

Next-Gen Agriculture: Enhancing Animal Welfare Through IoT, Edge Devices, and Artificial Intelligence

Benlahsiniya Maroua, Ait Abdelouahid Rachida, Marzak Abdelaziz

Department of Mathematics and Computer Science-Artificial Intelligence and Systems Laboratory-
Faculty of Sciences Ben M'Sik, Hassan II University of Casablanca, Casablanca, Morocco

Abstract—Agriculture has always been a field of innovation, evolving alongside technological advancements to boost productivity and address sector-specific challenges. Today, the Internet of Things (IoT), edge computing, and artificial intelligence (AI) are revolutionizing farming practices, particularly in livestock management. These technologies enable real-time monitoring of animal health, facilitate early disease detection, and optimize decision-making. Beyond immediate benefits, they contribute to more sustainable agriculture and improved animal welfare, issues of growing concern to consumers. However, adoption faces several hurdles, notably regarding cost, data management, and connectivity. Solutions are emerging to overcome these challenges, such as advanced technology integration, remote sensing, and improved digital infrastructure. In this article, we propose a hybrid Edge-Fog-Cloud architecture for the intelligent monitoring of animal welfare, based on a simulation using synthetic data (100,000 observations covering 100 simulated animals). In addition to standard classification metrics (accuracy, precision, recall, F1-score), the article presents a comprehensive factorial ablation study to determine the marginal contribution of each method, as well as indicators related to event-level detection (detection delay, alert fragmentation, false alert rate) and model performance in terms of resource usage and latency; the aim is to highlight the advantages of the proposed architecture and justify it with quantified, reproducible evidence, rather than relying solely on aggregate accuracy figures. The results presented here constitute a proof of concept based on simulated data, not a demonstration of agricultural performance under real-world conditions with live animals.

Keywords—Smart farming; animal welfare; IoT; edge computing; artificial intelligence; connected livestock farming

I. INTRODUCTION

Agriculture has always been at the forefront of technological innovation, adapting to advancements to improve production and address sector challenges. Today, with the advent of the Internet of Things (IoT), edge devices, and artificial intelligence (AI), a new agricultural revolution is underway. These technologies offer unprecedented opportunities to transform livestock farming practices and ensure animal welfare, an aspect often sidelined in the past.

Thanks to IoT and AI, farmers can monitor animal health in real time, detect diseases early, and even anticipate needs and behaviors [2], [3]. Furthermore, edge devices enable on-site data processing, thereby reducing latency and optimizing decision-making [4]. These advancements go beyond improving

animal health; they also help increase farm efficiency, ensure more sustainable production, and meet growing consumer expectations regarding animal welfare.

However, these technologies present major challenges (detailed in Section I) that still block their large-scale adoption. Following a state-of-the-art review of key technologies (Sections II through VI), this article proposes a hybrid Edge-Fog-Cloud architecture for monitoring animal and agricultural machinery to continuously collect data on the environment and welfare (Section VII), validates it through simulation including a factorial ablation study, event-level metrics, and a resource/latency characterization (Sections VIII and IX), discusses its limitations (Section X), and outlines prospects for embedded deployment (Section XI).

The proposed architecture in this study is validated solely through simulation using algorithmically generated synthetic data. The results should be viewed as a proof of concept demonstrating the internal consistency and relative merits of the design choices presented here, rather than as proof of performance under real-world livestock farming conditions; validation using real animal data is identified as the primary direction for future work (Sections X and XII).

Thus, next-generation agriculture, driven by IoT and AI, emerges as a response to contemporary challenges. It does more than just optimize productivity; it redefines how animals are raised, ensuring a more ethical and sustainable approach suited to the demands of the future.

II. TECHNOLOGICAL CONTEXT AND CHALLENGES OF MODERN AGRICULTURE

Modern agriculture faces major challenges linked to population growth, climate change, the weariness of natural resources, and the need to ensure sustainable food production. In this context, the integration of digital technologies, specifically the Internet of Things (IoT), artificial intelligence (AI), and edge computing, offers a strategic response to transform traditional agricultural systems into smart, adaptive, and resilient ones. As summarized by Friha et al., these tools "enable real-time monitoring of agro-environmental parameters and improve the precision of agricultural decision-making" [1]. AI, for its part, contributes significantly to "early disease detection and the objective assessment of animal welfare" [2], [3]. According to Srivastav and Das, edge computing enables "rapid, local processing of agricultural data, even in

environments with poor connectivity" [4]. Working in synergy, these technologies are redefining farming practices by boosting productivity, traceability, and farm resilience, while meeting growing societal demands regarding ethics, food safety, and environmental stewardship. However, as Vlaicu et al. point out, large-scale adoption "requires investment, reliable digital infrastructure, and appropriate human support" [3].

III. KEY TECHNOLOGIES: IOT, EDGE COMPUTING, AND ARTIFICIAL INTELLIGENCE

Digital technologies form the foundation of smart agriculture, particularly through the Internet of Things (IoT), edge computing, and artificial intelligence (AI). IoT enables the connection of sensors, embedded devices, and agricultural machinery to continuously collect data on the environment, crops, and livestock. These sensors measure parameters such as soil moisture, temperature, water consumption, and animal activity, thereby providing a rich dataset for decision-making. As Kamilaris and Prenafeta-Boldú note, "IoT constitutes an essential infrastructure for automating agricultural data collection on a large scale" [5]. Edge computing complements this approach by bringing data processing closer to the point of collection, which reduces latency, improves responsiveness, and decreases reliance on cloud infrastructure. Srivastav and Das highlight that this architecture "enables rapid local decisions, even in rural areas with poor connectivity" [4]. AI then comes into play to analyze the massive amounts of data generated by IoT, identifying complex patterns, predicting events (such as diseases or irrigation needs), and generating personalized recommendations. Wolfert et al. emphasize that "AI transforms agricultural data into actionable insights, enhancing the precision and sustainability of practices" [6]. Furthermore, Miller et al. (2025) demonstrate that integrating machine learning models into agricultural systems enables predictive crop monitoring and optimized resource management [7]. Together, these technologies form the backbone of Agriculture 4.0, making farms more autonomous, efficient, and resilient in the face of climatic and economic challenges.

IV. APPLICATION OF THESE TECHNOLOGIES TO SMART LIVESTOCK FARMING

The integration of the Internet of Things (IoT), artificial intelligence (AI), and edge computing into smart livestock farming is profoundly transforming animal management practices by placing data at the heart of decision-making. Thanks to connected devices such as biometric sensors, thermal cameras, directional microphones, and GPS collars, it is now possible to continuously monitor the health, behaviour, and environment of livestock. According to Neethirajan, these technologies enable "an objective assessment of animal welfare based on physiological and behavioral indicators" [2]. AI, coupled with these sensors, can detect early signs of stress, lameness, or respiratory disease by analyzing factors such as vocalizations, gait, or body temperature fluctuations [8]. Local data processing via edge devices allows for greater

As Wolfert et al. note, "interoperability between devices and platforms remains a major challenge, particularly in fragmented agricultural environments" [6]. Added to this is the issue of connectivity: in many rural areas, access to stable, high-speed

responsiveness by triggering alerts or automated actions without relying on a constant cloud connection, a feature particularly useful in rural areas with poor connectivity [4]. Furthermore, image recognition and computer vision systems are used to identify abnormal behaviors, such as aggression or feeding disorders, thereby facilitating individualized animal management [9]. These innovations extend beyond mere monitoring; they also enable the optimization of feeding, reproduction, pasture management, and biosecurity using predictive models and decision-support tools. In short, smart livestock farming is emerging as a key driver for enhancing productivity, animal welfare, and the sustainability of livestock systems.

V. BENEFITS FOR ANIMAL WELFARE AND SUSTAINABILITY

The adoption of digital technologies in agriculture, specifically the Internet of Things (IoT), artificial intelligence (AI), and edge computing, offers significant benefits for animal welfare and environmental sustainability. In livestock farming, the connected sensors described in Section III enable continuous monitoring of animal health and behaviour, thereby reducing stress associated with delayed human intervention and facilitating individualized, proactive herd management. Devices such as smart collars and thermal cameras detect early signs of illness, lameness, or heat stress, allowing for targeted treatments and a reduction in the preventive use of antibiotics [10]. This approach improves not only animal health but also the quality of livestock products.

From an environmental perspective, smart agriculture contributes to the more efficient use of natural resources. Automated irrigation systems, controlled by soil moisture sensors and predictive algorithms, can reduce water consumption by up to 30% while maintaining high yields [7]. Similarly, optimizing animal nutrition through the analysis of behavioural data helps minimize feed waste and reduce methane emissions [3]. AI also plays a role in crop planning and pasture rotation, contributing to soil regeneration and biodiversity preservation. Finally, the digital traceability of farming practices enabled by IoT meets societal expectations regarding transparency, ethics, and food safety [6]. By fostering agriculture that is more precise, responsive, and responsible, these technologies align fully with sustainable development goals.

VI. ADOPTION CONSTRAINTS AND TECHNOLOGICAL CHALLENGES

Despite promising advances in smart agriculture, the large-scale adoption of technologies such as the Internet of Things (IoT), artificial intelligence (AI), and edge computing remains hindered by a range of technical, economic, and organizational constraints. A major obstacle lies in the complexity of the required infrastructure: networking heterogeneous sensors, managing massive data flows, and integrating with existing systems demand specialized skills and substantial investment.

networks is limited, compromising the efficiency of IoT systems and the transmission of data to the cloud or edge devices [11].

Cybersecurity represents another critical issue. Connected devices are vulnerable to attacks, and protecting sensitive data

such as traceability information or animal biometric data has become a priority.

Recent research done on edge-enabled precision agriculture shows that, while data processing locally at the Edge is better for privacy than cloud-centric solutions, it does not replace robust security features to protect from data leakage or tampering in the Edge-Fog-Cloud pipeline [12].

Furthermore, a lack of standardization in communication protocols and data formats complicates the integration of multi-vendor solutions, limiting flexibility for farms. Finally, adopting these technologies requires upskilling farmers and technicians, who must be trained in the use, maintenance, and interpretation of digital tools. Without institutional support or suitable business models, these innovations risk exacerbating inequalities between well-equipped farms and those struggling to keep pace with the digital transition [3].

VII. EMERGING SOLUTIONS AND INTEGRATION PROSPECTS

To address the technical, economic, and organizational challenges hindering the adoption of digital technologies in agriculture, numerous emerging solutions aim to facilitate their gradual integration into farming systems. A major trend is the development of Edge AI, which involves running artificial intelligence models directly on edge devices (sensors, microcontrollers, smart cameras) without relying on the cloud. This approach reduces latency, enhances data privacy, and ensures real-time decision-making, even in environments with limited connectivity. According to Pal (2023), "Edge AI represents an effective response to bandwidth, security, and responsiveness constraints in agricultural IoT systems" [14]. Platforms such as NVIDIA Jetson Nano, Google Coral Edge TPU, and Qualcomm Edge Impulse now enable the deployment of lightweight AI models in the field, paving the way for robust, autonomous applications [15].

Concurrently, the rise of 5G and low-power, long-range networks (LoRaWAN, NB-IoT) is improving connectivity in rural areas, facilitating secure data transmission and device interoperability. This evolution is bolstered by the emergence of embedded AI microservices capable of over-the-air (OTA) updates and dynamic adaptation to field conditions [13]. Furthermore, digital cooperative initiatives and subscription models allow smallholders to access these technologies without bearing the full burden of investment costs. Finally, the convergence of AI, IoT, and Edge Computing is giving rise to AIoT (Artificial Intelligence of Things) systems capable of self-learning, predictive maintenance, and the continuous optimization of agricultural processes [16]. These developments outline a future of smarter, more inclusive, and resilient agriculture provided that public policy, training, and standardization support this technological transition.

A. Critical Positioning Relative to Previous Work and the Identified Gap

Sections I to VI have so far summarized the technological landscape of smart agriculture. In this subsection, we put our work in the context of specific previous work and not in the general trends we have discussed above.

Prior works on automated livestock behaviour analysis have mainly relied on computer vision. Lee et al. proposed a 3D camera-based system to pick up on aggressive behavior in pigs [8], and Hansen et al. developed a convolutional neural network for on-farm pig face recognition [9]. Both approaches can provide fine-grained behavioral or individual-level discrimination, but they come at the cost of camera infrastructure, sensitivity to lighting, occlusion, and pen layout, and the huge computational burden of running CNN inference, which typically requires GPU-class hardware at the Edge or Fog layer to be running in real time, not comparable to the lightweight statistical and shallow kernel approach we are using here. Conversely, video-based pipelines can recognize individual animals or specific postures, which is complementary but not a substitute for continuous monitoring of physiological states (body temperature, heart rate) of the entire herd.

The architecture proposed in this study aims to address a different aspect of this design space: instead of vision-based recognition of discrete behaviors, it performs continuous, lightweight statistical anomaly detection (Z-score, rolling standard deviation) directly on the low-dimensional physiological time-series at the Edge, and uses a non-linear classifier (RBF-kernel SVM) for the Fog layer, where we can use more compute but we need to keep the latency low. This design choice is intentional: it takes away the fine-grained behavioral resolution of vision-based methods but for much lower per-sensor compute, memory, and bandwidth (see Section IX-D), and does not capture purely visual or postural anomalies (e.g., lameness gaits) that a vision pipeline could detect. The gap that we aim to fill is not to cover the literature on livestock anomaly detection but to have an approach that is 1) combined with Edge-computed statistics into a Fog-level classifier rather than only using Edge and Fog as detectors and 2) quantifies the latency, accuracy, and resource trade-off of this fusion for a physiological (non-visual) sensor stream. In Sections VIII and IX, we explicitly quantify this trade-off against internal ablations and, where possible, we reference typical resource figures for CNN-based approaches for comparison, but we do not consider a head-to-head reimplementations of Lee et al. and Hansen et al. on this dataset because their models are designed for images not generated in our simulation, and this is pointed out in Section X.

VIII. PROPOSED EDGE AI-IOT ARCHITECTURE

The proposed architecture relies on a hybrid Edge-Fog-Cloud approach that integrates Internet of Things (IoT), Edge Computing, Fog Computing, and artificial intelligence technologies to ensure intelligent, real-time monitoring of animal welfare. This architecture aims to improve system responsiveness, reduce latency, and optimize the early detection of physiological and behavioral anomalies in smart livestock farming environments.

The first layer is the Edge layer, located directly at the farm site. It comprises various IoT sensors responsible for collecting physiological and environmental data, such as body temperature, ambient temperature, animal activity, and heart rate. This data is acquired in real-time via connected devices deployed on the animals or in their immediate environment. Communication is handled through wireless technologies suited

to agricultural environments, such as LoRaWAN, Wi-Fi, and 4G/5G networks (Fig. 1).

At the Edge level, initial local processing is performed to reduce data volume and enable a rapid response in the event of an anomaly. This layer performs several essential functions, including:

- Data acquisition,
- Noise filtering,
- Feature extraction (moving average, moving standard deviation, z-score),
- And initial local anomaly detection.

To improve detection speed, a statistical method based on the z-score is employed directly at the Edge level. This lightweight approach allows for the rapid identification of abnormal behaviors or measurements while limiting the computational load on embedded devices.

This baseline is calculated for each individual rather than as a whole animal, and the rolling mean and standard deviation of the z-score are calculated based on the previous history of each animal (Section VIII-B), so that the detector can adapt to the physiological baseline of each individual and not for the whole population. This decision directly accounts for the physiological variability of individual animals and is one of the reasons why the model applies to animals that are not seen during training (Section VIII-C, Section IX-A).

Hybrid Edge-Fog-Cloud Architecture for Intelligent Animal Welfare Monitoring




Fig. 1. Proposed hybrid Edge-Fog-Cloud architecture for intelligent animal welfare monitoring.

When an anomaly is detected locally, an immediate alert can be generated to ensure a quick response. These same statistical indicators (z-score, moving standard deviation) are not used solely for local alerts; they are also transmitted along with smoothed values to the Fog layer to enhance the classification performed there (see Section 8).

The preprocessed data is then transferred to the Fog layer, which acts as an intermediary between the Edge and the Cloud. This layer plays a pivotal role in data aggregation, decision fusion, and the execution of more advanced artificial intelligence algorithms. An RBF-kernel Support Vector Machine (SVM) model is integrated into this layer to confirm locally detected anomalies and improve the system's overall accuracy. The choice of a non-linear kernel is driven by the fact that the boundary separating normal and abnormal behaviors is not linearly separable within the physiological feature space employed.

The Fog layer also enables:

- Secure data aggregation,
- AI inference,
- Local decision-making,
- And temporary data management before synchronization with the Cloud.

For its part, the Cloud layer handles centralized storage, global data analysis, and large-scale monitoring of farming operations. Aggregated data from the lower layers is stored here

to facilitate advanced analysis, system performance tracking, and AI model updates. The computationally intensive training of the SVM model (see Section 9) also takes place at the Cloud level, after which the trained model is redeployed to the Fog layer for production inference. This layer also manages global dashboards, alerts, and user notifications.

Finally, decisions generated by the system are transmitted as alerts or intelligent recommendations to user interfaces accessible via web dashboards or mobile applications. This distributed architecture thus enables a combination of:

- Low latency through Edge processing,
- Improved accuracy via the SVM model enhanced by Edge features,
- Reduced network load,
- And intelligent, continuous monitoring of animal welfare.

The originality of this architecture lies in the hybrid integration of lightweight statistical detection based on Z-scores at the Edge level and SVM classification at the Fog level, leveraging both smoothed physiological values and Edge statistical indicators (Z-score, rolling standard deviation), thereby ensuring an effective trade-off between detection speed, classification accuracy, and computational resource optimization.

IX. EDGE AI-IOt IMPLEMENTATION AND SIMULATION

A. Experimental Simulation Pipeline




Fig. 2. Experimental Edge AI-IoT simulation pipeline for animal welfare monitoring.

Fig. 2 illustrates the experimental pipeline implemented to validate the Edge AI layer of the proposed architecture for intelligent animal welfare monitoring. The system relies on an Edge AI-IoT approach that enables the real-time collection, processing, and analysis of physiological and environmental data from the animals.

The experimental pipeline begins with a data collection phase using simulated sensors that measure four key parameters: body temperature, ambient temperature, animal activity, and heart rate across 100 animals, with 1,000 samples per animal (totaling 100,000 observations).

A significant methodological aspect concerns the generation of anomalous episodes: rather than drawing each unnatural sample independently, which would produce isolated point anomalies that are physiologically unrealistic, anomalies are simulated as episodes of varying duration (ranging from 15 to 60 consecutive samples per episode). This approach better reflects the reality of stress, fever, or lameness episodes, which unfold over a continuous period rather than at a single instant; it also prevents the Edge preprocessing step (moving average) from artificially diluting the anomaly signal.

B. Formal Data Generation Procedure

To make the synthetic data-generation process entirely reproducible, this subsection details the distributions, parameters, and randomization procedure that we described previously in a narrative.

Each animal has a physiological baseline in step 1. For each of the 100 simulated animals, we sample from a single baseline once.

Sensor	Baseline distribution	Unit
Body temperature	N (38.5, 0.3)	°C
Ambient temperature	N (20.0, 3.0)	°C
Activity	N (50.0, 8.0)	a.u.
Heart rate	N (65.0, 5.0)	bpm

Anomaly episode process (step 2). Anomalies are generated according to a two-state renewal process, rather than as i.i.d. point anomalies. At each time step without an anomaly, we initiate a new episode with probability p and draw its duration uniformly from the interval $[15, 60]$ (in number of samples). The probability p is calibrated analytically (rather than empirically tuned) to achieve the target ratio r of anomalous samples, using the stationary ratio of a two-state renewal process:

$$p = r / (\bar{L} \cdot (1 - r)), \text{ where } \bar{L} = (15 + 60) / 2 = 37.5 \text{ samples.}$$

For the target ratio $r = 0.18$ used throughout this work, this gives $p \approx 0.00585$ per time step. This formula and the resulting value of p were verified empirically to reproduce the target ratio (18.1% anomalous samples over 100,000 observations).

Measurement noise and anomaly amplitude (steps 3 and 4). Independent Gaussian measurement noise is added at each time step, and samples flagged as anomalous receive an additional shift justified by physiological considerations:

Sensor	Measurement noise σ	Anomaly offset (mean, σ)
Body temperature	0.35 °C	+0.9 °C ($\sigma = 0.5$)
Ambient temperature	0.8 °C	(environmental, not animal-driven)
Activity	9 a.u.	-10 a.u. ($\sigma = 7$)
Heart rate	6 bpm	+9 bpm ($\sigma = 6$)

The ambient temperature also follows a regular sinusoidal daily cycle (period = 1,000 samples, amplitude = 2°C) applied to each animal's specific baseline value; this cycle is independent of the anomaly process, as it depends on the

environment rather than the animal's state. Randomization and code availability. All random samplings use a fixed pseudo-random seed (NumPy generator, seed = 42) to ensure exact reproducibility. The complete generator code, along with the preprocessing (Edge), classification (Fog), ablation, and event metric calculation scripts used to produce all the quantitative results in this study, are provided as open-source Python code as part of the submission.

C. Animal-Level Dataset Partitioning for Evaluating Generalization

Since the 1,000 samples from a single animal are temporally linked (originating from the same baseline and generated from the same continuous time series), random splitting at the sample level would lead to information leakage from the training set to the evaluation set: two samples from the same animal, separated by only a few time steps, could end up in different sets despite being nearly identical. To avoid this, we partition the 100 animals at the animal level rather than the sample level, distributing the data according to a 70/15/15 ratio across the training, validation, and test sets (i.e., 70,000, 15,000, and 15,000 samples, respectively), ensuring that no animal contributes to more than one set. Consequently, every measure of accuracy, precision, recall, and F1-score presented in Section IX represents an evaluation based on data that are unseen at the animal level: the results demonstrate the classifier's ability to generalize to entirely unseen animals, rather than merely to unseen time points from animals used at least in part for its training.

X. EXPERIMENTAL RESULTS AND DISCUSSION

A. Simulation Results

The simulation was conducted on 100 animals with 1,000 samples per animal, totaling 100,000 observations, with anomaly episodes accounting for approximately 18% of the samples (compared to 12% in a purely point-based generation, as the duration of the episodes mechanically increases the total proportion of anomalous samples).

The data generated include:

- Body temperature,
- Ambient temperature,
- Animal activity,
- And heart rate.

Experimental results show that integrating the SVM model enhanced with Edge features into the Edge AI-IoT architecture enables significantly more reliable detection of behavioral and thermal anomalies compared to the initial approach based solely on smoothed raw physiological values (Fig. 3).

As mentioned in Table I, with an overall accuracy of 97.70%, a recall of 95.76%, and a precision of 91.86%, the system achieves an F1-score of 0.938, compared to 0.37 for the initial version of the pipeline (linear SVM based solely on smoothed physiological values and point anomalies).

Edge AI-IoT monitoring — Animal #22




Fig. 3. Real-time detection of physiological and behavioral anomalies using the proposed SVM-based Edge AI system.

TABLE I. SUMMARY OF SIMULATION RESULTS

Metric	Value
Number of simulated animals	100 (70 / 15 / 15 train / val / test, animal-level split)
Number of samples per animal	1 000
Total dataset size	100 000 observations
Overall SVM accuracy	97,70 %
Anomaly recall	95,76 %
Anomaly precision	91,86 %
Anomaly F1-score	0,938
Training time (offline, Cloud)	≈ 150,8 s
Prediction time (Fog inference)	≈ 0,19 ms/sample
Edge detection method	Z-score (per-animal baseline)
Classification type	Normal / Anormal

In subsection IX-B, we break down this improvement to determine the marginal contribution of each individual methodological change; this approach addresses the request of the first reviewer, who noted that the overall improvement (from 0.37 to 0.938) stemmed from three simultaneous changes, making it impossible to isolate the effect of each one.

B. Ablation Study

To isolate the individual contribution of each of the four methodological factors involved in the initial improvement (from 0.37 to 0.938) namely, the generation of anomaly episodes (point or episodic), the SVM kernel (linear or RBF), the enrichment of edge-related features (smoothed values only vs. smoothed values combined with rolling standard deviation and Z-score), and the decision threshold calibration (fixed vs. calibrated on the validation set) we conducted a full factorial ablation analysis covering $2^4 = 16$ configurations, while keeping the data split (training/validation/test) fixed at the animal level across all configurations. To facilitate the computations for these

16 configurations, we used a fixed random subsample of 30,000 training observations rather than the full set of 70,000 observations used for the main result presented in Table I; consequently, the absolute F1 values in this subsection are not directly comparable to those in Table I, although they share the same order of magnitude, which is the aspect of interest here (Table II).

TABLE II. MARGINAL EFFECT OF EACH FACTOR

Factor	Level A	Mean F1 (A)	Level B	Mean F1 (B)	Δ F1
Anomaly generation	Point	0.484	Episodic	0.753	+0.269
Edge feature set	Smoothed only	0.421	Enriched (+ rollstd + z-score)	0.816	+0.394
Threshold	Fixed (0.0)	0.585	Calibrated (val. F1-max)	0.652	+0.067
SVM kernel	Linear	0.628	RBF	0.609	-0.020

Enriching edge features is by far the most important factor (an average F1 score increase of 0.394, with low variance compared to the other three factors; standard deviation = 0.032), followed by the shift from generating point anomalies to episodic anomalies (+0.269 with high variance, standard deviation = 0.35, highlighting the critical importance of feature enrichment). Threshold calibration yields only a very slight and not always positive improvement (average +0.067); in two of the eight pairwise comparisons, calibration on the validation set slightly reduced the F1 score on the test set (we discuss the effect of variance later) (Fig. 4).

Regarding the chosen SVM kernel (linear vs. RBF), the average effect is negligible (-0.020) when the other three factors are considered. This is precisely what the architecture aims to achieve with the non-linear kernel in this simulation: it is primarily the edge-enriched feature vector that facilitates classification regardless of the classifier used rather than the use of a non-linear RBF kernel.




Fig. 4. Average marginal effects of individual factors on F1-score across the seven remaining factorial combinations.

As a further check, we replicated the configuration corresponding to the initial pipeline (point anomalies, linear kernel, smoothed-only features, calibrated threshold) and obtained an F1 score of 0.357, a value very close to the previously reported 0.37 and consistent with our baseline. The

configuration corresponding to the proposed final pipeline (episodic anomalies, RBF kernel, enriched features, calibrated threshold) also yielded a similar result, with an F1 score of 0.824 at the reduced training scale (30,000 samples) and 0.883 when trained on the full set (70,000 samples), as shown in Table I (the value obtained during the ablation study is expected to be slightly lower due to the reduced training scale).

C. Event-Level Evaluation

The precision, recall, and F1-score metrics calculated at the sample level (Table I) do not directly indicate the extent to which the set of welfare-related episodes is detected. This subsection presents metrics calculated at the event level on the test set (15 animals and 67 anomaly episodes constituting the ground truth): event-level recall (proportion of actual episodes with at least one overlapping alert), missed event rate, number of false alerts per 1,000 samples (sequences predicted as positive with no overlap with an actual episode), mean/median detection delay between the actual start of the episode and the first associated correct alert, and alert fragmentation (number of distinct alert sequences overlapping a single detected episode; a value of 1 indicates a continuous alert, whereas a value greater than 1 indicates an intermittent or unstable alert for the same episode).

TABLE III. EVENT-LEVEL PERFORMANCE COMPARISON: PROPOSED MODEL VS. LINEAR BASELINE

Metric	Linear baseline	Proposed (RBF, enriched, calibrated)
Event-level recall	100.0 %	100.0 %
Missed-event rate	0.0 %	0.0 %
False alarms / 1,000 samples	0.27	9.40
Mean detection delay (samples)	4.79	0.84
Median detection delay (samples)	4.0	0.0
Mean alert fragmentation	1.22	2.51
Episodes with fragmented (>1) alerts	16.4 %	65.7 %

Both models detect every reference episode at least once (100% event-level recall, 0% missed events) on this synthetic test set. Our model detects events much faster (median delay of 0 samples, meaning an alert is triggered at the very first abnormal sample, compared to 4 samples for the reference model), which is operationally significant (in our setup, a false negative, a missed well-being-related event, is costlier than a false positive). However, this speed comes at a cost that has not yet been quantified: our model generates more alerts per unit of time (9.40 vs. 0.27 per 1,000 samples), and alerts within a detected episode are more fragmented (65.7% of detected episodes trigger more than one burst of alerts, compared to 16.4% for the reference model). This poses a real challenge for field deployment: an alert stream that toggles on and off intermittently during a single physiological episode is harder to interpret than a continuous stream, an aspect that sample-level F1 scores fail to capture. We note that processing the alert stream at the "Fog" (network edge) layer using a short hysteresis window or a time-delay mechanism offers a potential solution

to this issue; we will address this in future work presented in Section X.

D. Resource Utilization and Latency Analysis

The Fog-level inference time of 0.19 ms/sample shown in Table I corresponds solely to the evaluation of the SVM decision functions; it does not include data acquisition, Edge-side preprocessing, network transmission, or alert transmission. Table IV breaks down the measured simulation times, the deployed model's memory footprint, and network latencies; the latter are derived from the literature and were not measured during this simulation, as no physical network or embedded hardware was used.

TABLE IV. LATENCY AND RESOURCE ANALYSIS BY ESTIMATION SOURCE

Pipeline stage	Estimate	Source
Edge preprocessing (smoothing, rolling std., z-score)	≈ 0.0013 ms/sample	Measured (vectorized simulation; indicative of arithmetic cost, not embedded-hardware timing)
Fog inference (SVM decision function)	≈ 0.19 ms/sample	Measured (this simulation, Table I)
Edge→Fog transmission (local network/gateway)	$\approx 1\text{--}10$ ms (typical LAN/short-range)	Literature-typical range; not measured
Cloud round-trip (cellular/WAN), for comparison	$\approx 50\text{--}200$ ms round-trip	Literature-typical cellular range; not measured
Sensing/acquisition (ADC sampling)	Not evaluated	Out of Scope: No Physical Hardware Implementation in This Work
Alert delivery to dashboard/mobile app	Not evaluated	Out of Scope: No Physical Hardware Implementation in This Work

Model memory footprint. The trained RBF-kernel SVM model retains 5,944 support vectors (out of 70,000 training samples, or approximately 8.5%) in the 12-dimensional augmented feature space, resulting—after serialization—in a

deployment artifact of 675 KB (including both the model and the feature normalizer). This size is well within the RAM capacity of a Fog-class gateway (e.g., Raspberry Pi-style hardware with hundreds of megabytes or even gigabytes of RAM), but it must be viewed in the context of the typical SRAM available on a highly resource-constrained microcontroller (such as the ESP32, which has approximately 520 KB of total SRAM shared between the operating system and application code); this confirms and quantifies the TinyML deployment challenge, the qualitative aspects of which were previously highlighted in Section XI.

E. Comparison of Edge-Only, Cloud-Only, and Edge-Fog-Cloud Configurations

To evaluate whether the proposed three-tier distribution is preferable to simpler configurations, Table V compares three scenarios: 1) "Edge-only," where only the Z-score-based detector issues alerts and no Fog or Cloud resources are involved in real-time detection; 2) "Cloud-only," where raw or lightly preprocessed data is transmitted over a wide-area network and the same SVM model is used for classification but evaluated centrally rather than at the Fog gateway; and 3) "Edge-Fog-Cloud."

Two points are explicitly clarified here rather than left implicit. First, accuracy rates are identical for the "Cloud-only" and "Edge-Fog-Cloud" configurations, as both utilize the same trained SVM classifier; the difference lies in where the classifier executes and the round-trip time, not in the model's quality. Second, the bandwidth-related argument favoring Edge preprocessing is modest in this specific application: the enriched feature vector transmitted from the Edge to the Fog (12 floating-point values) is not significantly lighter than the four raw sensor readings it replaces. Thus, unlike use cases involving images or video, the primary quantified benefit of the placement strategy proposed in this simulation lies in latency and offline resilience, rather than bandwidth reduction; this point is explicitly noted to avoid overstating an advantage that was not actually measured.

TABLE V. QUALITATIVE AND QUANTITATIVE EVALUATION OF ARCHITECTURAL CONFIGURATIONS

Configuration	Detection quality (F1)	Per-sample latency (detection path)	Resilience to connectivity loss	Note
Edge-only (Z-score)	0.268 (measured)	≈ 0.0013 ms (measured, Edge only)	Full no network dependency for alerts	Fast and resilient, but poor detection quality alone (Table III of the original submission)
Cloud-only (SVM, WAN round-trip)	0.883 (measured, same model)	$\approx 50\text{--}200$ ms (literature-typical WAN round-trip) + 0.19 ms inference	No detection stops entirely if connectivity is lost	Same accuracy as proposed, but detection latency dominated by network, and no offline operation
Proposed Edge-Fog-Cloud	0.883 (measured)	≈ 0.0013 ms (Edge) + 1–10 ms (local Fog link, literature-typical) + 0.19 ms (Fog inference)	Partial Fog continues real-time detection if Cloud/internet is lost; only training/aggregation is affected	Best measured trade-off: near Cloud-only accuracy, near Edge-only resilience and latency order-of-magnitude

XI. DISCUSSION

The results demonstrate the value of integrating Edge Computing and artificial intelligence into smart farming systems dedicated to animal welfare. Local data processing reduces latency and reliance on the Cloud while improving the speed of anomaly detection.

The hybrid approach combining the Z-score, rolling standard deviation, and an RBF-kernel SVM offers several advantages

now associated with a precise quantitative result, rather than mere qualitative assertions.

- Rapid detection of abnormal behaviors at the Edge,
- High recall (95.76%), prioritizing the early detection of critical situations,
- High precision (91.86%), limiting the number of false alerts sent to farmers,

- And improved real-time system responsiveness, with an inference time in the sub-millisecond range per sample, making it suitable for Fog deployment.

However, this gain in precision comes at a cost that must be explicitly noted: switching from a linear SVM to an RBF kernel significantly increases training time (≈ 150.8 s for 100,000 samples, compared to a few seconds for the initial linear model) and the complexity of the stored model. This cost remains acceptable within the proposed architecture, as training is performed offline in the Cloud (Section 7), while only inference, which is much less computationally intensive, is deployed at the Fog layer. Nevertheless, this factor must be considered when selecting Fog hardware and determining the model's retraining frequency in a production environment.

Furthermore, increasing the number of simulated animals and the volume of data strengthened the robustness of the experiment and allowed for an assessment of the architecture's ability to handle large data streams within a smart farming context.

Certain limitations remain:

- The use of simulated data: the generation of anomaly episodes (duration, physiological amplitude) relies on reasonable assumptions that were not calibrated against real animal data,
- The lack of actual deployment on a farm,
- The absence of a detailed energy assessment of the RBF-kernel SVM model on actual Fog hardware,
- The trade-off between increased accuracy and higher computational costs during training, which requires more precise quantification across different hardware profiles,
- And the lack of comprehensive experimental validation of the Cloud layer.
- Alert fragmentation: 65.7% of correctly detected episodes trigger multiple distinct bursts of alerts rather than a continuous alert (Section IX-C), which would be inconvenient for a livestock farmer receiving these alerts under real-world conditions; a post-processing step such as debouncing or hysteresis at the "Fog" layer represents a direct, low-cost mitigation solution to consider for future work;
- We did not perform a direct comparison with external vision-based methods (e.g., Lee et al. [8], Hansen et al. [9]) because those methods require input images derived from physiological sensor simulation data we do not generate (Section VI-B). This would require either a shared multimodal dataset or a new vision-based simulation pipeline; we leave this task for future work.
- The comparison between the "Edge"-only, "Cloud"-only, and "Edge-Fog-Cloud" approaches (Section IX-E) is based on network latency data typical of the literature rather than physical network hardware; it will need to be re-evaluated once field deployment has taken place (see also below) and future work has been completed.

Despite these limitations, the results demonstrate the value of integrating Edge AI, the fusion of Edge and Fog statistical features, and machine learning models for agricultural applications requiring intelligent, continuous, and real-time monitoring of animal welfare.

XII. TINYML DEPLOYMENT AND EDGE OPTIMIZATION

The growing adoption of Edge AI in smart agricultural environments has heightened interest in TinyML techniques, which enable the deployment of lightweight artificial intelligence models directly onto low-power embedded devices. In livestock monitoring applications, TinyML can help reduce latency, bandwidth consumption, and reliance on the cloud, while ensuring real-time anomaly detection at the Edge.

In the proposed architecture, the RBF-kernel SVM model used at the Fog layer offers superior performance but requires more memory than a linear model; this presents a specific challenge for TinyML deployment directly onto highly resource-constrained Edge devices (such as ESP32s or microcontrollers).

It is no longer merely a qualitative issue: Section IX-D presents the deployed model with 5,944 support vectors and a serialized memory footprint of 675 KB; although this size is suitable for "Fog"-class gateways, it far exceeds the total SRAM capacity (approximately 520 KB) typical of an ESP32-class microcontroller, making deployment at the most constrained "Edge" level impossible without appropriate model adjustment.

Optimization techniques such as reducing the number of support vectors, approximating the RBF kernel using random features, model quantization, or converting models to embedded formats (e.g., TensorFlow Lite for neural network-based variants) must be evaluated to determine the optimal trade-off between accuracy and memory footprint, depending on the layer (highly constrained Edge vs. Fog) targeted for inference.

Although this work focuses primarily on the simulation and validation of the proposed Edge AI architecture, future work will address the actual deployment of TinyML models on embedded Edge devices. This will involve evaluating inference performance, optimizing memory usage, and analyzing energy consumption for intelligent animal welfare monitoring systems.

XIII. CONCLUSION AND FUTURE OUTLOOK

This article proposed a hybrid Edge-Fog-Cloud architecture integrating IoT, Edge AI, and artificial intelligence technologies for the intelligent monitoring of animal welfare. The proposed approach relies on local anomaly detection based on the Z-score enhanced by a moving standard deviation and an RBF-kernel SVM model that jointly leverages Edge and Fog features to classify abnormal behaviors in smart livestock farming environments.

An experimental simulation based on realistic anomaly episodes was conducted to validate the proposed Edge AI pipeline. The results (accuracy 97.70%, recall 95.76%, precision 91.86%, F1-score 0.938) confirm the value of integrating Edge Computing, statistical feature fusion, and non-linear AI models for agricultural applications requiring real-time monitoring and low latency.

Future work will focus primarily on validating these results using real animal data, optimizing the SVM model for TinyML deployment on highly resource-constrained Edge devices, deploying the architecture in actual agricultural environments, comprehensively evaluating the energy consumption of Edge and Fog devices, and integrating advanced techniques such as Federated Learning and distributed AI models.

REFERENCES

- [1] Friha, O., Ferrag, M. A., Shu, L., Maglaras, L., & Chatterjee, M. (2021). Internet of Things for the Future of Smart Agriculture: A Comprehensive Survey of Emerging Technologies. *IEEE/CAA Journal of Automatica Sinica*, 8(4), 718-732. <https://doi.org/10.1109/JAS.2021.1003925>
- [2] Neethirajan, S. (2023). Artificial Intelligence and Internet of Things for Animal Welfare: A Review. *Frontiers in Animal Science*, 4. <https://doi.org/10.3389/fanim.2023.1166380>
- [3] Vlaicu, A. B., Szabo, C., & Dinu, I. (2024). Smart Livestock Farming: A Review of Sustainable Approaches Using IoT and AI. *Information*, 15(2), 84. <https://doi.org/10.3390/info15020084>
- [4] Srivastav, A. K., & Das, P. (2025). Edge Computing and AI in Agricultural IoT. In *Biotechnology and IoT in Agriculture and Food Production* (pp. 437-456). Springer. https://doi.org/10.1007/979-8-8688-1469-3_21
- [5] Kamilaris, A., & Prenafeta-Boldú, F. X. (2018). Deep learning in agriculture: A survey. *Computers and Electronics in Agriculture*, 147, 70-90. <https://doi.org/10.1016/j.compag.2018.02.016>
- [6] Wolfert, S., Ge, L., Verdouw, C., & Bogaardt, M.-J. (2017). Big Data in Smart Farming - A review. *Agricultural Systems*, 153, 69-80. <https://doi.org/10.1016/j.agsy.2017.01.023>
- [7] Miller, T., Mikiciuk, G., Durlík, I., et al. (2025). The IoT and AI in Agriculture: The Time Is Now - A Systematic Review of Smart Sensing Technologies. *Sensors*, 25(12), 3583. <https://doi.org/10.3390/s25123583>
- [8] Lee, J., Jin, L., & Park, D. (2020). Automated recognition of aggressive behavior in pigs using a 3D camera. *Computers and Electronics in Agriculture*, 169, 105240. <https://doi.org/10.1016/j.compag.2019.105240>
- [9] Hansen, M. F., Smith, M. L., Smith, L. N., Salter, M. G., Baxter, E. M., & Farish, M. (2018). Towards on-farm pig face recognition using convolutional neural networks. *Computers in Industry*, 98, 145-152. <https://doi.org/10.1016/j.compind.2018.03.032>
- [10] Largouët, C. et al. (2023). L'intelligence artificielle au service du bien-être animal. Institut Agro Rennes-Angers. <https://www.institut-agro-rennes-angers.fr>
- [11] Sapna, Gauttam, H., Chauhan, V., Pattanaik, K. K., Trivedi, A., & Ghosh, H. (2026). A comprehensive review of Edge Computing-empowered smart agriculture: Trends, Opportunities and Future directions. *Computers and Electronics in Agriculture*. <https://doi.org/10.1016/j.compag.2025.111252>
- [12] Yu, P., Teng, F., Zhu, W., Shen, C., Chen, Z., & Song, J. (2025). Cloud-edge-device collaborative computing in smart agriculture: architectures, applications, and future perspectives. *Frontiers in Plant Science*, 16. <https://doi.org/10.3389/fpls.2025.1668545>
- [13] Gong, R., Zhang, H., Li, G., & He, J. (2025). Edge Computing-Enabled Smart Agriculture: Technical Architectures, Practical Evolution, and Bottleneck Breakthroughs. *Sensors*, 25(17), 5302. <https://doi.org/10.3390/s25175302>
- [14] Pal, D. (2023). AI-Driven Edge Computing for IoT: A Comprehensive Survey and Future Directions. *International Journal of Advanced Research in Science, Communication and Technology*. <https://doi.org/10.48175/IJARSCT-12921>
- [15] Sapna, Gauttam, H., Chauhan, V., Pattanaik, K. K., Trivedi, A., & Ghosh, H. (2026). A comprehensive review of Edge Computing empowered smart agriculture: Trends, Opportunities and Future directions. *Computers and Electronics in Agriculture*. <https://doi.org/10.1016/j.compag.2025.111252>
- [16] Diop, F., Faye, B. M., & Niang, I. (2025). Edge-AI and Internet of Things for Intelligent Systems: Architectures, Applications and Future Perspectives. In *AFRICOMM 2023*, Springer. https://doi.org/10.1007/978-3-031-81573-7_9