig. 4.6: M5Stick C Plus (M5C+) device used for environmental data acquisition and on-device OTA classification.

Table 4.1: Key characteristics of the M5Stick C Plus (M5C+).

Feature	Description
Microcontroller	ESP32-PICO-D4 (240 MHz dual-core, 600 DMIPS)
Flash Memory	16 MB
RAM	8 MB (520 KB SRAM)
Display	0.96-inch TFT color display (80 × 160 resolution)
Battery	95 mAh Li-Po battery
Connectivity	Wi-Fi (802.11 b/g/n), Bluetooth 4.2 BLE
Expansion	Grove-compatible connector
Sensors	Accelerometer (MPU6886), IR transmitter
Input/Output	Built-in buttons (2), microphone, speaker
Interfaces	USB-C for charging and programming
Dimensions	48 mm × 24 mm × 13.5 mm
Weight	Approximately 14 g

Software

The proposed system adopts a hybrid software workflow that combines Python-based model development with C/C++ embedded deployment to support the complete pipeline from data preparation to on-device inference. Python was used for data preprocessing, model training, and quantization, leveraging cloud-based environments such as Google Colab to facilitate experimentation and hardware-agnostic optimization. Once trained, the selected model was quantized and converted to a lightweight TensorFlow Lite (TFLite) format, reducing memory footprint and inference latency to meet MCU constraints. For embedded implementation, the firmware was developed using PlatformIO, relying on C/C++ libraries compatible with the

ESP32 architecture. The quantized TFLite model was converted into a static C array and integrated directly into the firmware, enabling fully autonomous on-device inference without external dependencies. This end-to-end workflow from cloud-based training to embedded execution ensured fine-grained control over model size, computational load, latency, and power consumption. Compared to deterministic rule-based approaches, the integration of machine learning models on MCUs offers significant advantages for complex and dynamic phenomena such as OTA risk prediction. ML models are capable of capturing nonlinear relationships between environmental variables, adapting to new data distributions, and maintaining robustness in the presence of sensor noise. These properties make them particularly well suited for deployment in resource-constrained agricultural environments, where flexibility, scalability, and energy efficiency are critical.

Dataset Description

The dataset was designed to capture the relationship between harvesting conditions and OTA contamination risk, as defined by the evaluation function presented in Section 4.2.3. To reflect realistic operational scenarios, synthetic environmental data were generated within ranges derived from both chemical analysis and local climatic observations. Specifically, temperature values ranged from 28 °C to 35 °C, while relative humidity values varied between 35% and 70%, corresponding to typical harvesting conditions in the village of Bianco, Reggio Calabria, Italy. These ranges are consistent with the boundaries identified through the quantitative feasibility analysis, ensuring coherence between experimental measurements and model inputs. To emulate real harvesting practices, environmental conditions were associated with randomly generated time intervals spanning 30 to 360 minutes, reproducing realistic harvesting durations. This temporal component enables the model to capture the cumulative effect of environmental exposure on OTA development, providing a practical basis for contamination risk assessment. The resulting dataset includes the following features:

- TimeIn: Start time of the harvesting interval.
- TimeOut: End time of the harvesting interval.
- Tav: Average temperature over the interval.
- Hav: Average relative humidity over the interval.
- Alpha: Derived parameter combining environmental and temporal factors.
- Label: OTA risk class assigned according to the evaluation function.

Compared to the proof-of-concept study presented in [162], where minimum, maximum, and average values of temperature and humidity were treated as separate

inputs, the present dataset adopts a more compact representation based on average values (T_{avg} and H_{avg}). This simplification, combined with labeling grounded in chemical analysis rather than synthetic assumptions, results in a concise yet reliable dataset structure that is well suited for training lightweight models intended for on-device deployment.

Dataset Labeling Function

To enable supervised training of the proposed on-device AI model, a dataset labeling function was defined to estimate OTA concentration as a function of harvesting duration, temperature, and relative humidity. This function serves as a chemically grounded surrogate model, derived from analytical feasibility studies, and provides a systematic mechanism for assigning contamination risk labels to environmental data. The labeling process consists of two main stages: (i) estimation of OTA concentration using a growth-based model parameterized from experimental measurements, and (ii) assignment of discrete labels according to regulatory and risk-oriented concentration thresholds.

OTA Concentration Estimation

OTA concentration was estimated using an exponential growth formulation combined with weighted interpolation across multiple reference environmental conditions. Experimental OTA production curves obtained from analytical measurements under specific temperature and humidity settings (Fig 4.5) were used to parameterize the model. For a given reference condition, OTA concentration as a function of time is expressed as:

$$C_{\text{condition}}(t) = a \cdot e^{b \cdot t} + c \quad (4.1)$$

where:

- t denotes the harvesting duration in minutes;
- a , b , and c are curve-fitting parameters derived from chemical analysis;
- T and H represent temperature and relative humidity, respectively.

To generalize the estimation across a continuous range of environmental conditions, a weighting mechanism was introduced based on the proximity of a given sample's temperature T and humidity H to the experimentally defined reference values T_{ref} and H_{ref} . The weight associated with each reference condition is defined as:

$$w_{\text{condition}}(T, H) = \frac{\exp\left(-\frac{(T-T_{\text{ref}})^2}{10} - \frac{(H-H_{\text{ref}})^2}{100}\right)}{\sum_{\text{condition}} \exp\left(-\frac{(T-T_{\text{ref}})^2}{10} - \frac{(H-H_{\text{ref}})^2}{100}\right)} \quad (4.2)$$

The estimated OTA concentration for a given environmental condition and time interval is then computed as a weighted sum:

$$C(t, T, H) = \sum_{\text{condition}} w_{\text{condition}}(T, H) \cdot C_{\text{condition}}(t) \quad (4.3)$$

This formulation enables smooth interpolation between experimentally measured conditions while preserving consistency with analytically observed OTA growth dynamics.

Reference Environmental Conditions

The labeling function relies on four reference environmental conditions extracted from the analytical feasibility study:

- T28_H70: $T_{\text{ref}} = 28^\circ\text{C}$, $H_{\text{ref}} = 70\%$
- T30_H60: $T_{\text{ref}} = 30^\circ\text{C}$, $H_{\text{ref}} = 60\%$
- T34_H45: $T_{\text{ref}} = 34^\circ\text{C}$, $H_{\text{ref}} = 45\%$
- T35_H35: $T_{\text{ref}} = 35^\circ\text{C}$, $H_{\text{ref}} = 35\%$

These conditions represent realistic harvesting scenarios and serve as anchor points for interpolating OTA risk under intermediate environmental states.

Binary Classification Based on Regulatory Thresholds

For binary classification, samples were labeled according to the maximum allowable OTA concentration defined by European Union regulations. Concentrations less than or equal to 2 g/kg ($C \leq 2$) were classified as acceptable (class 0), while concentrations exceeding this threshold ($C > 2$) were classified as unacceptable (class 1):

$$\text{Label}_{\text{binary}}(C) = \begin{cases} 0, & \text{if } C \leq 2 \\ 1, & \text{if } C > 2 \end{cases} \quad (4.4)$$

This binary labeling scheme is directly aligned with Commission Regulation (EC) No. 1881/2006 [168], ensuring regulatory relevance for food safety assessment.

Multi-Class Risk Stratification

To support finer-grained contamination assessment, a multi-class labeling scheme was also defined, partitioning samples into five OTA risk levels:

$$\text{Label}_{\text{multi-class}}(C) = \begin{cases} 0, & \text{if } 0 \leq C \leq 1 \\ 1, & \text{if } 1 < C \leq 2 \\ 2, & \text{if } 2 < C \leq 3 \\ 3, & \text{if } 3 < C \leq 4 \\ 4, & \text{if } C > 4 \end{cases} \quad (4.5)$$

This stratification enables the model to distinguish progressive contamination levels, supporting both early warning and decision-making under varying risk scenarios.

Implementation and Role in Model Training

The proposed labeling function integrates OTA concentration estimation and risk classification into a single automated procedure, enabling consistent annotation of the environmental dataset used for model training and evaluation. By grounding the labeling process in analytical measurements rather than heuristic assumptions, the resulting dataset provides a reliable foundation for developing and validating on-device AI models capable of predicting OTA contamination risk under real harvesting conditions.

Dataset Structure

Using the labeling function described in Section 4.2.4, two distinct labeled datasets were generated to support different learning objectives within the OTA risk assessment framework: one targeting anomaly detection and one enabling probability-based risk classification.

Anomaly Detection Dataset (ADD)

The Anomaly Detection Dataset (ADD) adopts a binary classification scheme aligned with regulatory compliance requirements. Samples associated with OTA concentrations below the maximum allowable threshold for grape juice and wine production are labeled as safe (class 0), whereas samples exceeding this threshold are labeled as unsafe (class 1). This formulation is suitable for early warning systems and decision support scenarios where a clear safe/unsafe distinction is required. The dataset structure is summarized in Table 4.2.

Table 4.2: Structure of the Anomaly Detection Dataset (ADD).

Concentration of OTA	Concentration (mg/L)	Class	Production Safety
Below threshold	$0.00 \leq \text{Concentration} \leq 2.00$	0	Safe
Above threshold	$\text{Concentration} > 2.00$	1	Unsafe

Probability Classification Dataset (PCD)

To enable a more detailed assessment of contamination risk, a Probability Classification Dataset (PCD) was also defined. This dataset extends the labeling scheme to five classes, allowing the model to distinguish progressive OTA contamination levels. Classes 0 and 1 correspond to conditions compliant with regulatory limits, while classes 2 to 4 indicate increasing levels of non-compliance and risk severity. The complete class structure is reported in Table 4.3.

Table 4.3: Structure of the Probability Classification Dataset (PCD).

Concentration of OTA	Concentration (mg/L)	Class	Production Safety
Approaching zero	$0.00 \leq \text{Concentration} \leq 1.00$	0	Safe
Low	$1.00 < \text{Concentration} \leq 2.00$	1	Safe
Moderate	$2.00 < \text{Concentration} \leq 3.00$	2	Unsafe
High	$3.00 < \text{Concentration} \leq 4.00$	3	Unsafe
Critical	$\text{Concentration} > 4.00$	4	Unsafe

Dataset Generation Process

The dataset generation procedure was designed to emulate realistic grape harvesting conditions while ensuring consistency with the analytically derived OTA evaluation function. Let N denote the total number of samples to be generated. For each sample, a random harvesting start time ($TimeIn$, T_{in}) and end time ($TimeOut$, T_{out}) were selected within the interval between 3:00 AM and 8:00 PM, reflecting typical harvesting windows. The duration of each harvesting interval was computed in minutes as:

$$Time_{Difference} = T_{out} - T_{in}. \quad (4.6)$$

Within each interval, the following features were calculated:

- $Time_{Difference}$: harvesting duration in minutes;
- $T_{Average}$: average temperature over the interval;
- $H_{Average}$: average relative humidity over the interval.

Using these inputs, OTA concentration was estimated through the evaluation function defined in Section 4.2.4.

$$C(Time_{Difference}, T_{Average}, H_{Average}). \quad (4.7)$$

Each estimated concentration value was subsequently mapped to a discrete label according to the selected dataset structure (binary or multi-class). The final sample

vector consisted of $\{Time_{Difference}, T_{Average}, H_{Average}, \alpha, Label\}$, which was appended to the generated dataset. Here, the parameter α represents a tunable factor that can be adapted to plant-specific or contextual conditions. The complete dataset generation workflow is summarized in Algorithm 1.

Algorithm 1 Dataset generation algorithm

Number of samples N

Generated dataset *generatedDataset*

Initialize $\alpha \leftarrow 1$ {Adaptable to plant-specific conditions}

for $i = 1$ to N **do**

 Randomly select $Time_{In}$ and $Time_{Out}$

$Time_{Difference} \leftarrow Time_{Out} - Time_{In}$ (minutes)

 Compute $Temperature_{Avg}$ and $Humidity_{Avg}$ within the interval

$Concentration \leftarrow C(Time_{Difference}, Temperature_{Avg}, Humidity_{Avg}, \alpha)$

 Assign $Label$ based on $Concentration$ and dataset structure

 Append $\{Time_{Difference}, Temperature_{Avg}, Humidity_{Avg}, \alpha, Label\}$ to *generatedDataset*

end for

4.2.5 Model Development

Model Design and Training

Two lightweight neural network models were developed to address the OTA risk assessment tasks defined in this work: one dedicated to anomaly detection (ADD) and one to probability-based multi-class classification (PCD). Both models share the same core architecture, which was progressively refined from the initial proof-of-concept design introduced in [162] to improve accuracy, generalization capability, and computational efficiency. The final architecture, illustrated in Fig 4.7 from [96], was explicitly designed to satisfy the constraints of on-device deployment on resource-limited MCU platforms. Each model consists of an input layer followed by three fully connected hidden layers, each comprising 100 neurons. The input layer processes four features that capture the environmental and temporal conditions influencing OTA development: harvesting duration ($Time_{Difference}$), average temperature ($T_{average}$), average relative humidity ($H_{average}$), and the adaptive parameter α , which accounts for plant- or context-specific variability. Rectified Linear Unit (ReLU) activation functions were employed in the hidden layers to enable efficient learning of nonlinear relationships while avoiding vanishing gradient effects. The output layer configuration depends on the target task. For the ADD model, a single neuron with a sigmoid activation function was used to perform binary classification, producing a probability score representing the likelihood of exceeding the

regulatory OTA threshold. For the PCD model, a softmax activation function was adopted to generate a normalized probability distribution across five contamination risk classes. The ADD and PCD models contain approximately 20,902 and 21,205 trainable parameters, respectively, reflecting a compact design suitable for embedded inference. Both models were trained using the Adam optimizer and the sparse categorical cross-entropy loss function, which offer robust convergence properties for classification tasks. After empirical evaluation, the number of training epochs was fixed at 250, providing an effective balance between convergence speed and overfitting control. As reported in Tables 4.4 and 4.5, the ADD model achieved a peak training accuracy of 99.29% (loss = 0.219 at epoch 215) and a validation accuracy of 99.71% (validation loss = 0.012). The PCD model reached a training accuracy of 97.28% (loss = 0.707 at epoch 229) and a validation accuracy of 98.55% (validation loss = 0.060). The validation set represented 20% of the available data, which explains the slightly higher validation performance due to its reduced complexity. To further assess generalization capability, a five-fold cross-validation procedure was conducted, with results summarized in Table 4.6. Overall, this iterative refinement process resulted in two compact yet highly accurate models capable of supporting both binary anomaly detection and probability-based classification, while maintaining a computational footprint compatible with on-device execution.

Computational Complexity Analysis

To evaluate the suitability of the proposed models for real-time inference on embedded platforms, the computational complexity was estimated in terms of Floating Point Operations (FLOPs). For a fully connected layer, the number of FLOPs is given by:

$$FLOPs_{\text{layer}} = 2 \times (\text{Number of inputs} \times \text{Number of outputs})$$

ADD Model Calculation

- Input layer: 4 features
- Hidden layers: 3 layers with 100 neurons each
- Output layer: 1 neuron (binary classification)

$$\text{Layer 1: } 2 \times 4 \times 100 = 800$$

$$\text{Layer 2: } 2 \times 100 \times 100 = 20,000$$

$$\text{Layer 3: } 2 \times 100 \times 100 = 20,000$$

$$\text{Layer 4: } 2 \times 100 \times 1 = 200$$

Total FLOPs (ADD): 41,000

PCD Model Calculation

- Input layer: 4 features
- Hidden layers: 3 layers with 100 neurons each
- Output layer: 5 neurons (multi-class classification)

$$\text{Layer 1: } 2 \times 4 \times 100 = 800$$

$$\text{Layer 2: } 2 \times 100 \times 100 = 20,000$$

$$\text{Layer 3: } 2 \times 100 \times 100 = 20,000$$

$$\text{Layer 4: } 2 \times 100 \times 5 = 1,000$$

Total FLOPs (PCD): 41,800

The relatively low FLOP counts confirm that both models are computationally lightweight and suitable for real-time inference on microcontroller-based platforms such as the M5Stick C Plus, enabling fast execution with limited energy consumption.

Table 4.4: Training results for the ADD and PCD models.


Model	Epoch	Accuracy (%)	Loss	Parameters
ADD	215	99.29	0.219	20,902
PCD	229	97.28	0.707	21,205

Table 4.5: Validation results for the ADD and PCD models.

Model	Epoch	Validation Accuracy (%)	Validation Loss	Validation Split (%)
ADD	215	99.71	0.012	20%
PCD	229	98.55	0.060	20%

Table 4.6: Five-fold cross-validation results for the ADD and PCD models.

Model	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5	Mean Accuracy (%)
ADD	99.64	99.71	99.66	99.73	99.68	99.68%
PCD	98.40	98.55	98.50	98.47	98.53	98.49%

**Fig. 4.7:** Architecture of the probability-based classification (PCD) model.

Model Optimization and Embedded Deployment

Beyond model accuracy, particular attention was devoted to optimization strategies aimed at facilitating embedded deployment. During comparative analysis with the initial architecture proposed in [162], it was observed that several input features contributed marginally to predictive performance. A feature selection process was therefore applied, resulting in the retention of the four most informative variables: $Time_{Difference}$, $T_{average}$, $H_{average}$, and α . From the early design stages, the neural network architecture was developed with deployment on the M5C+ device as a primary constraint. This hardware-aware approach enabled careful tuning of layer sizes, parameter counts, and activation functions to fit within the device's limited memory and computational budget. Such alignment between model design and hardware characteristics is essential for effective edge AI solutions, as emphasized in [13]. Following training, the optimized models were successfully converted into a C-based implementation compatible with the embedded execution environment, without degradation in classification accuracy. These optimizations ensured that the final on-device implementation maintains a favorable balance between compactness, computational efficiency, and predictive reliability. Consequently, the proposed solution demonstrates the feasibility of deploying AI-driven OTA risk assessment models directly on low-power embedded systems for real-time agricultural monitoring.

Model Deployment

The deployment of the proposed anomaly detection and classification models on the M5C+ platform required a series of steps to ensure seamless integration within the embedded environment and reliable real-time operation. First, the trained neural

network models were converted into a C-compatible format suitable for execution on the MCU. This conversion was performed using TensorFlow Lite for MCU, a framework specifically designed for deploying machine learning models on devices with constrained memory and computational resources. The resulting TensorFlow Lite libraries were then integrated into the device firmware to enable on-device inference. Once converted, the trained models were stored in the flash memory of the M5C+, and the firmware was adapted to manage the complete inference pipeline. Specifically, the embedded software handles: (i) system activation triggered by the onboard IMU, (ii) periodic acquisition and storage of temperature and humidity data, (iii) real-time inference using the embedded neural network model, and (iv) transmission of inference results and raw sensor data to a remote server via the MQTT protocol. In addition, inference outcomes and key measurements are displayed on the built-in screen of the M5C+, providing immediate feedback to the user during harvesting operations. Resource utilization was evaluated after deployment to assess the feasibility of continuous operation under hardware constraints. As summarized in Table 4.7, the deployed firmware occupied approximately 33.8% of the available RAM and 39.0% of the flash memory, indicating a balanced use of system resources. Inference performance was experimentally measured using an oscilloscope, revealing an execution time of approximately 11–12 ms per inference and an average energy consumption of 2.75 mJ, as reported in Table 4.8. These results confirm that the proposed models satisfy real-time and low-power requirements, making them well suited for edge AI applications in agricultural environments.

Table 4.7: Memory occupation on the M5Stick C Plus device.

Memory Type	Used (Bytes)	Total (Bytes)	Usage (%)
RAM	110,740	327,680	33.80
Flash	511,153	1,310,720	39.00

Table 4.8: Inference time and energy consumption of the M5Stick C Plus implementation.

Metric	Measured Value	Unit
Inference Time	11–12	ms
Energy Consumption	2.75	mJ

Overall, the deployment results demonstrate that the optimized models achieve a favorable balance between computational efficiency, memory footprint, and predic-

tive accuracy, validating their suitability for embedded, resource-constrained agricultural monitoring systems.

4.2.6 Experimentation and Results

This section presents the experimental protocol and the results obtained for the proposed OTA risk assessment framework. The objective is to evaluate the effectiveness of the methodology under realistic conditions and to assess its ability to support contamination detection and classification in grape production. Experimental findings are analyzed in relation to regulatory compliance and compared with related approaches discussed in the state of the art.

Compliance Verification Using Grape Harvest Data from Reggio Calabria

To validate the practical applicability of the OTA concentration prediction function described in Section 4.2.3, the model was tested using real-world environmental data collected during grape harvesting activities in the village of Bianco, located in Reggio Calabria [169]. Temperature and humidity measurements were recorded on the same days when local farmers harvested white grapes [162], providing a realistic representation of field conditions. The collected data were used as inputs to the OTA prediction function to estimate contamination levels under actual harvesting scenarios. Fig 4.8 from [96] shows the resulting distribution of predicted OTA concentrations. The analysis indicates that all estimated values remained below the regulatory limit of 2 mg/L, corresponding to the maximum allowable concentration for grape juice and wine production. This experiment serves as a real-world validation of the proposed evaluation function, confirming its reliability in capturing the relationship between environmental conditions and OTA contamination risk. In addition, two synthetic test datasets were generated to systematically evaluate the performance of the ADD and PCD classification models.




Fig. 4.8: Distribution of OTA concentration predictions using real-world harvesting data from Reggio Calabria.

Test Set Generation

For the ADD model, a binary classification test set comprising 2000 samples per class was generated, resulting in a total of 4000 samples. For the PCD model, which performs five-class classification, 2000 samples were generated for each class, yielding a test set of 10,000 samples. These datasets were designed to uniformly cover the range of environmental conditions defined by the OTA evaluation function.

Test Set Distribution

Prior to performance evaluation, both test sets were analyzed to verify their distribution with respect to the OTA prediction function illustrated in Fig 4.5. Fig 4.11 from [96] depicts the distribution of ADD test samples along the predicted OTA concentration curve, while Fig 4.11 shows the corresponding distribution for the PCD dataset. The results confirm that the generated samples adequately span the function's domain, ensuring a representative evaluation of both models.




Fig. 4.9: ADD samples distribution




Fig. 4.10: PCD samples distribution

Fig. 4.11: Distribution of test set samples for the ADD and PCD models.

Model Testing

Model performance was evaluated using the generated test datasets. The ADD model achieved a test accuracy of 99.07%, correctly classifying 3963 out of 4000 samples, as illustrated in Fig 4.14 from [96], the PCD model achieved a test accuracy of 96.53%, correctly classifying 9653 out of 10,000 samples.




Fig. 4.12: ADD confusion matrix




Fig. 4.13: PCD confusion matrix

Fig. 4.14: Confusion matrices for the ADD and PCD models.

4.2.7 Results and Discussion

The proposed framework enables real-time OTA contamination assessment using two complementary learning strategies: binary anomaly detection (ADD) and multi-class probability-based classification (PCD). By exploiting temporal information together with environmental variables such as temperature and humidity, the system

delivers fast and reliable contamination risk predictions that are suitable for on-site decision-making during grape harvesting operations. The effectiveness of the proposed approach is confirmed by the high classification accuracies achieved by both models. The ADD model reached a test accuracy of 99.07%, while the PCD model achieved 96.53%, demonstrating strong generalization capability across both synthetic and field-related datasets. These results confirm that lightweight neural networks, when properly designed and trained, can effectively capture the nonlinear relationships governing OTA risk while remaining suitable for embedded deployment. Inference results are displayed directly on the M5C+ screen, providing immediate feedback to operators in the field and enabling preventive actions to avoid potentially contaminated grapes entering the production chain, as illustrated in Fig 4.15 from [96]. This real-time visualization enhances system usability and supports rapid, human-in-the-loop decision-making without reliance on external computing infrastructure.



Fig. 4.15: OTA classification results displayed on the M5C+ device, including the predicted class and inference time.

From a deployment perspective, the embedded implementation validated the feasibility of executing AI-based OTA risk assessment under strict hardware and energy constraints. The deployed models occupied 33.8% of the available RAM and 39.0% of the flash memory, while maintaining inference times between 11 ms and 12 ms and an average energy consumption of 2.75 mJ per inference. These values confirm that the proposed solution meets the real-time requirements of edge AI applications in agriculture. To further evaluate energy efficiency, two battery consumption experiments were conducted with the LCD display disabled and inference results transmitted via MQTT after each prediction. In the first experiment, inferences were performed every 30 minutes, corresponding to typical harvesting durations. Over a 6-hour period, battery charge decreased from 96.87% to 55%. In the second experiment, inferences were executed every 60 minutes, resulting in a battery decrease from 97.23% to 63.32% over the same duration. These results demon-

strate that continuous sensing, inference, and wireless communication impose a non-negligible energy cost, highlighting the importance of adaptive inference scheduling and power management strategies for long-term outdoor deployments. Figs 4.16 and 4.17 provide a detailed visualization of classification outcomes for both models. For the ADD model (Fig 4.16 from [96]), misclassifications were rare and primarily occurred near the regulatory threshold, where small variations in estimated OTA concentration can lead to class boundary transitions. For the PCD model (Fig 4.17 from [96]), most errors were observed near adjacent class boundaries, reflecting inherent ambiguity in samples with OTA concentrations close to threshold values. Despite these limitations, the overall distribution of correct classifications confirms the robustness of both models across the full range of evaluated conditions.




Fig. 4.16: Visual representation of ADD model classification results on the M5C+.




Fig. 4.17: Visual representation of PCD model classification results on the M5C+.

Beyond classification performance, the proposed system represents a significant shift from traditional OTA detection methodologies. Conventional laboratory-based techniques such as HPLC and LC–MS remain essential for quantitative confirmation

but are unsuitable for continuous, in-field monitoring due to cost, time, and infrastructural constraints. In contrast, the presented on-device AI framework provides a low-cost, portable, and autonomous solution that enables preventive risk monitoring directly at the production site. The validation of the OTA prediction function using real harvesting data collected in Reggio Calabria further strengthens the scientific reliability of the approach. Predicted OTA concentrations remained within regulatory limits, confirming the consistency between chemical analysis, mathematical modeling, and AI-based inference. This alignment supports the transition from simulation-driven experimentation to real-world deployment and demonstrates the practical viability of embedded intelligence for OTA risk management in viticulture.

4.2.8 Conclusion

This section presented the design, deployment, and experimental validation of a lightweight on-device AI framework for real-time OTA risk assessment in grapes. Building on prior exploratory work relying on synthetic data, the present study strengthened the approach through chemical analysis of grape samples and regulatory-aligned labeling, enabling the development of two compact models: ADD for binary compliance-oriented detection and PCD for multi-class risk stratification. The embedded implementation on the M5C+ demonstrated that accurate OTA risk prediction can be achieved with limited computational and energy resources, executing inference within 11–12 ms and consuming approximately 2.75 mJ per prediction. In addition to real-time on-device feedback, MQTT-based transmission enabled remote monitoring and integration into broader traceability workflows.

In summary, the main outcomes of the proposed OTA solution are:

- Design and validation of two lightweight neural network models (ADD and PCD) for OTA detection and risk classification.
- Definition of a chemically grounded evaluation and labeling function aligned with EU regulatory thresholds.
- Successful on-device deployment on a resource-constrained MCU platform (M5C+) with low latency and moderate memory usage.
- Demonstration of practical feasibility using field-related harvesting data and real-time communication via MQTT.

These results demonstrate that embedded AI can provide scalable and autonomous tools for preventive food safety monitoring in viticulture. The next section of this chapter extends the on-device AI perspective to agri-food process monitoring through a second solution focused on fermentation stage classification under embedded hardware constraints.

4.3 On-Device AI for Fermentation Stage Classification

4.3.1 Problem Statement and Motivation

Fermentation is a fundamental bioprocess in the agri-food sector, underpinning the production of beverages, dairy products, and plant-based foods. Despite its long-standing industrial adoption, effective fermentation monitoring remains challenging due to the intrinsic variability of biological systems, sensitivity to environmental conditions, and dynamic microbial activity. Conventional monitoring approaches are largely based on manual sampling and laboratory analyses, which are time-consuming, labor-intensive, and unsuitable for continuous, real-time quality assessment. These limitations hinder timely intervention and increase the risk of batch inconsistency and quality degradation. Recent advances in AI and IoT technologies have enabled data-driven fermentation monitoring through sensor networks, computer vision, and machine learning models [170, 171, 172]. Such approaches have demonstrated the potential to classify fermentation stages and predict process evolution in real time. However, many existing solutions rely on cloud-based computation or high-end hardware platforms, introducing latency, dependency on network connectivity, and limited applicability in small-scale or resource-constrained production environments. Edge AI offers a promising alternative by enabling inference directly on embedded devices, thereby reducing communication overhead, energy consumption, and response time while improving system autonomy. Nevertheless, deploying AI models on edge hardware introduces additional challenges, particularly in terms of robustness under real operating conditions. Models trained on static or curated datasets may exhibit degraded performance when exposed to live variations in illumination, viewpoint, and texture, as observed in real-time fermentation monitoring scenarios. This highlights the need for lightweight yet resilient on-device architectures that can maintain stable performance under uncontrolled conditions. In this context, fermentation monitoring represents a complementary application to environmental risk assessment within Chapter 4. While the OTA case study focuses on contamination risk prediction based on environmental parameters, fermentation monitoring addresses process-stage classification using visual cues. Both problems share common constraints real-time operation, limited hardware resources, and the need for autonomous decision support making fermentation an ideal second case study for evaluating the capabilities and limitations of on-device AI solutions in agri-food systems.

4.3.2 State of the Art in Fermentation Monitoring

The integration of IoT and AI technologies has transformed the monitoring, control, and optimization of fermentation processes. These digital solutions address long-standing challenges in traditional fermentation, such as limited real-time visibility, process variability, and dependence on manual supervision. Through continuous data acquisition and intelligent analytics, IoT- and AI-based systems enable predictive monitoring, process automation, and improved product quality in both artisanal and industrial fermentation environments. Early research has focused on using sensor networks to track essential fermentation parameters. Xu et al. [170] combined IoT devices, backpropagation (BP) neural networks, and expert systems to enhance lactic acid fermentation control. Real-time temperature, pH, and oxygen data were processed in the cloud to predict microbial growth and dynamically adjust conditions, improving efficiency and consistency. Similarly, Kimutai et al. [171] developed an IoT-based system for black tea fermentation using a Raspberry Pi camera and a CNN model named TeaNet. By applying a majority voting mechanism, the system achieved accurate classification of under-fermented, fermented, and over-fermented stages, improving the precision of quality assessment. In another study, An et al. [173] introduced hyperspectral imaging and data fusion to evaluate black tea taste quality. Using algorithms such as SPA and ACO for feature selection and regression models like PLSR, SVR, and ELM, the approach correlated spectral features with sensory evaluations, offering a non-destructive, real-time prediction method. Zhou et al. [172] proposed a multimodal sensor fusion strategy integrating computer vision and electronic nose (e-nose) systems to classify black tea fermentation stages, achieving 95.6% accuracy using SVM models. Bhargava et al. [174] extended this approach by applying computer vision and statistical feature extraction to tea fermentation, demonstrating that SVM classifiers reached 98.75% accuracy when deployed in real processing environments. Beyond tea, several works have explored IoT-based monitoring for other fermented products. Vovsahlik et al. [175] designed an Industry 4.0-compatible system for rice wine fermentation using Raspberry Pi controllers and multi-sensor modules (temperature, pH, humidity). Automated stirring improved homogeneity, stability, and sugar conversion rates. A subsequent work [176] introduced a 3D-printed stirrer prototype that enhanced mixing and reduced local pH fluctuations, demonstrating the potential of affordable hardware in process optimization. Buonocore et al. [177] applied a similar IoT framework to beer production, using ESP8266 MCUs and load-cell sensors to monitor CO₂ release through weight variation, effectively identifying fermentation phases and detecting anomalies. In cocoa fermentation, Abijaude et al. [178] developed the IoT Cocoa platform that integrates IoT sensors, BC, and smart contracts to secure

and trace data. The system monitored temperature, pH, and humidity while storing data immutably on BC ledgers, ensuring transparency and reproducibility. Wu et al. [179] applied IoT-based control to probiotic fermentation of oil tea cake, optimizing microbial growth through real-time temperature and pH regulation, and achieving a tea saponin degradation rate of 93.96%. Similarly, Saguidon et al. [180] proposed a low-cost, automated IoT system for cacao bean fermentation, maintaining stable temperature conditions using nichrome wire heating and Arduino-based feedback control to improve bean quality and flavor consistency. Recent works have also explored AI-driven and FL approaches to fermentation monitoring. Bowler et al. [181] employed ultrasonic sensors and LSTM-based models to monitor beer fermentation, integrating domain adaptation and federated learning to enhance prediction accuracy across different breweries without sharing private data. Other studies, such as [182], used low-cost IoT nodes equipped with CO₂ sensors and Wi-Fi modules to monitor wine fermentation, providing real-time feedback through ThingSpeak cloud dashboards. Faraco et al. [183] combined CO₂ and temperature sensing with clustering algorithms (DBSCAN, PCA, k-means) to characterize coffee fermentation aromas, replacing traditional sensory analysis with automated digital methods. Innovations have also targeted traditional fermentation industries. Triyason et al. [184] designed the PLAKDA IoT platform to monitor pH, salinity, and temperature in plara (fermented fish), reducing contamination risks and improving batch uniformity. Klaithin et al. [185] developed an Arduino- and ESP32-based bioreactor for biobutanol fermentation, providing cloud-integrated monitoring of temperature, pH, and dissolved oxygen with automated sample collection. Overall, these studies demonstrate the versatility and scalability of IoT- and AI-based monitoring systems across a wide range of fermented products. From small-scale artisanal production to industrial applications, such systems provide real-time, data-driven insights that improve product uniformity, reduce waste, and enhance traceability. The reviewed works are summarized in Table 4.9, which categorizes each study by technology, product type, and level of system maturity.

Table 4.9: Technologies and Products in Reviewed Studies

Cited Work	Technologies Used	Product	State of Solution
[170]	IoT sensors, BP neural networks, expert systems	Lactic acid fermentation	Experimental

Continued on next page

Cited Work	Technologies Used	Product	State of Solution
[171]	IoT, CNN, Raspberry Pi, majority voting	Black tea	Prototype
[173]	Hyperspectral imaging, SPA, ACO, PLSR, SVR, ELM	Black tea	Experimental
[172]	Computer vision, e-nose, GC-MS, SVM, SVR	Black tea	Experimental
[174]	Computer vision, PCA, SVM, k-NN, SRC	Black tea	Experimental
[175]	IoT sensors, Raspberry Pi, automated stirring	Rice wine	Experimental
[177]	IoT sensors, ESP8266, load cells	Beer	Prototype
[176]	IoT sensors, Raspberry Pi, 3D-printed stirrers	Rice wine	Experimental
[178]	IoT sensors, BC, smart contracts	Cocoa	Prototype
[179]	IoT sensors, pH, temperature control	Oil tea cake	Experimental
[181]	Ultrasonic sensors, LSTM, federated learning	Beer	Experimental
[182]	CO ₂ sensors, Arduino, ESP8266, ThingSpeak	Wine	Experimental
[180]	Temperature sensors, Arduino, heating control	Cacao beans	Prototype
[183]	CO ₂ and temperature sensors, PCA, DBSCAN, k-means	Coffee	Prototype
[184]	IoT sensors, wireless communication	Fermented fish (Plara)	Prototype
[185]	IoT sensors, Arduino, ESP32 MCU, cloud monitoring	Biobutanol	Prototype

4.3.3 Case Study Overview and Pipeline

Datasets

This case study focuses on developing an Edge-AI pipeline for classifying fermentation stages using visual features extracted from two datasets: a self-generated kefir fermentation dataset and a publicly available tea fermentation dataset. The kefir

dataset was generated to address the scarcity of open-access fermentation image datasets, while the tea dataset served as a benchmark for cross-domain validation of the proposed classification framework.

Kefir Dataset Generation

Following the literature review, a synthetic dataset generation phase was conducted to produce labeled images of kefir fermentation stages: under-fermented, fermented, and over-fermented. The process was designed to create visually distinct and representative samples suitable for model training and testing. Data generation was carried out using the *DataDreamer* tool, which enables prompt-based image synthesis to create realistic, domain-specific datasets for machine learning. Multiple iterations of prompt refinement were performed to achieve accurate visual representation of each stage. The final prompts are shown in Listing 4.1.

```
prompts = {"fermented-kefir": (
    "Close-up view of properly fermented kefir in a clear glass jar, "
    "highlighting the creamy, slightly thickened texture with fine bubbles "
    "and gentle curdling on the surface. Focus on the upper layer."
),
"over-fermented-kefir": (
    "Macro view of over-fermented kefir in a clear glass jar, showing distinct "
    "separation between watery whey and dense, grainy curds. Highlight the "
    "contrast between layers and the cracked surface."
),
"under-fermented-kefir": (
    "Detailed shot of under-fermented kefir, emphasizing the thin, smooth, "
    "milky texture with minimal curd formation, indicating early fermentation."
)
```

Listing 4.1: Final DataDreamer prompts for kefir dataset generation

In total, 1,200 images were generated per class, yielding a balanced dataset of 3,600 synthetic images for kefir fermentation classification. Representative samples for each class are displayed in Fig 4.18.



Fig. 4.18: Examples of generated kefir images using DataDreamer showing distinct class characteristics.

Data Preprocessing and Curation

To ensure data quality and model robustness, the generated kefir images underwent a structured preprocessing pipeline. In contrast, the tea fermentation dataset was already preprocessed and therefore used directly for model evaluation.

Step 1: Background Removal.

All synthetic images were processed using a U-Net-based segmentation model to remove non-relevant backgrounds, isolating the kefir content to focus on fermentation-specific visual features. Some samples were discarded when segmentation errors completely removed the content, reducing the dataset to approximately 900 images per class.

Step 2: Cropping and Manual Curation.

After segmentation, images were automatically cropped to focus on regions displaying clear fermentation indicators. Manual inspection was performed to exclude ambiguous samples and preserve only high-quality images representing each stage distinctly. The final dataset size per class is summarized in Table 4.10.

Table 4.10: Final dataset distribution after preprocessing and manual selection

Fermentation Stage	Number of Samples
Fermented	350
Over-fermented	295
Under-fermented	124

An example of the cropping process and the resulting high-quality samples is shown in Fig 4.19. This curated dataset forms the foundation for subsequent training and deployment of the Edge-AI classification model.




Fig. 4.19: Illustration of the cropping and cleaning process for preparing high-quality kefir fermentation images.

4.3.4 AI Model Design

Custom CNN Model

A custom CNN was implemented in TensorFlow/Keras as a lightweight baseline tailored for embedded inference. Images were prepared as described in the preprocessing pipeline and resized to either 128×128 or 224×224 to study the trade-off between computational cost and performance. Pixel values were normalized to $[0, 1]$, and labels were one-hot encoded for three classes. The architecture comprises four convolutional blocks with filter sizes $\{32, 64, 128, 128\}$. Each block includes Conv2D, Batch Normalization, and Max Pooling; Dropout layers are inserted to reduce overfitting. A Global Average Pooling layer precedes a dense layer with 128 ReLU units, followed by a softmax output layer for three-way classification (under-fermented, fermented, over-fermented). The model was trained with Adam (learning rate 1×10^{-4}) and categorical cross-entropy. Training ran up to 50 epochs with early stopping (patience 30) and model checkpointing on validation accuracy. This network serves as a compact baseline that is straightforward to quantize and deploy in edge pipelines.

Transfer Learning with EfficientNetB0 and MobileNet

To leverage strong image priors while retaining edge feasibility, transfer learning experiments were conducted with EfficientNetB0 (ImageNet pre-trained) and Mo-

bileNet variants. The backbone was unfrozen from the start to enable end-to-end fine-tuning on fermentation visual cues. On top of the backbone, the head includes Global Average Pooling, Dropout, a dense layer with 128 ReLU units, and a softmax output. Models were trained with Adam at a reduced learning rate 1×10^{-5} to stabilize fine-tuning of pre-trained weights. Sparse categorical cross-entropy was used with integer labels, early stopping (patience 5) and best-weight restoration, for up to 30 epochs. This approach typically improves generalization on limited, domain-specific data while remaining amenable to edge deployment when coupled with input downsizing and post-training quantization.

YOLOv10 for Whole-Image Classification

YOLOv10 [186] was adapted for image-level classification by assigning a single full-frame bounding box whose label corresponds to the fermentation stage. Training used the Ultralytics framework with standard augmentations and a three-class head. The model was trained for 100 epochs on the kefir dataset and 10 epochs on the tea dataset. Since localization was not required, evaluation focused on classification accuracy and throughput. This configuration preserves YOLOv10's real-time inference characteristics while simplifying the task to whole-image decisions. In practice, it also offers a migration path to future setups where region-level localization or multi-object scenes may be relevant without changing the family of tools.

Summary for Edge Deployment

Table 4.11: Qualitative comparison of model choices for edge deployment

Model	Strengths	Considerations
Custom CNN	Small footprint, simple to quantize, consistent latency at 128×128 inputs.	May underfit complex textures, performance can be sensitive to input size and augmentation.
EfficientNetB0 / MobileNet (TL)	Strong priors, good accuracy on limited data, flexible input scaling; widely supported in edge runtimes.	Requires careful LR and regularization; full fine-tuning increases compute, consider 8-bit quantization for on-device use.
YOLOv10 (image-level)	High throughput, unified path to future detection/localization, robust data augmentations built-in.	Heavier than minimal CNNs, classification head adapted from detection ensure calibration for three-class outputs.

For embedded deployment, all models can be exported to ONNX or TFLite with post-training quantization (e.g., dynamic range or full-integer with representative images). Input resolution should be chosen to meet device memory and latency targets, with 128×128 or 160×160 often providing an effective balance for camera-based edge setups.

subsection

4.3.5 Training and Offline Evaluation

Classification Models

To evaluate the impact of preprocessing and model selection on classification performance, several training configurations were tested using the generated and pre-processed kefir fermentation dataset. The experiments aimed to quantify how each improvement in data preparation and model architecture contributed to overall accuracy and convergence stability. The baseline experiment involved training the custom CNN directly on the raw, unprocessed dataset. This provided a reference for subsequent comparisons. When the same architecture was trained on images with background removal using the U-Net segmentation model, a modest improvement in validation accuracy was observed. The removal of irrelevant background elements helped the network focus more effectively on the kefir texture and surface features characteristic of each fermentation stage. Next, TL using the EfficientNetB0 architecture was tested. Leveraging pre-trained ImageNet weights improved feature generalization and reduced overfitting, leading to an additional increase in validation accuracy. However, the most substantial performance gain occurred when training the custom CNN on the manually cropped dataset, which emphasized key visual cues of each fermentation stage. This curated dataset provided clearer and more consistent examples, allowing the model to achieve stable and accurate classification. A comparison of validation accuracies for the different configurations is summarized in Table 4.12.

Table 4.12: Validation accuracy across different training configurations

Training configuration	Validation accuracy
Custom CNN on raw dataset	64%
Custom CNN on dataset with background removal (U-Net)	67%
Transfer Learning (EfficientNetB0)	70%
Custom CNN on cropped dataset	90%

The learning behavior of the best-performing configuration, the custom CNN trained on cropped images is illustrated in Fig 4.20. The training and validation accuracy and loss curves indicate smooth convergence with minimal overfitting, confirming that the preprocessing pipeline successfully enhanced model learning stability and feature discrimination.




Fig. 4.20: Training and validation accuracy and loss curves for the custom CNN trained on cropped images.

Detection and Classification Model (YOLOv10)

The YOLOv10 model was trained as a hybrid detection–classification framework for the same task. Each image was treated as a single bounding box covering the full frame, with the assigned label representing the corresponding fermentation stage. Training was performed using the Ultralytics implementation for 100 epochs on the kefir dataset and 10 epochs on the tea fermentation dataset. The training process simultaneously optimized the detection (box regression) and classification loss functions as described in [187]. Rapid convergence was observed across both datasets, confirming that the model effectively adapted to the simplified classification scenario. The results are summarized in Table 4.13.

Table 4.13: YOLOv10 training summary for kefir and tea fermentation datasets

Dataset	Epochs	Box loss (final)	Cls loss (final)
Kefir	100	0.034	0.056
Tea	10	0.040	0.063

These results indicate that YOLOv10 maintained low box and classification losses, confirming efficient training and robustness across datasets. Its high through-

put and low latency make it a strong candidate for real-time deployment on Edge AI devices in industrial fermentation monitoring systems.

4.3.6 Proposed Setup and Model Deployment

The proposed setup for real-time kefir fermentation stage classification is illustrated in Fig 4.21, showing the complete pipeline from model training to on-device deployment and inference on the Luxonis AI Camera PoE.




Fig. 4.21: Proposed system architecture for Edge AI kefir fermentation stage monitoring, including model training, deployment, and real-time on-device inference.

Camera Setup

The Edge AI camera is configured to stream the kefir fermentation process in real time while performing on-device inference using the deployed model. This continuous monitoring approach enables immediate classification of the fermentation stage without manual supervision, facilitating integration into automated traceability and quality control pipelines within agri-food systems.

Materials: Camera Used

Fig. 4.22: Luxonis AI Camera PoE integrated for real-time Edge AI fermentation monitoring.

The camera used in this setup, shown in Fig 4.22, is the Luxonis AI Camera PoE, an embedded AI-enabled device capable of executing NN inference locally without reliance on external computation. Its main features include:

- Real-time video streaming via Power-over-Ethernet (PoE) for stable operation.
- Support for on-device inference through the .blob format using the DepthAI framework.
- Compatibility with various connectivity interfaces for integration into IoT pipelines.
- Adequate resolution and frame rate for visually tracking fermentation progression.
- Built-in DepthAI platform optimized for real-time computer vision and AI inference.
- Proven deployment history in tasks such as on-device detection, age and gender estimation, and environmental monitoring.

These capabilities make the Luxonis AI Camera PoE suitable for continuous, automated observation of kefir fermentation, reducing the need for centralized computation or manual intervention.

4.3.7 Model Conversion

Before deployment, trained models must be converted into formats compatible with the Luxonis platform. Models are typically exported first to TFLite or ONNX to optimize for size and performance, then converted to the .blob format required for DepthAI execution. In this study, the trained Keras model was converted to TFLite

using Python, followed by conversion to the .blob format via the official Luxonis online converter [188]. The overall workflow is shown in Fig 4.23. This conversion pipeline enabled seamless deployment and execution of the model directly on the camera, allowing continuous real-time inference while streaming live video of the kefir fermentation process.




Fig. 4.23: Model conversion workflow: trained Keras model exported to TFLite and then to .blob format for deployment on the Luxonis AI Camera PoE.

4.3.8 Model Real-Time Inference and Discussion

Classification Models

After deployment, real-time inference was conducted using the Luxonis AI Camera PoE. During testing, the model frequently predicted the *fermented* class, with fewer predictions for the *under-fermented* and *over-fermented* stages. Confidence values fluctuated dynamically in response to variations in illumination, viewing angle, and kefir texture during live monitoring. During live streaming, the embedded model exhibited frequent misclassifications and reduced classification accuracy compared to offline evaluation. To simulate realistic observation conditions, sample images were displayed on a monitor and analyzed by the camera in real time. This setup validated the end-to-end functionality of the inference pipeline under controlled but representative conditions. Performance variability was largely attributed to the limited dataset size (around 120 samples per class) and the visual discrepancies between generated training images and live camera inputs. Additionally, minor accuracy degradation may occur due to quantization and conversion from Keras to

TFLite and subsequently to .blob format. Examples of live inference outputs are shown in Fig 4.24.




Fig. 4.24: Examples of real-time inference outputs during live monitoring showing predicted classes and fluctuating confidence values.

Detection and Classification Models (YOLOv10)

For comparison, the YOLOv10 model was also deployed on the Luxonis AI Camera PoE, treating each frame as a full-image detection problem. The model consistently identified the *under-fermented* and *fermented* kefir stages with confidence levels above 80–95%. Similar robustness was observed for the tea fermentation dataset. The primary limitation occurred in the *over-fermented* class, where occasional misclassifications were observed (Fig 4.25). This was primarily due to lower image quality and class imbalance during dataset generation. Nevertheless, YOLOv10 achieved stable performance and demonstrated strong suitability for unified detection–classification pipelines on edge devices.




Fig. 4.25: Real-time inference examples using YOLOv10. Captions indicate original and predicted classes. Misclassification occurred primarily for over-fermented kefir samples.

4.3.9 Discussion and Conclusion

This study presented the design and implementation of an Edge AI pipeline for real-time monitoring and classification of kefir fermentation stages. The study integrated computer vision, DL, and IoT technologies to address challenges in automated traceability and quality assessment within agri-food systems. The research began with an extensive review of fermentation monitoring and traceability frameworks, identifying key technological gaps in real-time and on-device analysis. To overcome the scarcity of open fermentation datasets, a synthetic dataset was generated using DataDreamer with carefully engineered prompts to capture visual distinctions among under-fermented, fermented, and over-fermented kefir samples. Comprehensive preprocessing, including background removal with U-Net and manual cropping, significantly enhanced data clarity and directly improved model performance during training. Two classification strategies were explored: a lightweight custom CNN and a transfer learning model based on EfficientNetB0. Comparative experiments revealed that model performance was highly dependent on dataset quality, with the cropped image dataset yielding a validation accuracy of up to 90%. However, when deployed on the Luxonis AI Camera PoE for real-time inference,

both models demonstrated inconsistencies caused by environmental variations such as illumination changes, camera angles, and surface texture dynamics. These results highlight the inherent limitation of static image-trained classification models when applied to live, uncontrolled scenarios. In contrast, the YOLOv10 classification framework achieved robust and stable real-time performance. Trained on the kefir dataset and validated using a publicly available tea fermentation dataset, YOLOv10 maintained high confidence levels during live inference, successfully distinguishing most fermentation stages with minimal misclassification. Occasional errors in the over-fermented kefir category were primarily attributed to the limited representativeness of the generated samples for that class. Despite this, YOLOv10 demonstrated greater adaptability and resilience to dynamic lighting and viewpoint changes compared to traditional classifiers. Overall, the findings confirm that while conventional CNN-based models are valuable for controlled, static image classification, they are not optimal for real-time edge deployment. Detection-based models such as YOLOv10 offer a more robust solution by combining efficient inference, scene adaptability, and scalability for industrial applications. These results emphasize the importance of validating models on both synthetic and real datasets to ensure generalizability across conditions. From a broader perspective, this work demonstrates the feasibility of integrating Edge AI into agri-food traceability and quality monitoring systems. The successful real-time implementation highlights YOLOv10's potential as a core component in automated, on-site fermentation tracking and anti-counterfeiting infrastructures, marking a step toward intelligent, self-sufficient food monitoring pipelines.

4.4 Conclusion

This chapter investigated the role of on-Device AI as a practical enabler for real-time decision support in agri-food systems operating under strict constraints in cost, energy, connectivity, and computational resources. Two complementary case studies were presented to demonstrate how embedded intelligence can be tailored to different stages of the agri-food chain: (i) contamination risk assessment in primary production, and (ii) process monitoring in food transformation. In the first case study, an AI-enhanced IoT framework was developed for OTA risk assessment in grape harvesting. By grounding the evaluation and labeling strategy in chemical analysis and aligning risk thresholds with European Union regulations, two lightweight neural models were designed to support both binary compliance-oriented detection (ADD) and multi-class risk stratification (PCD). The deployment on the M5Stick C Plus validated that accurate and responsive inference can

be achieved on a resource-constrained MCU, with low latency, moderate memory usage, and MQTT-based transmission enabling remote supervision and integration into traceability workflows. In the second case study, a vision-based Edge AI pipeline was proposed for fermentation stage classification. The work addressed the scarcity of publicly available fermentation image datasets through synthetic data generation and systematic preprocessing, and evaluated multiple model families to balance accuracy and deployability. The results highlighted the strong influence of dataset quality on classification performance and emphasized that robustness under real operating conditions is a central requirement for industrial monitoring. In this context, the YOLOv10-based formulation demonstrated improved stability for real-time inference compared to conventional static-image classifiers, supporting the adoption of detection-oriented architectures for resilient on-device monitoring. Overall, the chapter demonstrates that common design principles lightweight models, low-cost sensing, hardware aware optimization, and local inference can be transferred across heterogeneous agri-food problems. These findings support the broader adoption of embedded AI to enhance traceability, certification, and risk mitigation by enabling timely, autonomous, and scalable monitoring directly at the edge. Future work will focus on extending validation under diverse field conditions, improving model robustness to domain shifts, and integrating these on-device components into distributed traceability infrastructures spanning multiple products and production contexts.

RFID-Augmented Edge-AI Framework for End-to-End Agri-Food Traceability and Anti-Counterfeiting

5.1 Introduction

Ensuring product traceability, authenticity, and transparency across agri-food supply chains has become increasingly challenging in the context of globalized production and distribution systems [13]. Traceability plays a fundamental role in guaranteeing food safety, quality control, and regulatory compliance, however, its effective implementation across international and multi-actor supply chains remains complex [189]. At the same time, consumers are progressively more aware of the origin, quality, and composition of the food they consume, leading to growing expectations for reliable labeling and certification mechanisms [190]. This increased awareness has also amplified concerns related to food fraud, as counterfeit and adulterated products pose a serious threat to consumer trust and market integrity [191]. High-value agri-food products such as EVOO [192], honey [193], saffron [194], and wine [195] are particularly exposed, as counterfeit versions often closely mimic genuine products while failing to meet required quality or compositional standards. These practices not only damage brand reputation and consumer confidence but also undermine the efforts of legitimate producers who comply with strict traceability and quality requirements [196]. Despite the growing adoption of digital technologies such as RFID, the IoT, BC, and AI in the agri-food sector [197, 198, 199], most existing solutions remain largely focused on post-farm stages of the supply chain, including processing, logistics, and retail. As a result, traceability gaps persist at earlier stages, where product quality and safety are first determined. Even in intermediate phases such as transportation, traceability may be compromised, for instance, products requiring strict cold-chain conditions may be exposed to unsuitable temperatures, leading to quality degradation that is not accurately reflected in recorded data [200, 201]. Consequently, the quality estimated during earlier stages may differ significantly from the actual condition of the product reaching consumers or other stakeholders, potentially harming brand reputation and reducing trust. The

farm stage remains the most critical and, at the same time, the least digitally supported segment of the agri-food supply chain. Environmental conditions, cultivation practices, and temporal variability at this stage have a direct impact on product safety and quality, yet reliable and cost-effective traceability solutions are still limited. This issue is particularly relevant in grape harvesting, where variations in temperature and humidity can promote the development of OTA [202, 162, 96], a toxic secondary metabolite produced by certain fungal species. OTA poses serious health risks and can cause substantial economic losses in grape, wine, and juice production. Its presence is strictly regulated by the EU [203], making early monitoring of environmental factors during harvesting a crucial requirement for food safety and quality assurance. In Chapter 4 this thesis addressed this challenge by proposing and validating an AI-IoT system capable of performing real-time, on-device classification of OTA risk based on environmental data collected during the harvesting process. While the OTA-oriented framework presented in Chapter 4 demonstrated the feasibility of embedded intelligence at the farm level, it did not explicitly address the broader requirements of end-to-end TRAC. In particular, most existing traceability systems rely on centralized data transmission and storage architectures, exposing collected information to potential manipulation, loss, or inconsistency across different supply-chain stages, especially in low-connectivity or resource-constrained environments. Ensuring the integrity and authenticity of farm-level data therefore remains an open challenge. Building directly upon the AI-IoT framework introduced in Chapter 4 this chapter represents the main contribution of the thesis and extends the previous work toward a comprehensive solution for TRAC. Specifically, a low-cost RFID-enabled Edge-AI framework is proposed to securely bind processed information to the physical product at the earliest stage of the supply chain. The system integrates affordable sensors with a MCU capable of executing embedded AI models for on-device inference, thereby reducing reliance on continuous cloud connectivity and minimizing data transmission overhead. The inferred information is written directly into passive RFID tags, which act as portable and tamper-evident data carriers, and subsequently anchored to a simulated BC layer to demonstrate the feasibility of immutable record-keeping and anti-counterfeiting verification within a controlled environment. In addition, a web-based interface enables stakeholders and consumers to access and verify product information via QR codes, enhancing transparency and trust across the supply chain. The feasibility of the proposed framework is demonstrated through a laboratory-based case study using grape farm data previously introduced in [96], representing a challenging traceability scenario characterized by environmental variability and limited technological infrastructure.

5.2 Problem Definition, Design Requirements, and State-of-the-Art Positioning

Over the past two decades, agri-food TRAC systems have undergone a progressive evolution driven by the need to improve transparency, data reliability, and automation across increasingly complex supply chains. Early digital traceability solutions were primarily based on RFID technology, which enabled non-line-of-sight identification and automated product tracking [204]. By replacing paper-based documentation with electronic identifiers and sensor-assisted data collection, RFID systems represented a foundational step toward modern traceability infrastructures. Practical deployments demonstrated their effectiveness in monitoring environmental conditions during transportation and storage, for example, RFID temperature tags were successfully applied along the pineapple supply chain to track thermal variations with accuracy comparable to conventional methods [45], while other RFID-enabled frameworks showed improvements in quality control and logistical efficiency for perishable goods [43]. Despite these advantages, early RFID-based systems exhibited structural limitations. In particular, many solutions relied on costly active tags and lacked embedded intelligence, resulting in static data storage without contextual interpretation or adaptability across distributed supply-chain nodes. As a consequence, while identification and basic monitoring were improved, these systems were unable to provide meaningful, autonomous decision support or scalable analytics at the point of data generation. The introduction of the IoT marked a significant shift toward continuous, sensor-driven traceability. IoT-based architectures enabled real-time data acquisition from the field, increasing responsiveness and situational awareness across the supply chain. Multilayer IoT frameworks were proposed to support interoperability and data sharing among stakeholders [56], while smart farming systems combined environmental sensing, cloud-based data management, and consumer-facing interfaces such as QR codes to enhance transparency [60]. However, these architectures remained largely dependent on stable network connectivity and centralized data storage, which limits their applicability in rural or resource-constrained agricultural environments. Moreover, the separation between physical products and remotely stored digital records introduces vulnerabilities related to data loss, manipulation, or inconsistency across supply-chain stages. The integration of AI with IoT further advanced traceability systems by enabling autonomous data analysis and adaptive decision-making. AI-IoT architectures extend conventional sensing networks by embedding learning models at the edge or cloud level [129], allowing predictive and prescriptive capabilities for agricultural operations. ML models have been used to classify crop conditions and detect anomalies.

lies from environmental data streams [68], while more complex systems combining drones, neural networks, and image processing have been proposed for automated crop health assessment and field intervention [74]. Although these approaches significantly enhance monitoring accuracy and operational efficiency, they typically require substantial computational resources and reliable network access for model updates and data synchronization. This dependency poses challenges for low-cost deployments and limits scalability in remote agricultural settings. To address data integrity and trust issues, recent research has explored the integration of BC with AI and IoT, giving rise to AI-IoT-BC architectures. These hybrid systems combine intelligent data analytics with decentralized, tamper-resistant storage to enable transparent and verifiable traceability. Multi-layered solutions have been proposed in which IoT devices collect data, AI modules filter and select relevant information, and permissioned BC platforms such as Hyperledger Fabric ensure secure and immutable record-keeping [84]. Similarly, integrated AI-IoT-BC frameworks have been developed to enhance data security, optimize agricultural operations, and support real-time decision-making through predictive analytics [75]. Despite their conceptual robustness, such systems often remain cloud-centric, require continuous connectivity, and store information remotely from the physical product. As a result, the collected data do not physically follow the product along the supply chain, creating potential traceability discontinuities, particularly in offline or low-connectivity scenarios, as illustrated in Fig. 5.1 from [205]. From this analysis, a set of unresolved challenges emerges. Current traceability solutions struggle to (i) ensure persistent linkage between digital information and the physical product, (ii) operate reliably in resource-constrained and low-connectivity environments, (iii) reduce data volume while preserving meaningful information, and (iv) support autonomous decision-making directly at the edge. These limitations are particularly critical at the farm level, where product quality and authenticity are first established and where infrastructural constraints are most pronounced. Based on these observations, the design requirements for an advanced agri-food TRAC framework can be summarized as follows: the system must be low-cost and compatible with MCU-based platforms, it must support local data processing and on-device intelligence to minimize reliance on cloud connectivity, it must ensure that essential information physically follows the product across supply-chain stages, and it must provide mechanisms for data integrity, authenticity, and tamper resistance without imposing excessive computational or energy overhead. Building upon the technological advancements discussed above, this work proposes a practical and resilient architecture that integrates AI, IoT, simulated BC, and passive RFID tags to address the identified gaps. Unlike previous paradigms in which digital information remains detached from the physical

product, the proposed approach ensures product-linked traceability by embedding RFID identifiers at each supply-chain node. Environmental sensing is performed locally through IoT devices, while on-device AI enables immediate classification of collected data and autonomous decision-making at the edge. To optimize memory usage and communication overhead, only the resulting label representing a condensed and meaningful summary of multiple sensed parameters is written into the RFID tag, rather than storing raw data. In the current implementation, the user memory of each RFID tag is written once during harvesting, however, the architecture supports multi-write operations through controlled memory management strategies. Given the limited write cycles of passive UHF tags, overwriting is minimized and versioning mechanisms are foreseen to prevent rollback attacks. Finally, the simulated BC layer serves as a proof-of-concept for secure, tamper-evident data anchoring and verifiable provenance across supply-chain nodes.

5.3 Proposed RFID-Augmented Edge-AI Traceability Framework

5.3.1 Overall System Architecture

The proposed framework [205] is designed to strengthen TRAC at the farm level, where supply chains are most vulnerable to information gaps, data manipulation, and loss of provenance. The architecture targets the grape harvesting process as a representative use case, spanning pre-harvest, harvest, and post-harvest phases and extending to consumer-level verification. The overall workflow is illustrated in Fig. 5.2 from [205]. The system integrates five main components: (i) an edge device with embedded intelligence, (ii) RFID-enabled harvest containers, (iii) a backend platform secured through BC-inspired mechanisms, (iv) a web-based traceability interface, and (v) a QR-code verification layer. At the edge, a MCU equipped with environmental sensors performs local data acquisition and on-device inference using an embedded AI model. While the OTA risk assessment task described in Chapter 4 serves as the reference application, the architecture is inherently task-agnostic and can support alternative embedded-AI functions depending on the targeted agri-food scenario. Each harvest container is equipped with a passive RFID tag that acts as a secure and portable memory unit. The tag's Electronic Product Code (EPC) uniquely identifies the product unit and serves as the primary key for traceability. During harvesting, the EPC is scanned at the beginning of the process to initialize a traceability session and again at completion to bind the inferred information to the correct physical unit. In addition to EPC-based identification, the immutable Tag Identifier (TID) may be leveraged as a hardware-level reference to further strengthen resistance against tag cloning [206].




Fig. 5.1: Comparative overview of agri-food traceability architectures and the positioning of the proposed AI-IoT-BC framework with passive RFID tags.




Fig. 5.2: Workflow of the proposed RFID-Edge-AI system during the harvest process.

5.3.2 End-to-End Traceability Workflow

During harvesting, environmental measurements are continuously collected at the edge and processed locally by the embedded AI module. Instead of transmitting raw sensor data, the system generates a compact and semantically meaningful representation of the observed conditions, which is written directly into the user memory of the RFID tag at the end of the harvest session. This approach ensures that essential information is physically embedded within the product container, significantly reducing data volume and limiting dependence on continuous network connectivity. An RFID reader subsequently scans both the EPC and user memory content and uploads the corresponding records to the backend platform. The backend manages structured data storage and anchors cryptographic hashes of the uploaded records to a BC-inspired layer to ensure immutability and detect tampering. A web-based application (AgriTraceX) provides stakeholders with access to harvest records and traceability information, while a QR-code interface enables consumers and regulators to verify product authenticity and provenance from farm to market. In this framework, and Chapter of the thesis data refer to the raw environmental measurements collected during harvesting, whereas information denotes the processed and validated results generated by on-device intelligence and stored in the RFID tag. This distinction is central to the proposed architecture, as it enables efficient, product-linked traceability without exposing large volumes of raw data to transmission or storage vulnerabilities.

5.3.3 Edge Intelligence and On-Device Data Processing

Edge intelligence constitutes a foundational element of the proposed framework. By executing inference directly on the MCU, the system enables autonomous decision-making at the point of data generation, reducing latency, preserving privacy, and minimizing reliance on cloud-based processing. The embedded AI module operates on constrained hardware resources while producing actionable outputs suitable for traceability and certification purposes. The OTA risk assessment model introduced in Chapter 4 is used in this work as a representative embedded-AI task, however, the framework does not depend on any specific model architecture or application domain. Instead, it provides a general mechanism for transforming raw sensor streams into validated information that can be securely bound to physical products. This abstraction allows the same architecture to be reused for other agri-food monitoring tasks, such as quality grading, contamination risk estimation, or compliance verification.

5.3.4 RFID-Based Product–Data Binding Strategy

A key novelty of the proposed framework lies in its RFID-based product–data binding strategy. Passive UHF RFID tags are used not only for identification but also as portable memory units that store AI-generated information directly on the product container. This design ensures that traceability information physically follows the product across supply-chain stages, overcoming the limitations of architectures in which data remain detached in remote databases. In the current implementation, the user memory of each RFID tag is written once at the completion of harvesting. Nevertheless, the architecture supports multi-write strategies, either by selectively overwriting memory fields or by appending new information starting from the last written address. Given the limited write cycles of passive RFID tags, overwriting operations are minimized, and versioning mechanisms are foreseen to prevent rollback or unauthorized modification of stored information.

5.3.5 Blockchain-Inspired Data Integrity and Provenance Layer

To enhance data integrity and support anti-counterfeiting, a BC-inspired layer is integrated into the backend platform. Each uploaded record is hashed and logically linked to previous records, forming an immutable chain that enables tamper detection and verifiable provenance. While the current implementation operates in a centralized environment without distributed consensus, it effectively demonstrates core BC properties such as immutability, transparency, and auditability. By anchoring cryptographic hashes of RFID-linked records to this layer, the framework provides an additional level of trust without imposing the computational and connectivity overhead associated with fully decentralized BC deployments. This approach is particularly suitable for agri-food environments characterized by limited infrastructure and resource constraints, while remaining compatible with future extensions toward distributed or permissioned BC systems.

5.4 Hardware and Software Implementation

5.4.1 Hardware

The proposed traceability framework relies on a set of lightweight and portable hardware components designed for deployment in field environments with limited infrastructure. The selected components support on-device data acquisition, embedded intelligence, and product-linked information storage while maintaining low power consumption and operational simplicity.

The core hardware elements are summarized as follows:




Fig. 5.3: Prototype setup showing the Arduino Nano 33 BLE, SHT40 Sensor, EM4325 RFID tag, and the battery placed behind the sensor.

- Environmental sensor (SHT40): a digital temperature and humidity sensor used to monitor environmental conditions during harvesting operations.
- MCU (Arduino Nano 33 BLE): the MCU acts as the edge computing device responsible for sensor data acquisition, execution of embedded AI inference, and interaction with the RFID tag [207]. It integrates an ARM Cortex-M4 processor with 1 MB of flash memory and 256 KB of RAM, providing sufficient computational resources for real-time inference on constrained devices. During experimental validation, the board was powered via USB and an external battery; however, the architecture supports fully autonomous operation through integrated power sources and energy harvesting solutions.
- RFID Tag (EM4325): each harvest container is equipped with a passive UHF RFID tag providing a unique EPC for identification and additional user memory for storing AI-generated information. While operating primarily in passive mode, the tag also supports semi-passive operation when externally powered, enabling controlled write operations to user memory.
- MCU–Sensor Interface: an I²C wired connection between the MCU and the environmental sensor enables reliable real-time data acquisition. The sampling interval is configurable according to the application requirements, ranging from tens of minutes to several hours [96].
- MCU–RFID Interface: a wired interface between the MCU and the RFID tag allows direct read and write access to the tag user memory, ensuring that information generated on-device is securely stored on the physical product.

- **RFID Reader and Antenna:** during development and validation, a fixed RFID reader equipped with a dedicated antenna and the Universal Reader Assistant software was used to test EPC detection, user-memory access, and MCU–tag communication.
- **Handheld RFID Terminal (ORCA-50):** for operational use, a handheld RFID terminal [208] was employed to read EPCs and user-memory content from harvest containers and upload records to the backend platform.

The integration of the MCU, sensor, RFID tag, and power supply within the prototype setup is illustrated in Fig. 5.3 from [205]. The components are arranged to ensure compact wiring, mechanical robustness, and portability in a proof-of-concept configuration suitable for field deployment.

5.4.2 Software

The software stack supports end-to-end traceability by integrating embedded firmware, mobile applications, cloud services, and web-based interfaces. At the edge, custom firmware running on the MCU orchestrates sensor data acquisition, execution of embedded AI inference, and controlled updates of RFID user memory. This firmware layer constitutes a core element of the system, enabling autonomous operation without continuous network connectivity. On the reader side, a mobile application developed for the handheld RFID terminal manages EPC and user-memory scanning, performs data formatting, and uploads harvested records to the backend infrastructure. The backend platform is implemented using Supabase, which provides a managed PostgreSQL database, authentication mechanisms, and serverless functions for secure access control and real-time synchronization. A full-stack web application, named AgriTraceX, was developed using React, TypeScript, and Tailwind CSS to provide visualization and verification services. Through this interface, stakeholders can explore traceability records, inspect historical data, and generate QR codes for consumer-facing verification. The QR-code mechanism allows regulators and end users to retrieve product information directly, supporting transparency and anti-counterfeiting throughout the supply chain.

5.4.3 Embedded AI Module

The proposed traceability framework integrates an embedded AI module to enable on-device transformation of raw environmental measurements into validated information suitable for secure storage and traceability. Rather than introducing a new learning model, the framework directly relies on the AI model presented and validated in Chapter 4, where the full methodological details including chemical

analysis, model formulation, training procedure, and performance evaluation are provided. Within the context of the present chapter, the embedded AI module is treated as a reusable functional component of the overall architecture. Its role is to execute real-time inference on resource-constrained MCU hardware and to generate a compact, application-specific output that can be persistently associated with the physical product through RFID tag memory. This design choice avoids transmitting raw sensor data and instead stores only synthesized information, thereby reducing communication overhead and limiting exposure to data manipulation. By abstracting the AI model from the traceability infrastructure, the proposed framework remains independent of the specific learning task. Although the OTA risk assessment model described in Chapter 4 is used as a reference implementation, the same architecture can accommodate alternative embedded AI models targeting different agri-food monitoring objectives, such as quality classification, contamination detection, or compliance assessment. This modularity ensures that advances in embedded intelligence can be incorporated without modifying the core TRAC mechanisms.

5.5 Experimental Case Study: Grape Harvest Traceability

This section describes the experimental implementation of the proposed RFID-augmented Edge-AI traceability framework using grape harvesting as a representative case study. The objective is to detail how the proposed architecture was instantiated and executed in a controlled environment, demonstrating the feasibility of secure, product-linked traceability at the farm level. The evaluation of system performance and outcomes is intentionally deferred to a dedicated Results section. The embedded AI model employed during the experiments is the same as that presented and validated in Chapter 4 and is therefore not re-discussed here.

5.5.1 Case Study Context

Grape harvesting was selected as an experimental context due to its sensitivity to environmental conditions and its relevance to food safety, quality certification, and anti-counterfeiting practices. The harvesting phase represents the earliest point at which traceability information can be generated and bound to physical products. In this case study, individual harvest boxes were treated as traceable units, each equipped with a passive RFID tag used to store and transport validated information along the supply chain.




Fig. 5.4: Overall system architecture.

5.5.2 Experimental Setup

Table 5.1: Wiring configuration between MCU, sensor, and RFID tag.

Interface	Signal	Function
I ² C (Sensor–MCU)	SDA	Data line for communication.
	SCL	Clock line for synchronization.
	3V3	3.3 V power supply.
	GND	Common ground.
SPI (RFID Tag–MCU)	MISO	Tag → MCU data transfer.
	MOSI	MCU → Tag data transfer.
	SCLK	Serial clock signal.
	CS	Chip select for tag activation.
	VDD	Power supply for tag.
	GND	Common ground.

To realize the proposed architecture, a laboratory prototype was developed that replicates the operational conditions of grape harvesting while allowing controlled experimentation. The setup includes an environmental sensing unit, an edge computing device, RFID-enabled harvest boxes, a handheld RFID reader, and a backend platform for data management and verification. The overall system architecture and interactions between components are illustrated in Fig. 5.4 from [205]. Environmental data acquisition and on-device inference were performed using the same sensing configuration and embedded intelligence pipeline described in Chapter 4, ensuring consistency between the AI model and the traceability framework. The MCU periodically acquires temperature, humidity, and time-related information through a wired I²C interface, while a wired SPI interface enables read and write access to the RFID tag user memory. The wiring configuration between the MCU, sensor, and RFID tag is summarized in Table 5.1. A handheld RFID terminal was employed to scan EPCs and user-memory content and to synchronize collected records with the backend platform.

5.5.3 Implementation Workflow

The experimental workflow follows three main phases. Prior to harvesting, each empty harvest box is initialized by scanning the EPC of its attached RFID tag. The EPC serves as a unique identifier and is transmitted to the backend platform, where a corresponding product entry is created and anchored to the integrity layer. This step ensures that each physical container is registered and traceable before data collection begins. During the harvesting phase, the edge device executes periodic sensing




Fig. 5.5: Screenshot from the mobile reader application.

and on-device inference cycles. Instead of transmitting raw sensor measurements, the embedded AI module produces a compact representation of the observed conditions. At the end of the harvesting session, this synthesized information is written into the user memory of the RFID tag attached to the harvest box. A second RFID scan retrieves both the EPC and the updated user memory content, which are then uploaded to the backend and linked to the previously initialized product record. In the final phase, traceability verification is enabled through a QR-code mechanism. A QR code associated with the product allows stakeholders and consumers to access the corresponding backend record via a web-based interface. This process enables transparent inspection of harvest-related information while preserving the associ-

5.5.4 Data Encoding and Storage Strategy

Table 5.2: Encoding of supply chain and harvest data stored in the EM4325 RFID user memory.

Code	Example	Description	Size (chars)
V	V01	Payload version	2
S	PH	Stage (e.g., PH = Post-Harvest)	2
P	Grapes	Product type	12
SC	juicing-Grapes	Supply chain type	20
PR	ExampleName	Producer	12
F	Farm84923	Farm ID	12
L	4468	Lot number	6
D1	250920250800	Harvest start (DDMMYYYYHHMM)	12
D2	250920251300	Harvest end (DDMMYYYYHHMM)	12
Q	0/1/2	Quality (0 = good, 1 = mid, 2 = bad)	1
RF	0/1	Red-flag (1 = flagged, 0 = safe)	1
DM	RESERVED	Reserved field for future metadata	2+ 292 (the rest)
Example	V01SPHPGRAPESSCJUICING-GRAPESPREXAMPLENAMEFFARM84923 L4468D1250920250800D2250920251300Q2RF1DM		92

Given the limited memory capacity of passive UHF RFID tags, a compact encoding scheme was implemented to store essential traceability information within the user memory. Each data field is represented by a short prefix and fixed-length value, allowing product identifiers, harvest timestamps, quality indicators, and metadata to be embedded efficiently. The adopted encoding format ensures that critical information physically accompanies the product while remaining compatible with backend synchronization and verification processes. An overview of the encoding structure is provided in Table 5.2.

5.5.5 Backend Integration and Verification Mechanisms

As shown in Fig. 5.6 from [205], all records collected through RFID scans are synchronized with the backend platform, where structured storage and integrity verification are performed. Each record is linked to its corresponding EPC and anchored to a BC-inspired data structure that supports immutability and tamper detection. A web-based application provides interfaces for record visualization, auditing, and




Fig. 5.6: AgriTraceX front-end displaying product registration and supply-chain tracking information linked to RFID tags.

QR-code generation, enabling downstream stakeholders to verify product provenance and traceability without direct access to the underlying infrastructure.

5.6 Results

This section reports the experimental results obtained from the implementation and execution of the proposed RFID-augmented Edge-AI traceability framework. The focus is placed on system-level performance, on-device feasibility, RFID interaction, operational efficiency, and data integrity mechanisms. The evaluation of the embedded AI model itself including training methodology and classification accuracy is fully detailed in Chapter 4 and is therefore not repeated here.

5.6.1 On-Device Execution and Resource Utilization

The complete embedded pipeline comprising environmental sensing, on-device inference, and RFID tag memory updates was successfully deployed on the Arduino Nano 33 BLE platform. The integration confirmed that real-time processing and traceability operations can be executed entirely at the edge without reliance on continuous network connectivity.

Inference Latency

On-device inference was evaluated to assess suitability for real-time operation on resource-constrained hardware. Each inference operation, triggered immediately after sensor data acquisition, required less than 1 ms. Measured inference latency ranged between 650 and 700 μ s per prediction. This low latency confirms that embedded inference does not introduce perceptible delays and is compatible with downstream tasks such as RFID memory access and data synchronization.

Memory Footprint

The compiled firmware integrating sensor acquisition, embedded inference, and RFID communication occupied a moderate portion of the MCU resources. As summarized in Table 5.3, program storage utilized 37% of the available flash memory, while dynamic memory consumption reached 48% of the available RAM. These results demonstrate that the proposed framework comfortably fits within the constraints of low-cost MCU platforms, leaving sufficient margin for future extensions.

Table 5.3: On-device resource utilization on Arduino Nano 33 BLE Sense

Resource	Usage	Maximum Available
Flash	365,800 bytes (37%)	983,040 bytes
RAM	127,832 bytes (48%)	262,144 bytes

5.6.2 RFID Read and Write Performance

The performance of RFID operations was evaluated by analyzing both the MCU-to-tag writing process and the tag-to-reader data retrieval.

MCU-to-Tag Writing Performance

Writing performance was measured by evaluating the time required to store the complete encoded information payload into the EM4325 RFID tag user memory using a wired interface. Using the encoding structure defined in Table 5.2, the measured writing time per byte was approximately 241.5 μ s. This result confirms fast and reliable memory access, ensuring that information can be securely written at the end of harvesting without disrupting operational workflows.

Reader Application Performance

RFID reading performance was evaluated using a portable UHF RFID reader controlled by a custom Android application. The application reads EPC and user memory content, converts hexadecimal data into ASCII format, and uploads decoded records to the backend platform. The full read–decode–store pipeline operated consistently and without data loss. Preliminary tests were also conducted to assess read reliability under varying tag orientations. Experiments were performed in an isolated environment without nearby metallic interference. When the reader antenna and tag were aligned face-to-face, the maximum reliable read range reached approximately 1.5 m at an RF output power of 36 dBm, with a received RSSI of approximately –42 dBm. Similar performance was observed when the tag orientation

was tilted by $\pm 45^\circ$, indicating negligible sensitivity to orientation within the tested range. These results confirm stable read performance suitable for handheld harvesting applications.

5.6.3 Operational Efficiency and Energy Consumption

Operational efficiency was evaluated by analyzing the timing and energy consumption of the sensing–inference–RFID workflow during harvesting. Each operational cycle consists of sensor data acquisition, on-device inference, and RFID memory writing, all of which are completed within a few seconds. Given that harvesting sessions may last up to six hours, the cumulative overhead introduced by the system is negligible and does not interfere with standard field operations. In the current configuration, environmental data are sampled every 30 minutes. Each cycle includes approximately 10,000 μs for sensor acquisition, 700 μs for inference, and 22,218 μs for writing a 92-byte payload to the RFID tag, resulting in a total active time of roughly 33 ms per cycle. During this active phase, the MCU draws approximately 2.5 mA. For the remainder of the interval, the device operates in low-power sleep mode, consuming approximately 1 mA. Based on these measurements, the average current draw over a full cycle is close to 1 mA, as the system remains in sleep mode for more than 99% of the time. Using a compact 50 mAh lithium-ion battery, the device achieves approximately 50 hours of autonomy, corresponding to one or two complete harvesting sessions without recharging. These results confirm the feasibility of lightweight, battery-powered deployment in real harvesting environments. From a cost perspective, the MCU board used in the prototype costs between €15 and €25, the environmental sensor approximately €1–2, and passive RFID tags approximately €1–2 per unit. The RFID reader, priced between €400 and €700, is a shared component across multiple harvest units, resulting in a low amortized cost per container. Further cost reductions are expected through the development of a dedicated application-specific printed circuit board.

5.6.4 BC-Based Data Integrity Evaluation

A simulated BC environment was implemented locally to validate data integrity and traceability mechanisms under controlled conditions. The evaluation focused on metrics that remain meaningful in a non-distributed setup, namely record anchoring, hash validation, data integrity checks, and storage growth. Metrics such as transaction throughput and confirmation latency were deliberately excluded, as they are not representative in a simulated environment without consensus mechanisms. As summarized in Table 5.4, all records were successfully anchored and

validated, achieving a 100% success rate. Tampering attempts were consistently detected through hash mismatch verification, and storage growth remained minimal, with 2.11 KB corresponding to two recorded blocks. These results demonstrate that the proposed BC-inspired layer effectively supports record immutability and verifiable provenance, serving as a proof of concept for secure traceability.

Table 5.4: BC performance metrics considered in the simulated local environment.

Metric	Result Average
Record Anchoring & Hash Validation	100.0%
Data Integrity Test	Tamper Detected
Storage Growth	2.11 KB (2 blocks)

5.6.5 End-to-End Traceability Verification




Fig. 5.7: Screenshot of BC information displayed after scanning the product QR code.

End-to-end traceability was validated through the QR-code verification mechanism. After scanning a product's QR code, the system retrieves information previously written to the RFID tag, read by the RFID reader, and anchored to the backend integrity layer. Fig. 5.7 from [205] illustrates the information displayed through the BC interface, including harvest timing, producer identity, quality indicators, and red-flag status. This verification confirms correct operation of the complete traceability chain from edge data acquisition and on-device inference to RFID storage, backend anchoring, and consumer-facing visualization demonstrating the feasibility of secure, product-linked traceability across early supply-chain stages.

5.7 Discussion and Limitations

This section discusses the implications of the experimental results presented in the previous section, with particular emphasis on traceability robustness, security properties, and practical deployment constraints. While the Results section focused on measurable system behavior, the present discussion interprets these findings in the broader context of agri-food TRAC, and highlights the limitations that must be addressed to transition from a laboratory prototype to real-world deployment.

5.7.1 Security and Threat Mitigation Considerations

Beyond functional performance, the proposed RFID-Edge-AI framework was analyzed with respect to potential security threats affecting data authenticity, integrity, and trust across the traceability workflow. Several attack vectors commonly associated with RFID- and IoT-based systems were considered, along with corresponding mitigation strategies. A primary concern in RFID-enabled traceability systems is tag cloning, swapping, or unauthorized reuse. To mitigate this risk, the framework employs a dual-scan mechanism: an initial scan verifies that the RFID tag user memory is empty prior to data writing, while a second scan at the end of harvesting securely binds the generated information to the same physical tag. This approach ensures continuity between data acquisition and product identification. In addition, the architecture supports the use of the TID, a factory-programmed and immutable identifier embedded within each RFID chip. By storing the EPC-TID pair during initial registration and validating it during subsequent scans, the system can detect tag cloning or substitution attempts, thereby strengthening resistance against physical-layer counterfeiting. Data tampering during communication between the MCU, RFID reader, and backend represents another critical threat. This risk is mitigated through BC-inspired anchoring of all uploaded records. By chaining cryptographic hashes, any unauthorized modification of stored data results in a detectable

Table 5.5: Threat Analysis and Mitigation Strategies in the RFID–Edge-AI Pipeline

Threat	Preconditions	Likelihood	Impact	Mitigation (linked to Fig. 5.4 steps)
Tag swap	Physical access to harvest box	Medium	High	Dual-scan verification (Steps 1.a, 2.b) and EPC consistency check in backend (Step 2.d)
Relay attack	Proximity to reader during scan	Low	High	Timestamping and location logging at backend (Step 2.c)
RF shielding	Tag close to metal or moisture	High	Medium	Repeated scans and signal-strength monitoring (Step 2.b)
Sticker removal / reuse	Tag physically detached and reused	Medium	High	Password lock and EPC-to-batch link in backend (Step 2.d)
Password brute-force	Knowledge of tag password scheme	Low	Medium	Strong password policy and limited write attempts (Step 2.a)
EPC collision	Two tags with same EPC	Very Low	Medium	Backend EPC uniqueness check during registration (Step 1.b)

inconsistency, ensuring data integrity even in the presence of partial system compromise. At the RFID layer, unauthorized modification of tag content can be further prevented by locking the user memory with password protection, restricting write access to authorized devices only. The current implementation assumes a trusted backend environment, as the BC layer operates in a simulated and centralized configuration. While this setup is sufficient for validating immutability and traceability properties, a production deployment using a distributed or permissioned BC would further enhance resilience against backend manipulation and single-point failures. A structured threat analysis mapping attack scenarios to mitigation strategies is summarized in Table 5.5, providing a systematic overview of the system’s defensive posture. Overall, the framework does not claim absolute tamper resistance partic-

ularly given the use of passive UHF RFID tags, but instead provides best-effort authenticity verification through the combined use of physical safeguards, controlled writing mechanisms, and backend integrity checks. This layered approach aligns with practical security requirements in agri-food supply chains, where cost and deployability constraints limit the use of heavyweight cryptographic solutions at the field level.

5.7.2 System-Level Evaluation and Architectural Limitations

From a system perspective, the results confirm that integrating embedded intelligence, RFID-based product binding, and backend anchoring effectively addresses several limitations of existing traceability architectures. The direct association between edge-generated information and the physical product significantly reduces opportunities for data falsification and traceability discontinuities, particularly in low-connectivity environments. Moreover, the lightweight nature of the implementation lowers the barrier to adoption for small- and medium-sized producers, who are often excluded from complex and cloud-dependent traceability solutions. Nevertheless, several limitations emerged during the prototype evaluation. First, the BC component operates in a simulated environment and therefore does not capture the latency, scalability, and coordination overhead inherent to fully distributed deployments. While sufficient for demonstrating immutability and auditability, this setup does not reflect the operational complexity of real BC networks, which must be considered in future implementations. Second, the prototype relies on wired connections between the MCU, sensor, and RFID tag. Although this configuration ensures stability in laboratory conditions, it may introduce installation and durability challenges in agricultural environments exposed to moisture, dust, and mechanical stress. Future designs should explore more robust packaging or wireless alternatives while preserving energy efficiency. Third, although the edge device operates autonomously, the RFID reader still requires network connectivity to upload records to the backend. This dependency may limit system performance in areas with intermittent or unreliable internet access. Buffering strategies or delayed synchronization mechanisms could mitigate this constraint in future deployments. Finally, the experimental validation focused primarily on the harvesting stage. While the architecture is inherently modular and designed to span the entire agri-food supply chain including processing, packaging, and distribution scaling the system across heterogeneous nodes introduces additional challenges related to interoperability, synchronization, and data consistency. Addressing these challenges will be essential for large-scale adoption.

5.7.3 Data Protection and Privacy Considerations

Data protection and privacy were addressed through a combination of architectural and operational choices. Only processed indicators generated by on-device intelligence such as quality flags or risk labels are written to RFID tag memory, while raw sensor data remain confined to controlled environments or backend storage. This strategy reduces data volume, limits exposure of sensitive farm information, and mitigates privacy risks associated with large-scale IoT deployments. Backend access is governed through role-based permissions, allowing different stakeholders to view only the information relevant to their role within the supply chain. Consumer-facing QR codes expose a restricted subset of non-sensitive data, ensuring transparency without compromising proprietary or personal information. Additional safeguards include versioning control at the RFID and backend levels to prevent rollback attacks, whereby outdated or manipulated records could otherwise be reintroduced. Together, these measures balance transparency, traceability, and privacy, ensuring that the system complies with practical data protection requirements while maintaining trust among producers, regulators, and consumers.

5.8 Conclusion

This chapter presented a complete RFID-augmented Edge-AI framework aimed at strengthening traceability, authenticity verification, and data integrity in agri-food supply chains, with particular emphasis on the farm level, where traceability gaps are most critical. Building upon the embedded intelligence framework introduced in Chapter 4, this work extends the scope from localized risk assessment to end-to-end, product-linked traceability and anti-counterfeiting. By integrating on-device intelligence with passive RFID storage and a BC-inspired anchoring layer, the proposed system establishes a persistent and verifiable link between processed environmental information and the physical product at the earliest stage of the supply chain. While OTA risk estimation during grape harvesting was used as the validation scenario, the architecture itself is task-agnostic and modular, allowing adaptation to different sensing, classification, or compliance objectives across multiple supply-chain stages, including processing, storage, packaging, and distribution. The experimental evaluation confirmed the practical feasibility of the proposed approach under embedded and operational constraints. On-device sensing, inference, and RFID data preparation were executed within a short active window of approximately 33 ms per measurement cycle, resulting in low average power consumption and enabling battery-powered operation suitable for field deployment. The embedded implementation demonstrated efficient use of computational resources, leaving sufficient

headroom for additional functionality while maintaining real-time responsiveness. From a traceability perspective, the RFID-based product–data binding mechanism proved effective in ensuring that validated information physically follows the product. The dual-scan interaction model enabled intuitive and reliable initialization and retrieval of traceability records without requiring specialized technical expertise. Moreover, the transformation of raw sensor data into compact, high-level information directly at the edge significantly reduced storage and communication overhead, preventing unnecessary data accumulation and improving system scalability. Overall, this chapter demonstrates that combining embedded intelligence, RFID-based portable storage, and BC-inspired integrity mechanisms provides a viable and low-cost solution for secure, transparent, and energy-efficient traceability in agri-food systems. The proposed framework addresses key limitations of existing approaches and constitutes the main contribution of this thesis toward practical, trustworthy, and deployable traceability and anti-counterfeiting infrastructures.

Conclusions and Future Work

This thesis addressed the growing need for advanced TRAC solutions in AFSCs while explicitly accounting for the practical constraints of low-cost, low-power, and low-capacity devices. Current state-of-the-art approaches often depend on continuous cloud connectivity, extensive data storage, and centralized processing infrastructures, which limit their applicability in real-world agricultural contexts, particularly for small and medium-sized producers and rural environments. To overcome these limitations, this work proposed a data-centric and system-level approach that combines on-device intelligence, efficient data reduction, and secure identification mechanisms within a unified traceability framework.

The data distillation techniques presented in Chapter 3 demonstrated that substantial reductions in data volume can be achieved directly at the edge while preserving the information required for ML tasks. Experimental results showed that the proposed techniques achieved compression levels of up to 90% across both numerical sensor datasets and agri-food image datasets while maintaining comparable classification performance. Such reductions further decrease communication overhead, memory usage, and storage requirements, enabling efficient on-device inference and making learning feasible on MCU-based platforms. In addition, the reduced dataset size facilitates incremental and continual learning, allowing models to be updated progressively as new data become available without exceeding the limited memory capacity of embedded devices. This capability directly addresses one of the most critical bottlenecks of edge-based traceability systems, namely the feasibility of executing and adapting AI workloads on resource-constrained hardware without sacrificing reliability. Furthermore, by performing preprocessing and inference locally, the system reduces dependence on permanent cloud connectivity and improves robustness in low-bandwidth or intermittent-network scenarios.

The on-device AI studies presented in Chapter 4 were validated using complementary and application-appropriate strategies. Specifically, the fermentation stage classification tasks for kefir and tea were validated using publicly available datasets,

ensuring reproducibility and comparability with existing approaches. In contrast, the OTA risk assessment case study was validated through chemical analysis, providing ground-truth measurements and confirming the reliability of the proposed AI-based predictions under real environmental conditions. Building upon this foundation, the thesis demonstrated that AI models can be effectively embedded and executed on heterogeneous edge hardware, including low-power MCU and intelligent vision cameras. Across the proposed case studies, neural network models, including YOLOv10, were successfully deployed on MCU-based sensing platforms and smart cameras, enabling real-time inference directly at the point of data acquisition. These implementations confirm that meaningful analytics can be executed locally on resource-limited hardware, avoiding the latency, privacy concerns, and operational costs associated with remote processing.

The work then introduced an RFID-augmented Edge-AI traceability framework that integrates on-device data processing, low-cost sensing, RFID-based product identification, and a BC-inspired data integrity layer. System-level evaluation showed that the complete operational cycle—including environmental data acquisition, feature processing, on-device inference, and RFID memory writing—can be executed in approximately 33 ms, demonstrating near real-time responsiveness suitable for field operations. In contrast, the dual RFID scans used for product initialization and final verification require only a few seconds and represent the dominant temporal overhead of the process. This timing analysis confirms that the intelligence layer itself introduces negligible delay and that the framework can operate seamlessly during harvesting and processing activities without disrupting workflows. Importantly, because the essential information is stored directly on the RFID tag, the system remains functional even in offline or low-connectivity environments, enabling reliable operation directly in the field. From a comparative perspective, the proposed architecture offers clear advantages over conventional cloud-centric IoT or AI-based traceability systems. Unlike approaches that depend on continuous data transmission and remote inference, the proposed solution achieves lower latency, reduced bandwidth usage, improved energy efficiency, and deterministic reaction times by keeping computation at the edge. Moreover, the physical–digital binding enabled by RFID ensures that critical information travels with the product itself, addressing a key weakness of existing systems where data remain detached from the physical item as discussed in Chapter 5. Beyond performance evaluation, this thesis also enabled the identification of several practical vulnerabilities and system limitations that inform future improvements. The study highlighted constraints related to the limited memory capacity and finite write endurance of passive RFID tags, potential tag cloning or unauthorized duplication risks, sensitivity of AI models to environmental vari-

ability and dataset bias, and the trade-off between BC robustness and edge-device resource consumption. While mitigation strategies such as cryptographic hashing, backend anchoring, and physical–digital coupling reduce these risks, these findings demonstrate that secure and trustworthy traceability requires careful co-design of hardware, software, and data management layers. Explicitly identifying these weaknesses contributes to a more realistic and deployable system design rather than an idealized solution.

While the results obtained are promising, several research directions remain open. Although in Chapter 3 data distillation demonstrated strong efficiency gains, the current implementations are primarily based on average-based and histogram-based techniques. Future work will therefore focus on developing more robust and semantically informed distillation methods that preserve discriminative features while remaining computationally lightweight. In addition, in Chapter 5 the BC component was intentionally implemented as a blockchain-inspired layer to accommodate connectivity constraints. Future developments will integrate fully decentralized infrastructures, smart contracts, and lightweight consensus mechanisms suitable for edge environments. Further work will also explore stronger physical–digital binding mechanisms, including tamper-evident packaging, multimodal authentication, and sensor fingerprints, and will evaluate large-scale deployments across diverse agri-food sectors to assess scalability and long-term reliability. Overall, this thesis demonstrates that effective, trustworthy, and scalable agri-food traceability systems can be achieved without relying on high-cost infrastructure or continuous cloud connectivity. This work establishes a practical and deployable foundation for next-generation traceability and certification systems while clearly identifying the remaining challenges that motivate future research.

References

1. V. Dadi, S. R. Nikhil, R. S. Mor, T. Agarwal, and S. Arora, "Agri-food 4.0 and innovations: Revamping the supply chain operations," *Production Engineering Archives*, vol. 27, no. 2, pp. 75–89, 2021.
2. European Parliament and Council, "Regulation (EC) No 178/2002 of 28 January 2002 laying down the general principles and requirements of food law, establishing the European Food Safety Authority and laying down procedures in matters of food safety." *Official Journal of the European Communities*, vol. L31, pp. 1–24, Jan. 2002. [Online]. Available: [urlhttps://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32002R0178](https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32002R0178).
3. A. Mostaccio, G. M. Bianco, G. Marrocco, and C. Occhiuzzi, "Rfid technology for food industry 4.0: A review of solutions and applications," *IEEE Journal of Radio Frequency Identification*, vol. 7, pp. 145–157, 2023.
4. J. E. Hobbs, D. Bailey, D. L. Dickinson, and M. Haghiri, "Traceability in the canadian red meat sector: do consumers care?," *Canadian Journal of Agricultural Economics*, vol. 53, no. 1, pp. 47–65, 2005.
5. A. E. Wilcock and K. A. Boys, "Reduce product counterfeiting: An integrated approach," *Business Horizons*, vol. 57, no. 2, pp. 279–288, 2014.
6. J. Xie, C. Wan, A. Tolón Becerra, and M. Li, "Streamlining traceability data generation in apple production using integral management with machine-to-machine connections," *Agronomy*, vol. 12, no. 4, p. 921, 2022.
7. S. Violino, F. Pallottino, G. Sperandio, S. Figorilli, L. Ortenzi, F. Tocci, S. Vasta, G. Imperi, and C. Costa, "A full technological traceability system for extra virgin olive oil," *Foods*, vol. 9, no. 5, p. 624, 2020.
8. D. I. Onwude, G. Chen, N. Eke-Emezie, A. Kabutey, A. Y. Khaled, and B. Sturm, "Recent advances in reducing food losses in the supply chain of fresh agricultural produce," *Processes*, vol. 8, no. 11, p. 1431, 2020.
9. S. Charan and A. N. Panghal, "Importance of traceability in food supply chain for brand protection and food safety systems implementation," *Annals of Biology*, vol. 34, no. 2, pp. 111–118, 2018.
10. U. Mc Carthy, I. Uysal, R. Badia-Melis, S. Mercier, C. O'Donnell, and A. Ktenioudaki, "Global food security – issues, challenges and technological solutions," *Trends in Food*

- Science & Technology*, vol. 77, pp. 11–20, 2018.
11. X. Shang, R. Zhang, X. Zhu, and Q. Zhou, “Design theory, modelling and the application for the internet of things service,” *Enterprise Information Systems*, vol. 10, no. 3, pp. 249–267, 2016.
 12. M. Murphy, R. Carey, and L. Alexandra, “Building the resilience of agri-food systems to compounding shocks and stresses: A case study from melbourne, australia,” *Frontiers in Sustainable Food Systems*, vol. 7, p. 1130978, 2023.
 13. M. Merenda, C. Porcaro, and D. Iero, “Edge machine learning for ai-enabled iot devices: A review,” *Sensors*, vol. 20, no. 9, p. 2533, 2020.
 14. J. Qian, L. Ruiz-Garcia, B. Fan, J. I. R. Villalba, U. McCarthy, B. Zhang, Q. Yu, and W. Wu, “Food traceability system from governmental, corporate, and consumer perspectives in the european union and china: A comparative review,” *Trends in Food Science & Technology*, vol. 99, pp. 402–412, 2020.
 15. S. Geng, X. Liu, and R. Beachy, “New food safety law of china and the special issue on food safety in china,” *Journal of Integrative Agriculture*, vol. 14, no. 11, pp. 2136–2141, 2015.
 16. Y.-n. Wu and J.-s. Chen, “Food safety monitoring and surveillance in china: Past, present and future,” *Food Control*, vol. 90, pp. 429–439, 2018.
 17. Q. Tang, J. Li, M. Sun, J. Lv, R. Gai, L. Mei, and L. Xu, “Food traceability systems in china: The current status of and future perspectives on food supply chain databases, legal support, and technological research and support for food safety regulation,” *Bio-science Trends*, vol. 9, no. 1, pp. 7–15, 2015.
 18. J.-P. Ding, J.-K. Huang, X.-P. Jia, J.-F. Bai, S. Boucher, and M. Carter, “Direct farm, production base, traceability and food safety in china,” *Journal of Integrative Agriculture*, vol. 14, no. 11, pp. 2380–2390, 2015.
 19. EC Regulation and others, “No 1924/2006 of the european parliament and of the council of 20 december 2006 on nutrition and health claims made on foods,” *Official Journal of the European Union*, vol. 12, pp. 3–18, 2007.
 20. I. Pádua, A. Moreira, P. Moreira, F. M. de Vasconcelos, and R. Barros, “Impact of the regulation (eu) 1169/2011: Allergen-related recalls in the rapid alert system for food and feed (rasff) portal,” *Food Control*, vol. 98, pp. 389–398, 2019.
 21. G. C. Smith, J. D. Tatum, K. E. Belk, J. A. Scanga, T. Grandin, and J. N. Sofos, “Traceability from a us perspective,” *Meat Science*, vol. 71, no. 1, pp. 174–193, 2005.
 22. E. H. Golan, B. Krissoff, and F. Kuchler, “Food traceability: One ingredient in a safe and efficient food supply,” *Amber Waves*, 2004.
 23. E. H. Golan, B. Krissoff, F. Kuchler, L. Calvin, K. E. Nelson, and G. K. Price, “Traceability in the us food supply: Economic theory and industry studies,” tech. rep., U.S. Department of Agriculture, Economic Research Service, 2004.
 24. S. Dey, S. Saha, A. K. Singh, and K. McDonald-Maier, “Foodsqrblock: Digitizing food production and the supply chain with blockchain and qr code in the cloud,” *Sustainability*, vol. 13, no. 6, p. 3486, 2021.

25. A. S. Rajasekaran, M. Azees, and F. Al-Turjman, "A comprehensive survey on blockchain technology," *Sustainable Energy Technologies and Assessments*, vol. 52, p. 102039, 2022.
26. H. Li, B. Zhang, L. Zhang, Y. Xue, M. He, and C. Ren, "A food traceability framework for dairy and other low-margin products," *IBM Journal of Research and Development*, vol. 60, no. 5/6, pp. 10–1, 2016.
27. S. Qiao, Z. Wei, and Y. Yang, "Research on vegetable supply chain traceability model based on two-dimensional barcode," in *Proceedings of the 6th International Symposium on Computational Intelligence and Design*, vol. 1, pp. 317–320, 2013.
28. X.-t. Yang, J.-p. Qian, J. Li, Z.-t. Ji, B.-l. Fan, B. Xing, and W.-y. Li, "A real-time agro-food authentication and supervision system on a novel code for improving traceability credibility," *Food Control*, vol. 66, pp. 17–26, 2016.
29. L. M. Abenavoli, F. Cuzzupoli, V. Chiaravalloti, and A. R. Proto, "Traceability system of olive oil: A case study based on the performance of a new software cloud," *Agronomy Research*, vol. 14, no. 4, 2016.
30. R. Guido, G. Mirabelli, E. Palermo, and V. Solina, "A framework for food traceability: Case study—italian extra-virgin olive oil supply chain," *International Journal of Industrial Engineering and Management*, vol. 11, no. 1, pp. 50–60, 2020.
31. P. Vizza, G. Tradigo, P. Veltri, P. Lambardi, C. Garofalo, F. M. Caligiuri, G. Caligiuri, and P. H. Guzzi, "Tracking agricultural products for wellness care," in *Proceedings of the IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pp. 2064–2067, 2018.
32. A. Yadav, P. R. Kumar, B. Bhasuran, and I. R. Oviya, "A novel approach for classifying dna barcodes using ensemble nlp models," in *Proceedings of the International Conference on Research Methodology in Knowledge Management, Artificial Intelligence and Telecommunications Engineering (RMKMATE)*, pp. 1–5, 2023.
33. M. Y. Stoeckle, C. C. Gamble, R. Kirpekar, G. Young, S. Ahmed, and D. P. Little, "Commercial teas highlight plant dna barcode identification successes and obstacles," *Scientific Reports*, vol. 1, no. 1, p. 42, 2011.
34. I. Bruni, A. Galimberti, L. Caridi, D. Scaccabarozzi, F. De Mattia, M. Casiraghi, and M. Labra, "A dna barcoding approach to identify plant species in multiflower honey," *Food Chemistry*, vol. 170, pp. 308–315, 2015.
35. T. Gorini, V. Mezzasalma, M. Deligia, F. De Mattia, L. Campone, M. Labra, and J. Frigerio, "Check your shopping cart: Dna barcoding and mini-barcoding for food authentication," *Foods*, vol. 12, no. 12, p. 2392, 2023.
36. Y. Wu, M. Li, Y. Yang, L. Jiang, M. Liu, B. Wang, and Y. Wang, "Authentication of small berry fruit in fruit products by dna barcoding method," *Journal of Food Science*, vol. 83, no. 6, pp. 1494–1504, 2018.
37. B. Nath, F. Reynolds, and R. Want, "Rfid technology and applications," *IEEE Pervasive Computing*, vol. 5, no. 1, pp. 22–24, 2006.
38. N. S. S. Shobha, K. S. P. Aruna, M. D. P. Bhagyashree, and K. S. J. Sarita, "Nfc and nfc payments: A review," in *Proceedings of the International Conference on ICT in Business, Industry and Government (ICTBIG)*, pp. 1–7, 2016.

39. L. Kumari, K. Narsaiah, M. K. Grewal, and R. K. Anurag, "Application of rfid in agri-food sector," *Trends in Food Science & Technology*, vol. 43, no. 2, pp. 144–161, 2015.
40. C. Costa, F. Antonucci, F. Pallottino, J. Aguzzi, D. Sarriá, and P. Menesatti, "A review on agri-food supply chain traceability by means of rfid technology," *Food and Bioprocess Technology*, vol. 6, pp. 353–366, 2013.
41. F. Tian, "An agri-food supply chain traceability system for china based on rfid & blockchain technology," in *Proceedings of the 13th International Conference on Service Systems and Service Management (ICSSSM)*, pp. 1–6, 2016.
42. J. Ouyang, A. Sharif, Y. Yan, and Q. H. Abbasi, "Small conformal uhf rfid tag antenna for labelling fruits," in *Proceedings of the IEEE Asia-Pacific Microwave Conference (APMC)*, pp. 33–35, 2019.
43. F. Gandino, B. Montrucchio, M. Rebaudengo, and E. R. Sanchez, "Analysis of an rfid-based information system for tracking and tracing in an agri-food chain," in *Proceedings of the 1st Annual RFID Eurasia*, pp. 1–6, 2007.
44. F. Gandino, B. Montrucchio, M. Rebaudengo, and E. R. Sanchez, "On improving automation by integrating rfid in the traceability management of the agri-food sector," *IEEE Transactions on Industrial Electronics*, vol. 56, no. 7, pp. 2357–2365, 2009.
45. C. Amador, J.-P. Emond, and M. C. do Nascimento Nunes, "Application of rfid technologies in the temperature mapping of the pineapple supply chain," *Sensing and Instrumentation for Food Quality and Safety*, vol. 3, pp. 26–33, 2009.
46. A. Vergara, E. Llobet, J. L. Ramírez, P. Ivanov, L. Fonseca, S. Zampolli, A. Scorzoni, T. Becker, S. Marco, and J. Wöllenstein, "An rfid reader with onboard sensing capability for monitoring fruit quality," *Sensors and Actuators B: Chemical*, vol. 127, no. 1, pp. 143–149, 2007.
47. K. H. Eom, M. C. Kim, S. Lee, and C. W. Lee, "The vegetable freshness monitoring system using rfid with oxygen and carbon dioxide sensor," *International Journal of Distributed Sensor Networks*, vol. 8, no. 6, p. 472986, 2012.
48. P. B. Gandino, C. Bernardi, F. Demartini, B. Montrucchio, M. Rebaudengo, and E. R. Sanchez, "Agri-food traceability management using a rfid system with privacy protection," tech. rep., Dipartimento di Automatica e Informatica, 2020.
49. Y. Zhu, W. Gao, L. Yu, P. Li, Q. Wang, Y. Yang, and J. Du, "Research on rfid-based anti-counterfeiting system for agricultural production," in *Proceedings of the World Automation Congress*, pp. 351–353, 2010.
50. D. Pignini and M. Conti, "Nfc-based traceability in the food chain," *Sustainability*, vol. 9, no. 10, p. 1910, 2017.
51. M. Conti, "Food traceability in fruit and vegetables supply chain," in *Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS)*, pp. 1–5, 2020.
52. M. Conti, "Evo-nfc: Extra virgin olive oil traceability using nfc suitable for small-medium farms," *IEEE Access*, vol. 10, pp. 20345–20356, 2022.
53. M. P. Caro, M. S. Ali, M. Vecchio, and R. Giaffreda, "Blockchain-based traceability in agri-food supply chain management: A practical implementation," in *Proceedings of the IoT Vertical and Topical Summit on Agriculture (IOT Tuscany)*, pp. 1–4, 2018.

54. G. Xu, W. Shen, and X. Wang, "Applications of wireless sensor networks in marine environment monitoring: A survey," *Sensors*, vol. 14, no. 9, pp. 16932–16954, 2014.
55. V. K. Quy, N. V. Hau, D. V. Anh, N. M. Quy, N. T. Ban, S. Lanza, G. Randazzo, and A. Muzirafuti, "Iot-enabled smart agriculture: Architecture, applications, and challenges," *Applied Sciences*, vol. 12, no. 7, p. 3396, 2022.
56. A. Corallo, R. Paiano, A. L. Guido, A. Pandurino, M. E. Latino, and M. Menegoli, "Intelligent monitoring internet of things based system for agri-food value chain traceability and transparency: A framework proposed," in *Proceedings of the IEEE Workshop on Environmental, Energy, and Structural Monitoring Systems (EESMS)*, pp. 1–6, 2018.
57. Y. Lu, P. Li, and H. Xu, "A food anti-counterfeiting traceability system based on blockchain and internet of things," *Procedia Computer Science*, vol. 199, pp. 629–636, 2022.
58. D. Davchev, L. Kocarev, A. Carbone, V. Stankovski, and K. Mitreski, "Blockchain-based distributed cloud/fog platform for iot supply chain management," in *Proceedings of the International Conference on Advanced Computing, Electronics and Electrical Technology (CEET)*, 2018.
59. S. Balamurugan, A. Ayyasamy, and K. S. Joseph, "Iot-blockchain driven traceability techniques for improved safety measures in food supply chain," *International Journal of Information Technology*, vol. 14, no. 2, pp. 1087–1098, 2022.
60. K. Wongpatikaseree, P. Kanka, and A. Ratikan, "Developing smart farm and traceability system for agricultural products using iot technology," in *Proceedings of the IEEE/ACIS 17th International Conference on Computer and Information Science (ICIS)*, pp. 180–184, 2018.
61. W. Hong, Y. Cai, Z. Yu, and X. Yu, "An agri-product traceability system based on iot and blockchain technology," in *Proceedings of the IEEE International Conference on Hot Information-Centric Networking (HotICN)*, pp. 254–255, 2018.
62. P. Visconti, R. de Fazio, R. Velázquez, C. Del-Valle-Soto, and N. I. Giannoccaro, "Development of sensors-based agri-food traceability system remotely managed by a software platform for optimized farm management," *Sensors*, vol. 20, no. 13, p. 3632, 2020.
63. A. Fatima, G. V. Sharma, P. V. K. Reddy, and V. Harini, "Agricultural products traceability system," in *Hybrid and Advanced Technologies* (S. P. J. Christydass, N. Nurhayati, and S. Kannadhasan, eds.), pp. 408–412, CRC Press, 2025.
64. N. N. Misra, Y. Dixit, A. Al-Mallahi, M. S. Bhullar, R. Upadhyay, and A. Martynenko, "Iot, big data, and artificial intelligence in agriculture and food industry," *IEEE Internet of Things Journal*, vol. 9, no. 9, pp. 6305–6324, 2020.
65. M. Trabelsi, E. Casprini, N. Fiorini, and L. Zanni, "Unleashing the value of artificial intelligence in the agri-food sector: Where are we?," *British Food Journal*, vol. 125, no. 13, pp. 482–515, 2023.
66. S. Valarezo-Plaza, J. Torres-Tello, K. D. Singh, S. J. Shirtliffe, S. Deivalakshmi, and S.-B. Ko, "A novel optiuniquenola crop yield prediction on edge devices," *IEEE Transactions on AgriFood Electronics*, vol. 2, no. 2, pp. 436–444, 2024.

67. Z. Dakhia, M. Russo, and M. Merenda, "Ai-enabled iot for food computing: Challenges, opportunities, and future directions," *Sensors*, vol. 25, no. 7, p. 2147, 2025.
68. S. Violino, L. Ortenzi, F. Antonucci, F. Pallottino, C. Benincasa, S. Figorilli, and C. Costa, "An artificial intelligence approach for italian evoo origin traceability through an open source iot spectrometer," *Foods*, vol. 9, no. 6, p. 834, 2020.
69. G. Alfian, M. Syafrudin, U. Farooq, M. R. Ma'arif, M. A. Syaekhoni, N. L. Fitriyani, J. Lee, and J. Rhee, "Improving efficiency of rfid-based traceability system for perishable food by utilizing iot sensors and machine learning model," *Food Control*, vol. 110, p. 107016, 2020.
70. P. W. Khan, Y.-C. Byun, and N. Park, "Iot-blockchain enabled optimized provenance system for food industry 4.0 using advanced deep learning," *Sensors*, vol. 20, no. 10, p. 2990, 2020.
71. H. Chen, Z. Chen, F. Lin, and P. Zhuang, "Effective management for blockchain-based agri-food supply chains using deep reinforcement learning," *IEEE Access*, vol. 9, pp. 36008–36018, 2021.
72. G. S. Rao, I. S, L. P. G, S. S, and V. R. S, "Effective blockchain-based management of agri-food supply chains using deep reinforcement learning," in *Proceedings of the 8th International Conference on Trends in Electronics and Informatics (ICOEI)*, pp. 534–538, 2025.
73. D. Uike, M. Villavicencio-Mendoza, A. Gehlot, D. Verma, J. Panduro-Ramirez, and E. Susanto, "Blockchain integration with ai for agriculture product supply chain," in *Proceedings of the International Conference on System Modeling & Advancement in Research Trends (SMART)*, pp. 774–778, 2022.
74. N. P. Bhatta and N. Thangadurai, "Utilization of iot and ai for agriculture applications," *International Journal of Engineering and Advanced Technology*, vol. 8, pp. 2731–2735, 2019.
75. R. S. Alonso, I. Sittón-Candanedo, Ó. García, J. Prieto, and S. Rodríguez-González, "An intelligent edge-iot platform for monitoring livestock and crops in a dairy farming scenario," *Ad Hoc Networks*, vol. 98, p. 102047, 2020.
76. R. B. Ayed, M. Hanana, S. Ercisli, R. Karunakaran, A. Rebai, and F. Moreau, "Integration of innovative technologies in the agri-food sector: The fundamentals and practical case of dna-based traceability of olives from fruit to oil," *Plants*, vol. 11, no. 9, p. 1230, 2022.
77. D. Santos, F. Gonçalves, G. Reis, M. Santos, M. Saraiva, P. Durães, M. Maximiano, R. Gomes, V. Távora, and O. Remédios, "Using llms to bridge the gap between consumers and blockchain on a agro-food traceability platform: an architectural proposal," in *Procedia Computer Science*, vol. 256, pp. 319–326, 2025.
78. W. Zhang, E. Chen, M. Anderson, S. Thompson, and J. Lee, "Multimodal deep learning-based intelligent food safety detection and traceability system," *International Journal of Management Science Research*, vol. 8, no. 3, pp. 71–76.
79. H. Feng, X. Wang, Y. Duan, J. Zhang, and X. Zhang, "Applying blockchain technology to improve agri-food traceability: A review of development methods, benefits and challenges," *Journal of Cleaner Production*, vol. 260, p. 121031, 2020.

80. J. Lin, Z. Shen, A. Zhang, and Y. Chai, "Blockchain and iot based food traceability for smart agriculture," in *Proceedings of the 3rd International Conference on Crowd Science and Engineering*, pp. 1–6, 2018.
81. J. Jaiyen, S. Pongnumkul, and P. Chaovalit, "A proof-of-concept of farmer-to-consumer food traceability on blockchain for local communities," in *Proceedings of the International Conference on Computer Science and Applications in Agriculture (ICOSICA)*, pp. 1–5, 2020.
82. X. Yang, M. Li, H. Yu, M. Wang, D. Xu, and C. Sun, "A trusted blockchain-based traceability system for fruit and vegetable agricultural products," *IEEE Access*, vol. 9, pp. 36282–36293, 2021.
83. A. Marchese and O. Tomarchio, "A blockchain-based system for agri-food supply chain traceability management," *SN Computer Science*, vol. 3, no. 4, p. 279, 2022.
84. M. T. Gaudio, S. Chakraborty, and S. Curcio, "Agri-food supply-chain traceability: A multi-layered solution," in *Proceedings of the IEEE International Conference on DASC/Pi-Com/CBDCCom/CyberSciTech*, pp. 1–5, 2022.
85. G. Baralla, A. Pinna, and G. Corrias, "Ensure traceability in european food supply chain by using a blockchain system," in *Proceedings of the IEEE/ACM 2nd International Workshop on Emerging Trends in Software Engineering for Blockchain (WETSEB)*, pp. 40–47, 2019.
86. L. Wang, L. Xu, Z. Zheng, S. Liu, X. Li, L. Cao, J. Li, and C. Sun, "Smart contract-based agricultural food supply chain traceability," *IEEE Access*, vol. 9, pp. 9296–9307, 2021.
87. Q. Yao and H. Zhang, "Improving agricultural product traceability using blockchain," *Sensors*, vol. 22, no. 9, p. 3388, 2022.
88. L. Cocco, K. Mannaro, R. Tonelli, L. Mariani, M. B. Lodi, A. Melis, M. Simone, and A. Fanti, "A blockchain-based traceability system in agri-food sme: Case study of a traditional bakery," *IEEE Access*, vol. 9, pp. 62899–62915, 2021.
89. S. Dinde and S. Shirgave, "Improved food traceability for restaurant customers using blockchain technology," in *Proceedings of the International Conference on Advanced Technologies (ICONAT)*, pp. 1–7, 2023.
90. S. Rajput, A. Jadhav, J. Gadge, D. Tilani, and V. Dalgade, "Agricultural food supply chain traceability using blockchain," in *Proceedings of the International Conference on Innovative Trends in Information Technology (ICITIIT)*, pp. 1–6, 2023.
91. A. Shahid, A. Almogren, N. Javaid, F. A. Al-Zahrani, M. Zuair, and M. Alam, "Blockchain-based agri-food supply chain: A complete solution," *IEEE Access*, vol. 8, pp. 69230–69243, 2020.
92. K. Salah, N. Nizamuddin, R. Jayaraman, and M. Omar, "Blockchain-based soybean traceability in agricultural supply chain," *IEEE Access*, vol. 7, pp. 73295–73305, 2019.
93. J. Hao, Y. Sun, and H. Luo, "A safe and efficient storage scheme based on blockchain and ipfs for agricultural products tracking," *Journal of Computers*, vol. 29, no. 6, pp. 158–167, 2018.
94. I. Hasan, M. M. Habib, Z. Mohamed, and V. Tewari, "Integrated agri-food supply chain model: An application of iot and blockchain," *American Journal of Industrial and Business Management*, vol. 13, no. 2, pp. 29–45, 2023.

95. F. A. Sunny, P. Hajek, M. Munk, M. Z. Abedin, M. S. Satu, M. I. A. Efat, and M. J. Islam, "uniqueain applications," *IEEE Access*, vol. 10, pp. 59155–59177, 2022.
96. M. R. Sebti, Z. Dakhia, S. Carabetta, R. Di Sanzo, M. Russo, and M. Merenda, "Real-time classification of ochratoxin a contamination in grapes using ai-enhanced iot," *Sensors (Basel, Switzerland)*, vol. 25, no. 3, p. 784, 2025.
97. R. Singh and S. S. Gill, "Edge ai: A survey," *Internet of Things and Cyber-Physical Systems*, vol. 3, pp. 71–92, 2023.
98. A. Polino, R. Pascanu, and D. Alistarh, "Model compression via distillation and quantization," *arXiv preprint arXiv:1802.05668*, 2018.
99. B. Jacob, S. Kligys, B. Chen, M. Zhu, M. Tang, A. Howard, H. Adam, and D. Kalenichenko, "Quantization and training of neural networks for efficient integer-arithmetic-only inference," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2704–2713, 2018.
100. Y. Jiang, S. Wang, V. Valls, B. J. Ko, W.-H. Lee, K. K. Leung, and L. Tassiulas, "Model pruning enables efficient federated learning on edge devices," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 34, no. 12, pp. 10374–10386, 2022.
101. S. Han, J. Pool, J. Tran, and W. Dally, "Learning both weights and connections for efficient neural network," *Advances in neural information processing systems*, vol. 28, 2015.
102. J. Gou, B. Yu, S. J. Maybank, and D. Tao, "Knowledge distillation: A survey," *International Journal of Computer Vision*, vol. 129, no. 6, pp. 1789–1819, 2021.
103. X. Deng, H. Zhong, R. Ai, F. Feng, Z. Wang, and X. He, "Less is more: Improving llm alignment via preference data selection," *arXiv preprint arXiv:2502.14560*, 2025.
104. M. Marion, A. Üstün, L. Pozzobon, A. Wang, M. Fadaee, and S. Hooker, "When less is more: Investigating data pruning for pretraining llms at scale," *arXiv preprint arXiv:2309.04564*, 2023.
105. A. G. Accettola and M. Merenda, "Dataset distillation as an enabling technique for on-device training in tinymml for iot: an rfid use case," in *2023 8th International Conference on Smart and Sustainable Technologies (SpliTech)*, pp. 1–4, IEEE, 2023.
106. M. R. Sebti, A. Accettola, R. Carotenuto, and M. Merenda, "Dataset distillation technique enabling ml on-board training: Preliminary results," in *Annual Meeting of the Italian Electronics Society*, pp. 379–384, Springer, 2023.
107. M. R. Sebti, Z. Dakhia, A. Macheda, D. Iero, M. Russo, and M. Merenda, "On-device dataset distillation: The mnist use case and initial experiments on agri-food product classification," in *2025 10th International Conference on Smart and Sustainable Technologies (SpliTech)*, pp. 1–6, IEEE, 2025.
108. M. R. Sebti, K. Laabassi, M. Russo, and M. Merenda, "Lightweight on-device pre-processing for agri-food classification," in *Proceedings of the IEEE Conference on Automation, Future Energy and Electronics (CAFE 2025)*, (Italy), IEEE, 2025. Accepted for publication.
109. M. Merenda, V. Mazzullo, M. Princi, A. Martino, R. Carotenuto, and D. Iero, "Evaluation of obdii data contribution in tiny machine learning based driving behaviour mon-

- itoring,” in *Proceedings of the 2022 7th International Conference on Smart and Sustainable Technologies (SpliTech)*, (Split, Croatia), pp. 1–6, IEEE, July 2022.
110. G. Hinton, O. Vinyals, and J. Dean, “Distilling the knowledge in a neural network,” *arXiv preprint arXiv:1503.02531*, 2015.
 111. B. Salem *et al.*, “Principal component analysis (pca).,” *La Tunisie Medicale*, vol. 99, no. 4, pp. 383–389, 2021.
 112. R. N. Youngworth, B. B. Gallagher, and B. L. Stamper, “An overview of power spectral density (psd) calculations,” *Optical manufacturing and testing VI*, vol. 5869, pp. 206–216, 2005.
 113. P. Musa, F. Al Rafi, and M. Lamsani, “A review: Contrast-limited adaptive histogram equalization (clahe) methods to help the application of face recognition,” in *2018 third international conference on informatics and computing (ICIC)*, pp. 1–6, IEEE, 2018.
 114. L. Yann, “Mnist handwritten digit database,” *ATT Labs.*, 2010.
 115. S. S. Kadam, A. C. Adamuthe, and A. B. Patil, “Cnn model for image classification on mnist and fashion-mnist dataset,” *Journal of scientific research*, vol. 64, no. 2, pp. 374–384, 2020.
 116. I. Ozkan, M. Koklu, and R. Saraçoğlu, “Classification of pistachio species using improved k-nn classifier,” *Health*, vol. 23, p. e2021044, 2021.
 117. D. Singh, Y. S. Taspinar, R. Kursun, I. Cinar, M. Koklu, I. A. Ozkan, and H.-N. Lee, “Classification and analysis of pistachio species with pre-trained deep learning models,” *Electronics*, vol. 11, no. 7, p. 981, 2022.
 118. M. Koklu, “Pistachio dataset.” <https://www.muratkoklu.com/datasets/>, n.d. Accessed on [29/04/2025].
 119. K. Laabassi, M. A. Belarbi, S. Mahmoudi, S. A. Mahmoudi, and K. Ferhat, “Wheion based on a deep learning approach,” *Journal of the Saudi Society of Agricultural Sciences*, vol. 20, no. 5, pp. 281–289, 2021.
 120. M. Sohan, T. Sai Ram, and C. V. Rami Reddy, “A review on yolov8 and its advancements,” in *Proceedings of the International Conference on Data Intelligence and Cognitive Informatics*, (Singapore), pp. 529–545, Springer, 2024.
 121. M. Kanaga, K. Saipriya, and N. M. Meenachi, “A comparative study of deep learning-based segmentation techniques: U-net, segnet, and deeplabv3 for breast cancer detection by thermal imaging,” 2025.
 122. C. Yang, F. Zheng, D. Z. Al-Hamid, P. H. J. Chong, and P. Lam, “Deep learning-based advertisement replacement on dynamic background videos,” *Journal of Image and Graphics*, vol. 13, no. 1, 2025.
 123. Y. Tao, S. Xie, F. Xu, A. Liu, Y. Wang, D. Chen, Y. Pan, L. Huang, D. Peng, X. Wang, and Z. Yuan, “Ochratoxin a: Toxicity, oxidative stress and metabolism,” *Food and Chemical Toxicology*, vol. 112, pp. 320–331, 2018.
 124. S. Abdul Rahman, A. N. Izzah, R. M. Salim, N. Muhamad, A. Hamdan, and Y.-H. Leong, “A scoping review on fungus and mycotoxin studies in the building’s environment: Mycotoxin analysis by mass spectrometry,” *International Journal of Analytical Chemistry*, vol. 2024, p. 8581029, 2024.

125. R. Abedi-Firoozjah, H. Ebdali, M. Soltani, P. Abdolahi-Fard, M. Heydari, E. Assadpour, M. Azizi-Lalabadi, F. Zhang, and S. M. Jafari, "Nanomaterial-based sensors for the detection of pathogens and microbial toxins in the food industry: A review on recent progress," *Coordination Chemistry Reviews*, vol. 500, p. 215545, 2024.
126. W. Mu, G. A. Kleter, Y. Bouzembrak, E. Dupouy, L. J. Frewer, F. N. Radwan Al Natour, and H. J. P. Marvin, "Making food systems more resilient to food safety risks by including artificial intelligence, big data, and internet of things into food safety early warning and emerging risk identification tools," *Comprehensive Reviews in Food Science and Food Safety*, vol. 23, p. e13296, 2024.
127. J.-F. Yeh, K.-M. Lin, C.-Y. Lin, and J.-C. Kang, "Intelligent mango fruit grade classification using alexnet-spp with mask r-cnn-based segmentation algorithm," *IEEE Transactions on Agrifood Electronics*, vol. 1, pp. 41–49, 2023.
128. M. Merenda, C. Porcaro, and F. G. Della Corte, "Led junction temperature prediction using machine learning techniques," in *Proceedings of the 2020 IEEE 20th Mediterranean Electrotechnical Conference (MELECON)*, (Piscataway, NJ, USA), pp. 207–211, IEEE, 2020.
129. M. Merenda, G. Cimino, R. Carotenuto, F. G. Della Corte, and D. Iero, "Edge machine learning techniques applied to rfid for device-free hand gesture recognition," *IEEE Journal of Radio Frequency Identification*, vol. 6, pp. 564–572, 2022.
130. A. Fotia, A. Macheda, M. R. Sebti, C. Nunnari, and M. Merenda, "Design of a portable water pollutants detector exploiting ml techniques suitable for iot devices integration," in *Proceedings of the International Conference on Applications in Electronics Pervading Industry, Environment and Society*, (Genova, Italy), pp. 358–363, Springer Nature, 2023.
131. C. K. Yap and K. A. Al-Mutairi, "A conceptual model relationship between industry 4.0—food-agriculture nexus and agroecosystem: A literature review and knowledge gaps," *Foods*, vol. 13, p. 150, 2024.
132. E. P. on Contaminants in the Food Chain (EFSA CONTAM Panel), D. Schrenk, M. Bignami, L. Bodin, J. K. Chipman, J. del Mazo, B. Grasl-Kraupp, C. Hogstrand, L. Hoogenboom, J.-C. Leblanc, *et al.*, "Risk to human health related to the presence of perfluoroalkyl substances in food," *EFSA journal*, vol. 18, no. 9, p. e06223, 2020.
133. A. Nugroho, "The impact of food safety standard on indonesia's coffee exports," *Procedia Environmental Sciences*, vol. 20, pp. 425–433, 2014.
134. H. Egmond, R. Schothorst, and M. Jonker, "Regulations relating to mycotoxins in food: Perspectives in a global and european context," *Analytical and Bioanalytical Chemistry*, vol. 389, pp. 147–157, 2007.
135. World Health Organization, "Mycotoxins." <https://www.who.int/news-room/fact-sheets/detail/mycotoxins>, 2024. Accessed: 2025-10-25.
136. J. Varga, K. Rigó, J. Téren, and Á. Mesterházy, "Recent advances in ochratoxin research i. production, detection and occurrence of ochratoxins," *Cereal research communications*, vol. 29, no. 1, pp. 85–92, 2001.
137. E. Streit, G. Schatzmayr, P. Tassis, E. Tzika, D. Marin, I. Taranu, C. Tabuc, A. Nicolau, I. Aprodu, O. Puel, and I. P. Oswald, "Current situation of mycotoxin contamination

- and co-occurrence in animal feed—focus on europe,” *Toxins*, vol. 4, no. 10, pp. 788–809, 2012.
138. L. Pinotti, M. Ottoboni, C. Giromini, V. Dell’Orto, and F. Cheli, “Mycotoxin contamination in the EU feed supply chain: A focus on cereal byproducts,” *Toxins*, vol. 8, no. 2, p. 45, 2016.
 139. G. Schatzmayr and E. Streit, “Global occurrence of mycotoxins in the food and feed chain: facts and figures,” *World Mycotoxin Journal*, vol. 6, no. 3, pp. 213–222, 2013.
 140. IARC, *Ochratoxin A*, vol. 56. Lyon, France: International Agency for Research on Cancer, 1993.
 141. P. Battilani, N. Magan, and A. Logrieco, “European research on ochratoxin a in grapes and wine,” *International Journal of Food Microbiology*, vol. 111, pp. S2–S4, 2006.
 142. C. Jiang, J. Shi, and C. Zhu, “Fruit spoilage and ochratoxin a production by *Aspergillus carbonarius* in the berries of different grape cultivars,” *Food Control*, vol. 30, no. 1, pp. 93–100, 2013.
 143. M. Bau, G. Castella, M. R. Bragulat, and F. J. Cabañes, “Rflp characterization of *Aspergillus niger* aggregate species from grapes from europe and israel,” *International Journal of Food Microbiology*, vol. 111, pp. S18–S21, 2006.
 144. S. L. Leong, A. D. Hocking, and E. S. Scott, “*Aspergillus* species producing ochratoxin a: isolation from vineyard soils and infection of semillon bunches in australia,” *Journal of Applied Microbiology*, vol. 102, no. 1, pp. 124–133, 2007.
 145. A. Medina, R. Mateo, L. López-Ocaña, F. M. Valle-Algarra, and M. Jiménez, “Study of spanish grape mycobiota and ochratoxin a production by isolates of *Aspergillus tubingensis* and other members of *Aspergillus* section *Nigri*,” *Applied and Environmental Microbiology*, vol. 71, no. 8, pp. 4696–4702, 2005.
 146. J. D. Palumbo, T. L. O’Keeffe, S. J. Vasquez, and N. E. Mahoney, “Isolation and identification of ochratoxin a-producing *Aspergillus* section *Nigri* strains from california raisins,” *Letters in Applied Microbiology*, vol. 52, no. 4, pp. 330–336, 2011.
 147. R. Serra, A. Lourenço, P. Alípio, and A. Venâncio, “Influence of the region of origin on the mycobiota of grapes with emphasis on *Aspergillus* and *Penicillium* species,” *Mycological Research*, vol. 110, no. 8, pp. 971–978, 2006.
 148. D. Spadaro, S. Patharajan, A. Lorè, A. Garibaldi, and M. L. Gullino, “Ochratoxigenic black species of *Aspergilli* in grape fruits of northern italy identified by an improved pcr-rflp procedure,” *Toxins*, vol. 4, no. 2, pp. 42–54, 2012.
 149. D. Mitchell, R. Parra, D. Aldred, and N. Magan, “Water and temperature relations of growth and ochratoxin a production by *Aspergillus carbonarius* strains from grapes in europe and israel,” *Journal of Applied Microbiology*, vol. 97, no. 2, pp. 439–445, 2004.
 150. A. Esteban, M. L. Abarca, M. R. Bragulat, and F. J. Cabañes, “Effects of temperature and incubation time on production of ochratoxin a by black *Aspergilli*,” *Research in Microbiology*, vol. 155, no. 10, pp. 861–866, 2004.
 151. S. M. Romero, V. F. Pinto, A. Patriarca, and G. Vaamonde, “Ochratoxin a production by a mixed inoculum of *Aspergillus carbonarius* at different conditions of water activity and

- temperature," *International Journal of Food Microbiology*, vol. 140, no. 2–3, pp. 277–281, 2010.
152. Food and Agriculture Organization of the United Nations, "Fao website." <https://www.fao.org/>, 2025. Accessed: 25 October 2025.
153. W.-J. Kong, J.-Y. Li, F. Qiu, J.-H. Wei, X.-H. Xiao, Y. Zheng, and M.-H. Yang, "Development of a sensitive and reliable high performance liquid chromatography method with fluorescence detection for high-throughput analysis of multi-class mycotoxins in coix seed," *Analytica Chimica Acta*, vol. 799, pp. 68–76, 2013.
154. J. Skarkova, V. Ostry, F. Malir, and T. Roubal, "Determination of ochratoxin a in food by high performance liquid chromatography," *Analytical Letters*, vol. 46, pp. 1495–1504, 2013.
155. A. M. Timperio, P. Magro, G. Chilosi, and L. Zolla, "Assay of ochratoxin a in grape by high-pressure liquid chromatography coupled on-line with an esi-mass spectrometry," *Journal of Chromatography B*, vol. 832, pp. 127–133, 2006.
156. M. Ventura, I. Anaya, F. Broto-Puig, M. Agut, and L. Comellas, "Two-dimensional thin-layer chromatographic method for the analysis of ochratoxin a in green coffee," *Journal of Food Protection*, vol. 68, pp. 1920–1922, 2005.
157. M. Qu, Y. He, W. Xu, D. Liu, C. An, S. Liu, G. Liu, and F. Cheng, "Array-optimized artificial olfactory sensor enabling cost-effective and non-destructive detection of mycotoxin-contaminated maize," *Food Chemistry*, vol. 456, p. 139940, 2024.
158. L. Zhang, Q. Yang, and Z. Zhu, "The application of multi-parameter multi-modal technology integrating biological sensors and artificial intelligence in the rapid detection of food contaminants," *Foods*, vol. 13, p. 1936, 2024.
159. N. S. Jusoh, N. A. Mohd Said, F. Salam, N. Awaludin, A. A. Abdul Kadir, and M. R. Ismail, "An integrated handheld biosensing device with iot system for in-situ and rapid detection of multi-mycotoxins," in *Proceedings of the 2022 4th International Conference on Smart Sensors and Application (ICSSA)*, (Kuala Lumpur, Malaysia), pp. 11–16, IEEE, 2022.
160. M. Akhtar and T. Feng, "Iot-based detection of molded bread and expiry prediction using machine learning techniques," *EAI Endorsed Transactions on Creative Technologies*, vol. 9, p. 173972, 2018.
161. Y. Kim, S. Kang, O. S. Ajani, R. Mallipeddi, and Y. Ha, "Predicting early mycotoxin contamination in stored wheat using machine learning," *Journal of Stored Products Research*, vol. 106, p. 102294, 2024.
162. M. R. Sebti, S. Carabetta, M. Russo, and M. Merenda, "Ochratoxin a growth probability estimation in wine production using ai-powered iot devices," in *2023 IEEE Conference on AgriFood Electronics (CAFE)*, pp. 152–156, IEEE, 2023.
163. K. Makhoulf, Z. Hmidi, L. Kahloul, S. Benhrazallah, and T. Ababsa, "On the forecasting of body temperature using iot and machine learning techniques," in *Proceedings of the 2021 International Conference on Theoretical and Applicative Aspects of Computer Science (ICTAACS)*, (Skikda, Algeria), pp. 1–6, IEEE, 2021.

164. C. Rocha Rosa, C. Magnoli, M. E. Fraga, A. M. Dalcero, and D. M. N. Santana, "Occurrence of ochratoxin a in wine and grape juice marketed in rio de janeiro, brazil," *Food Additives and Contaminants*, vol. 21, pp. 358–364, 2004.
165. M. Miranda, J. Varela-Aldás, and G. Palacios-Navarro, "Comparison of blockly vs arduino ide for programming education using m5stack core2 esp32 iot development kit," in *Human Systems Engineering and Design (IHSED2021): Future Trends and Applications*, vol. 21, pp. 677–684, USA: AHFE Open Access, 2021.
166. P. D. Prasetyo Adi, A. Kitagawa, and J. Akita, "Finger robotic control use m5stack board and mqtt protocol based," in *Proceedings of the 2020 7th International Conference on Information Technology, Computer, and Electrical Engineering (ICITACEE)*, (Semarang, Indonesia), pp. 24–25, IEEE, Sept. 2020.
167. Y. B. Jonggo'man, "Automatic foldron t-shirt with proportional integral derivative method and m5stack for iot-server," *Internet of Things and Artificial Intelligence Journal*, vol. 1, pp. 81–99, 2021.
168. European Commission, "Regulation (eu) 2022/1370 amending regulation (ec) no 1881/2006 as regards maximum levels of ochratoxin a in certain foodstuffs," *Official Journal of the European Union*, vol. L 206, pp. 1–7, 2022.
169. Time and Date, "Time and date." <https://www.timeanddate.com>, 2024. Accessed: 4 February 2024.
170. Z. Xu, J. Ye, L. Su, and Z. Liu, "Growth and metabolism control of lactic acid fermentation based on internet of things," in *Journal of Physics: Conference Series*, vol. 1757, p. 012158, IOP Publishing, 2021.
171. G. Kimutai, A. Ngenzi, S. Rutabayiro Ngoga, R. C. Ramkat, and A. Förster, "An internet of things (iot)-based optimum tea fermentation detection model using convolutional neural networks (cnns) and majority voting techniques," *Journal of Sensors and Sensor Systems*, vol. 10, no. 2, pp. 153–162, 2021.
172. Q. Zhou *et al.*, "Monitoring black tea fermentation quality by intelligent sensors: Comparison of image, e-nose, and data fusion," *Food Bioscience*, 2023.
173. A. *et al.*, "Evaluation of the black tea taste quality during fermentation process using image and spectral fusion features," *Fermentation*, 2023.
174. A. Bhargava, A. Bansal, V. Goyal, and A. Shukla, "Machine learning & computer vision-based optimum black tea fermentation detection," *Multimedia Tools and Applications*, 2023.
175. J. Vošahlík and J. Hart, "Measurability of quality in fermentation process of rice wine by iot in the field of industry 4.0," 2021.
176. J. Vošahlík, "Iot and measurement of fermentation process of rice wine.," *Agronomy Research*, vol. 21, no. 3, 2023.
177. D. Buonocore, G. Ciavolino, S. Dello Iacono, and C. Liguori, "Online identification of beer fermentation phases," *Fermentation*, vol. 10, no. 8, p. 399, 2024.
178. J. Abijaude, P. Sobreira, L. Santiago, and F. Greve, "Improving data security with blockchain and internet of things in the gourmet cocoa bean fermentation process," *Sensors*, vol. 22, no. 8, p. 3029, 2022.

179. M. Wu, W. Liu, and S. Zheng, "Probiotic fermentation environment control under intelligent data monitoring," *SLAS technology*, p. 100153, 2024.
180. J. R. Saguidon, J. T. Dellosa, and R. N. Mendoza, "Internet-of-things (iot)-based automated temperature monitoring and control system for enhanced cacao fermentation quality," in *2024 6th International Conference on Electrical, Control and Instrumentation Engineering (ICECIE)*, pp. 1–7, IEEE, 2024.
181. A. L. Bowler, M. P. Pound, and N. J. Watson, "Domain adaptation and federated learning for ultrasonic monitoring of beer fermentation," *Fermentation*, vol. 7, no. 4, p. 253, 2021.
182. E. Cañete-Carmona, J.-J. Gallego-Martínez, C. Martín, M. Brox, J.-J. Luna-Rodríguez, and J. Moreno, "A low-cost iot device to monitor in real-time wine alcoholic fermentation evolution through co2 emissions," *IEEE Sensors Journal*, vol. 20, no. 12, pp. 6692–6700, 2020.
183. R. Faraco-Filho, J. V. De Castro, F. O. Barino, D. E. Campos, and A. B. Dos Santos, "Monitoring and characterization technology for coffee fermentation aromas: A data-driven approach," *IEEE Sensors Journal*, 2024.
184. T. Triyason and P. Setthawong, "Plakda-an iot platform for the production of mekong basin styled fermented fish (plara)," in *2016 2nd International Conference on Science in Information Technology (ICSITech)*, pp. 203–208, IEEE, 2016.
185. E. Klaithin, V. Matimapa-Kay, W. Daosud, and Y. Laoonguthai, "Effective an iot-based arduino design for automated bioreactor control and sample collection in biobutanol production," *IEEE Sensors Journal*, 2024.
186. M. A. R. Alif and M. Hussain, "Yolov1 to yuniqueof yolo variants and their application in the agricultural domain," *arXiv preprint arXiv:2406.10139*, 2024.
187. A. Wang, H. Chen, L. Liu, K. Chen, Z. Lin, J. Han, *et al.*, "Yolov10:uniqueect detection," *Advances in Neural Information Processing Systems*, vol. 37, pp. 107984–108011, 2024.
188. Luxonis, "Luxonis blob converter."
url<https://blobconverter.luxonis.com>, 2025. Accessed: 2025-07-02.
189. S. P. Plakantara and A. Karakitsiou, "Transforming agrifood supply chains with digital technologies: A systematic review of safety and quality risk management," *Operations Research Forum*, vol. 6, no. 3, p. 113, 2025.
190. E. Verna, G. Genta, and M. Galetto, "Enhanced food quality by digital traceability in food processing industry," *Food Engineering Reviews*, pp. 1–25, 2025.
191. K. Giannakas and A. Yiannaka, "Food fraud: Causes, consequences, and deterrence strategies," *Annual Review of Resource Economics*, vol. 15, no. 1, pp. 85–104, 2023.
192. E. Casadei, E. Valli, F. Panni, J. Donarski, J. F. Gubern, P. Lucci, and T. G. Toschi, "Emerging trends in olive oil fraud and possible countermeasures," *Food Control*, vol. 124, p. 107902, 2021.
193. D. Bose and M. Padmavati, "Honey authentication: A review of the issues and challenges associated with honey adulteration," *Food Bioscience*, vol. 61, p. 105004, 2024.
194. M. Mehdizadeh, A. Omid, S. Morya, and Z. Abideen, "Combating saffron fraud: A systematic review of adulteration practices, detection technologies, recommendations and challenges," *Critical Reviews in Food Science and Nutrition*, pp. 1–17, 2025.

195. J. O. Pesme, "Tracing and tracking wine bottles: Protecting consumers and producers," in *BIO Web of Conferences*, vol. 68, p. 03028, EDP Sciences, 2023.
196. D. Skalkos, "Prospects, challenges and sustainability of the agri-food supply chain in the new global economy ii," *Sustainability*, vol. 15, no. 16, p. 12558, 2023.
197. R. Colella, A. Arciello, G. Grassi, and M. Merenda, "Memristor-based circuits and architectures enabling next-generation neuromorphic rfid systems," *IEEE Journal of Radio Frequency Identification*, vol. 9, pp. 384–394, 2025.
198. G. Martino, M. T. Bevacqua, L. Catarinucci, and M. Merenda, "Towards multi-frequency and multi-view localization via uhf-rfid passive tags," in *Proceedings of the IEEE International Conference on RFID Technology and Applications (RFID-TA)*, pp. 78–80, 2024.
199. R. E. Venturini, "Technological innovations in agriculture: The application of blockchain and artificial intelligence for grain traceability and protection," *Brazilian Journal of Development*, vol. 11, no. 3, p. e78100, 2025.
200. A. Hamdi-Asl, H. Amoozad-Khalili, R. Tavakkoli-Moghaddam, and M. Hajiaghahi-Keshteli, "Toward sustainability in designing an agricultural supply chain network: A case study on palm date," *Scientia Iranica*, vol. 31, no. 18, pp. 1691–1709, 2024.
201. N. T. Nha Trang, T. T. Nguyen, H. V. Pham, T. T. A. Cao, T. H. Trinh Thi, and J. Shahreki, "Impacts of collaborative partnership on the performance of cold supply chains of agriculture and foods: Literature review," *Sustainability*, vol. 14, no. 11, p. 6462, 2022.
202. Z. Li, A. Jallow, S. Nidiaye, Y. Huang, Q. Zhang, P. Li, and X. Tang, "Improvement of the sensitivity of lateral flow systems for detecting mycotoxins: Up-to-date strategies and future perspectives," *Comprehensive Reviews in Food Science and Food Safety*, vol. 23, p. e13255, 2024.
203. D. V. Čepo, M. Pelajić, I. V. Vrček, A. Krivohlavek, I. Žuntar, and M. Karoglan, "Differences in the levels of pesticides, metals, sulphites and ochratoxin a between organically and conventionally produced wines," *Food Chemistry*, vol. 246, pp. 11–18, 2017.
204. M. Merenda, D. Iero, and F. G. Della Corte, "Cmos rf transmitters with on-chip antenna for passive rfid and iot nodes," *Electronics*, vol. 8, no. 12, p. 1448, 2019.
205. M. R. Sebti, A. Arciello, M. Russo, and M. Merenda, "An rfid-augmented information retrieval technique for ai-enabled iot devices ensuring agri-food traceability and anti-counterfeiting," *IEEE Journal of Radio Frequency Identification*, 2026.
206. K. Fan, W. Jiang, H. Li, and Y. Yang, "Lightweight rfid protocol for medical privacy protection in iot," *IEEE Transactions on Industrial Informatics*, vol. 14, no. 4, pp. 1656–1665, 2018.
207. Arduino, "Arduino Nano 33 BLE Sense Technical Specifications." <https://docs.arduino.cc/hardware/nano-33-ble/>, 2025. Accessed: Oct. 22, 2025.
208. Shenzhen RodinBell Technology Co., Ltd., "Orca 50 Handheld Data Terminal." <https://www.rodinbell.com/en/ProductDetail.html?ID=12423>, 2025. Accessed: Oct. 22, 2025.