Computers and Electronics in Agriculture 245 (2026) 111559

Contents lists available at ScienceDirect

Computers and Electronics in Agriculture

journal homepage: www.elsevier.com/locate/compag

Deep learning-based UAV framework for automated morphological and growth analysis of feedlot cattle

Jianglong Yan [a], Everton C. Tetila [b], Liang Zhao [c,*], Rian C. Gonçalves [b], Letícia F. Castanheiro [d], Lucas P. Valem [a], Jayme G.A. Barbedo [d]

[a] Institute of Mathematics and Computer Science (ICMC), University of São Paulo (USP), Av. Trabalhador São-carlense, 400, São Carlos, 13566-590, Brazil
[b] Federal University of Grande Dourados, Rodovia Dourados/Itahum, Km 12, Cidade Universitária, Dourados, 79804-970, Brazil
[c] Faculty of Philosophy, Sciences and Letters at Ribeirão Preto (FFCLRP), University of São Paulo (USP), Av. Bandeirantes, 3900, Ribeirão Preto, 14040-901, Brazil
[d] Embrapa Digital Agriculture, Av. Dr. André Tosello, 209, Cidade Universitária, Campinas, 13083-886, Brazil

ARTICLE INFO

Keywords:
UAV imagery
Cattle growth monitoring and modeling
Morphological feature extraction
Automatic cattle selection
Precision livestock farming

ABSTRACT

Accurate and continuous monitoring of cattle development is essential to optimize feedlot management and determine the most profitable slaughter period. Traditional weighing procedures are labor-intensive, intrusive, and difficult to scale, highlighting the need for automated, non-invasive longitudinal monitoring frameworks. In this study, we propose a UAV-based computer vision framework that integrates high-resolution aerial imagery with deep learning techniques to track the morphological development of beef cattle throughout the production cycle, rather than directly estimating body weight. A subset of images was annotated to fine-tune YOLOv11s for cattle detection, while instance-level body contours were extracted using SAM 2.1 without additional training. An automated filtering strategy retained only well-posed standing animals, enabling consistent extraction of morphological features such as body length, width, and Width-to-Length ratio (W/L), a scale-invariant proxy of body condition over time. Longitudinal, population-level analysis of Nellore cattle monitored over 112 days revealed a clear sigmoidal growth pattern, well described by a logistic model, and the growth curve uncovered three biological phases: lag, rapid development, and plateau, allowing the framework to qualitatively determine an appropriate economic marketing window. These results demonstrate the practical value of the proposed framework for scalable, data-driven precision livestock management in commercial feedlot environment.

1. Introduction

Beef cattle farming is a key sector of the Brazilian agricultural economy, supporting both domestic consumption and international beef exports (Lathuillière et al., 2025; Costa Jr et al., 2025). In feedlot systems, determining the suitable selling or slaughter time is crucial for maximizing profitability. Traditionally, this decision relies on periodic weighing, a process that requires physical restraint, is labor-intensive, and can induce stress in the animals, potentially affecting welfare and growth performance. These limitations have motivated increasing interest in automated, non-invasive, and high-throughput methods capable of monitoring cattle morphology and growth without the need for physical handling (Aquilani et al., 2022; Besler et al., 2024; Curti et al., 2023).

Recent advancements in computer vision and deep learning have enabled image-based livestock monitoring. Early studies using 2D images captured by fixed cameras demonstrated the feasibility of estimating cattle weight under controlled conditions. For example, Weber et al. (2020) obtained dorsal images of Nelore and Girolando cattle and performed manual segmentation to extract morphological features. However, these approaches are constrained by the need for manual annotation and their sensitivity to environmental factors.

Regarding UAV-based monitoring, Los et al. (2023) estimated weight by analyzing 3D models made with RGB imagery and video from UAV, and estimated height and weight from LiDAR data. Similarly, Wang et al. (2025) used UAV-based LiDAR to evaluate weight across different postures. However, the use of LiDAR data from one-way scanning may result in incomplete point cloud for body parts farther from the flight path reducing the accuracy of volume-based weight estimation for certain postures (e.g., cattle lying and heads up). In addition, LiDAR data require more complex and computationally intensive data processing.

* Corresponding author.
E-mail address: zhao@usp.br (L. Zhao).
URL: https://cemeai.icmc.usp.br/zhao-liang/ (L. Zhao).

https://doi.org/10.1016/j.compag.2026.111559
Received 2 December 2025; Received in revised form 6 February 2026; Accepted 9 February 2026
Available online 12 February 2026

Mobile device-based approaches have also been explored. Bai et al. (2025) proposed the LaWE model, a lightweight framework for smartphone-acquired images that integrates multiscale keypoint detection with deep regression to estimate body measurements, achieving over 97% accuracy on Horqin Yellow cattle under controlled conditions. However, its generalization to commercial feedlots remains limited due to variability in lighting, posture, and crowding. In commercial environments, Giannone et al. (2025) developed a YOLOv8-based system for individual dairy cow identification and feeding behavior analysis, achieving 85% precision and 62% recall (F1 score 0.72) at IoU 0.5, demonstrating the feasibility of real-time deployment. Afridi et al. (2024) further showed that integrating multimodal data (RGB, depth, and segmentation) significantly improves the robustness of cattle weight estimation using vision-transformer-based segmentation models. Other approaches include the method of Xu et al. (2024), which combined semantic segmentation with ResNet-101-D and Squeeze-and-Excitation modules, and the PointNet++-based 3D modeling approach of Hou et al. (2023) using LiDAR point clouds, which achieved 95.1% segmentation accuracy but required specialized sensors and extensive preprocessing.

A positive correlation between cattle body morphology and body weight has been widely reported in the literature (Firdaus et al., 2023; Xiong et al., 2023). However, existing studies predominantly rely on close-range images of individual animals acquired under controlled or semi-controlled conditions. In contrast, our approach operates on large-scale UAV imagery covering extensive areas, where multiple animals are observed simultaneously and scene complexity is substantially higher.

Despite these advances, population-level growth trajectory modeling based on longitudinal morphological data remains largely unexplored in commercial feedlot systems. Prior work typically focuses on small numbers of animals and emphasizes individual-level measurement or weight estimation. By shifting the focus to population-level dynamics, our framework enables characterization of collective growth patterns and identification of economically relevant marketing windows without requiring individual animal identification or tracking.

In this study, we propose a UAV-based cattle monitoring framework deployed in a commercial feedlot in Laguna Carapã, Mato Grosso do Sul, Brazil. Twelve UAV missions were conducted to acquire high-resolution top-view images approximately every ten days throughout the finishing period of Nelore cattle.

Specifically, the objectives of this work are threefold: The technical objective of this study is to develop a two-stage deep learning pipeline that integrates YOLOv11s (Jocher and Qiu, 2024) as a primary cattle detector trained on a limited manually annotated subset with SAM 2.1 (Ravi et al., 2024) for instance segmentation, enabling large-scale automatic extraction of morphological measurements, such as body length and width, from UAV imagery without exhaustive manual labeling. Building on these outputs, the analytical objective is to model population-level growth dynamics by fitting logistic growth functions to aggregated morphological indices, thereby characterizing distinct phases of morphological development during feedlot finishing. Finally, the decision-support objective is to leverage the inferred growth dynamics to identify economic saturation points and potential marketing windows, providing qualitative support for determining appropriate slaughter periods without relying on direct body-weight prediction.

Because the UAV missions covered large feedlot areas without reliable mechanisms for re-identifying the same individuals across different acquisition dates, precise longitudinal tracking at the single-animal level was not feasible. Therefore, the proposed growth modeling is performed at the population level, where temporal trends are inferred from aggregated morphological measurements of filtered cattle instances automatically detected by the fine-tuned YOLOv11s model at each time point.

The major contributions of this work are as follows:

- Creation of the *NelloreBeefCattleDataset*, comprising 904 high-quality UAV images, of which 370 were selected and annotated, to support model training and performance evaluation.
- A fully automatic and high-throughput UAV-based monitoring pipeline in which a fine-tuned YOLOv11s detector enables large-scale, non-invasive instance discovery and SAM 2.1 performs precise morphometric segmentation, allowing population-level longitudinal analysis of cattle growth under real commercial feedlot conditions.
- Technical innovations: a semi-supervised annotation and inference strategy where YOLOv11s trained on a limited labeled subset is used to automatically localize cattle in the remaining unlabeled images, followed by SAM 2.1 segmentation and posture-based instance filtering for reliable morphometric extraction without human intervention.
- More importantly, a growth modeling using logistic function to uncover cattle developmental phases and appropriate marketing periods.

This integrated framework provides a scalable, practical, and non-invasive solution for precision livestock monitoring, reduces labor requirements, and supports data-driven decision-making to enhance feedlot profitability.

2. Methodology

This section presents the UAV-based computer vision pipeline used for cattle monitoring, covering detection, segmentation, filtering, morphological measurement, and growth curve modeling. Each subsection summarizes the corresponding methodological component.

2.1. Proposed pipeline

Monitoring individual cattle growth trajectories in large-scale feedlot operations presents substantial technical challenges: animals are visually similar, densely grouped, and exhibit dynamic postures that complicate accurate biometric measurement. Conventional approaches relying on periodic manual weighing are labor-intensive, stressful to animals, and lack the high temporal resolution required for precision management decisions. To overcome these limitations, we developed an integrated computer vision pipeline that combines Unmanned Aerial Vehicle (UAV) imagery, deep learning-based object detection and segmentation, and temporal growth modeling, enabling non-invasive, scalable, and frequent monitoring of individual cattle morphology.

The proposed system consists of seven stages shown by Fig. 1:

1. UAV Image Acquisition: Capture aerial views in high-resolution images at regular intervals throughout the monitoring period.
2. Dataset Preparation and Manual Annotation: A subset of the collected images was manually annotated and used to fine-tune YOLOv11s for cattle detection.
3. YOLO Fine-tuning for Cattle Detection: A pre-trained YOLOv11s model was fine-tuned using the annotated dataset to detect individual cattle in complex feedlot scenes. This fine-tuning yields a transferable detection backbone that can be efficiently adapted to similar feedlot environments in future deployments.
4. Instance Segmentation using SAM Model: Detect the precise pixel-level body masks using SAM 2.1 for each cattle bounding box.
5. Morphological Feature Extraction: Compute segmentation-based morphological indices (e.g., length, width, and Width-to-Length ratio (W/L)) and remove noise or artifacts.
6. Temporal Analysis and Growth Curve Fitting: Fit individual growth trajectories using time-series morphological measurements.

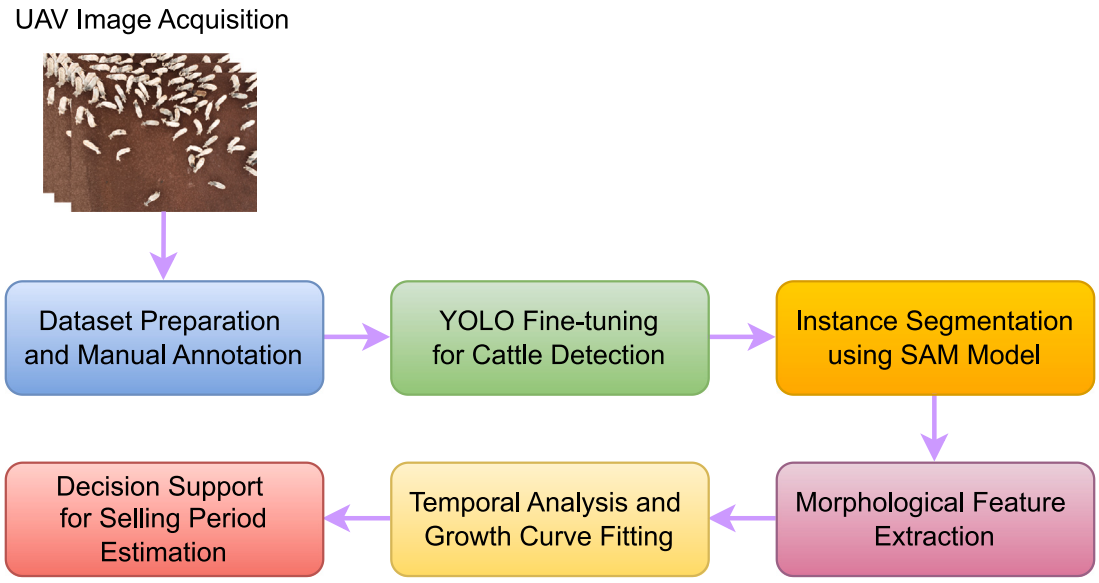

Fig. 1. Overview of the proposed cattle growth monitoring pipeline.

7. Decision Support for Selling Period Estimation: Estimate the appropriate marketing window based on growth-curve parameters and developmental saturation points.

This modular pipeline offers several methodological advantages and enables automated, high-throughput morphometric extraction in large-scale commercial feedlot environments. First, the two-stage detection-then-segmentation approach (YOLOv11s followed by SAM 2.1) decouples localization from fine-grained boundary delineation, enabling efficient processing of high-resolution imagery while maintaining segmentation accuracy. Second, leveraging foundation models (YOLOv11s pre-trained on ImageNet and SAM 2.1 applied zero-shot) reduces the demand for extensive domain-specific training data, facilitating rapid deployment and transferability across feedlot settings. Third, extracting morphological features from segmentation masks provides a non-invasive proxy for growth status, eliminating the need for physical weighing infrastructure. Finally, the temporal growth modeling framework converts discrete morphological measurements into continuous growth trajectories, enabling predictive analytics and proactive management interventions.

2.2. Data collection

A field study was carried out at the Fazenda Campanário feedlot in Laguna Carapã, Mato Grosso do Sul, Brazil (22° 47′8″S; 55° 3′57″W), which houses approximately 14,600 cattle. Data acquisition was performed using a multirotor UAV (DJI Phantom 4 Advanced) equipped with a 20-megapixel Sony camera (4864 × 3648 pixels). All flights were conducted at a fixed altitude of 15 m to maintain consistent spatial resolution. Twelve flight operations were carried out at approximately 10-day intervals. Image acquisition was performed using a 90 degree camera angle, with a 50% overlap between consecutive images. Data collection was conducted in the morning, between 7:00 and 11:00 a.m., under varying weather conditions, including cloudy, sunny and rainy scenarios. The acquisition platform operated at a speed of 4 m/s. Monitoring was performed over a 112-day period, starting on July 10, 2024 (Fig. 2). The average age of the cattle at the beginning of the feedlot period was 18 months. This study exclusively uses still images.

In these experiments, 904 high-quality images were collected. From this pool, 370 images were randomly selected and manually annotated with bounding boxes covering the full body extent of each visible animal. These annotated samples were used to fine-tune the YOLOv11s detector, whereas the remaining unannotated images were reserved for automatic inference using the fine-tuned model, thereby substantially reducing the manual labeling workload.

Table 1
Dataset Statistics.

	Images	Instances	Standing	Lying
Training	273	2113	1735	378
validation	48	430	335	95
Test	49	420	335	95
Total	370	2963	2385	576

Manual bounding box annotations were performed on these images. Each visible animal in the training set was labeled with a bounding box covering the entire body extent. The annotated dataset was then randomly split into training (80%), validation (10%), and test (10%) subsets using stratified sampling to preserve representative variability in individual animals and imaging conditions. The training and validation subsets were used during the training process for model optimization and hyperparameter tuning, whereas the test subset was held out throughout training and employed exclusively for the final performance evaluation.

Detailed statistics of each subset, including the number of images, total cattle instances, and posture class distribution (standing and lying), are summarized in Table 1.

2.3. Object detection

We selected YOLOv11s (Jocher and Qiu, 2024) as the backbone detector for cattle detection, initialized with COCO pre-trained weights (Lin et al., 2014). The choice is driven by the need to balance accuracy, computational efficiency, and real-time capability in large-scale ranch monitoring scenarios.

Table 2 summarizes representative object detection models commonly adopted in the recent literature. Importantly, the table reports *standard COCO benchmark results from their original publications* rather than results from our dataset. The comparison helps to clarify the design considerations that motivated our choice of YOLOv11s. YOLOv11s incorporates several architectural refinements, including improved C3k2 feature extractors and C2PSA attention modules that strengthen spatial sensitivity and computational efficiency (Jocher and Qiu, 2024). These improvements enable YOLOv11s to achieve

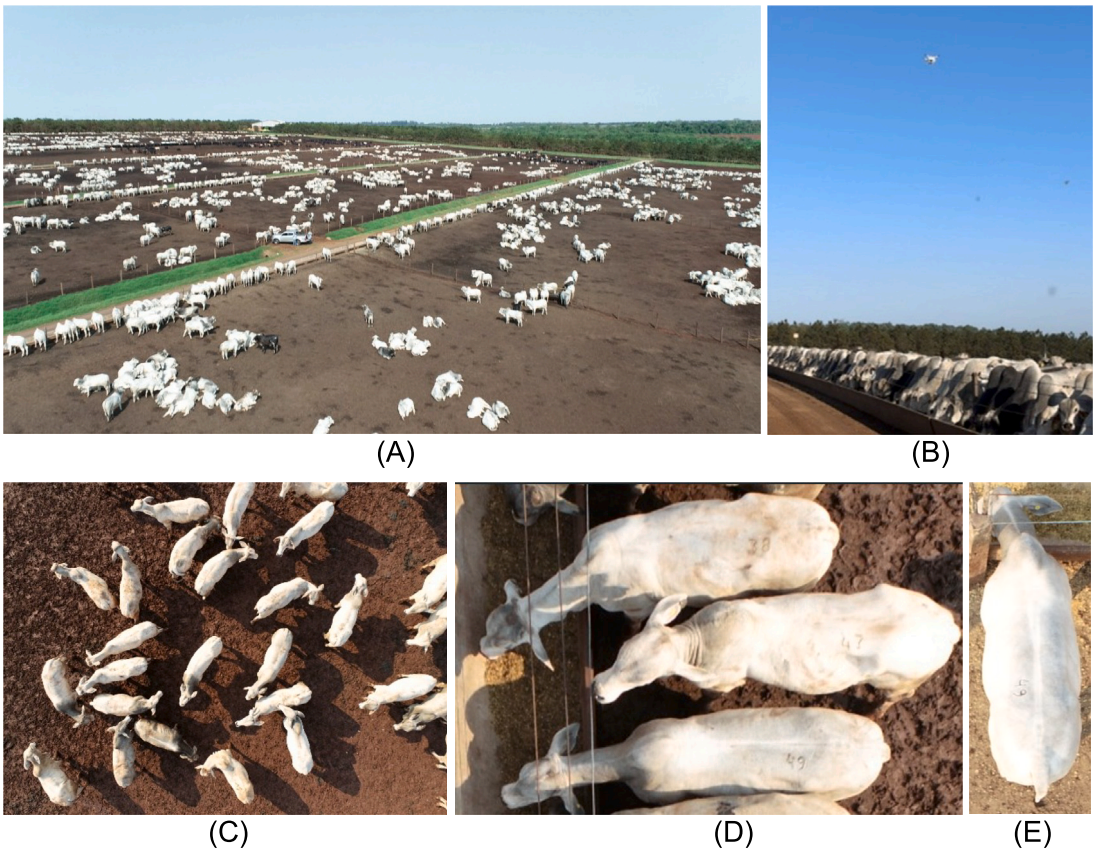

Fig. 2. (A) Aerial view of the feedlot at Fazenda Campanário; (B) UAV operation at approximately 15 m altitude; (C) representative top-view images acquired during the flights; (D) and (E) enlarged images.

Table 2
Background benchmark of representative object detection architectures on the COCO validation set, reproduced from the respective original publications. All metrics are reported from the literature and are not experimental results of this study. FPS values correspond to TensorRT-optimized inference where available.

Model	Params (M)	FLOPs (G)	**mAP** (%)	FPS	NMS-Free
Faster R-CNN (Ren et al., 2015)	41.5	–	42.0	~15	No
DETR-R50 (Carion et al., 2020)	41.3	–	42.0	~28	Yes
RT-DETR-R50 (Zhao et al., 2024)	32.0	–	53.1	108	Yes
YOLOv8s (Jocher et al., 2023)	11.2	28.6	44.7	120	No
YOLOv10s (Wang et al., 2024)	8.0	24.5	46.3	122	Yes
YOLOv11s	**9.4**	**21.5**	**46.7**	**145**	No

competitive accuracy while remaining lightweight enough for scalable deployment.

Our choice of YOLOv11s as the backbone detector is motivated by three key observations. First, two-stage detectors such as Faster R-CNN (Ren et al., 2015), despite reasonable accuracy, offer limited throughput (~15 FPS), which is unsuitable for continuous UAV-based monitoring. Second, transformer-based models like RT-DETR (Zhao et al., 2024) eliminate NMS and deliver strong accuracy but generally exhibit slower convergence and weaker performance on small, distant objects—an important factor in aerial livestock imagery. Third, YOLOv11s provides measurable improvements over YOLOv8s, achieving higher mAP (+2.0%), reduced parameters (–16%), lower FLOPs (–25%), and faster inference (+20%) (Jocher and Qiu, 2024). These gains result from more efficient feature aggregation modules and an optimized neck design.

2.4. Instance segmentation

The Segment Anything Model (SAM) is a foundation model designed for promptable image segmentation, capable of generating high-quality object masks from various input prompts (points, boxes, or text) (Kirillov et al., 2023). Unlike traditional segmentation models that require task-specific training, SAM's zero-shot learning paradigm enables direct deployment across diverse domains without fine-tuning. SAM 2.1 (Ravi et al., 2024), the latest iteration, extends these capabilities with enhanced architectural components, including streaming memory attention for video consistency and improved mask decoder efficiency, achieving state-of-the-art performance on zero-shot segmentation benchmarks.

Instance segmentation of cattle in real-world farm environments faces a critical data bottleneck: the scarcity of pixel-level annotated masks. While bounding box annotations for object detection are relatively straightforward to acquire, creating precise segmentation masks requires labor-intensive manual delineation of irregular cattle body contours—especially challenging given variations in posture, occlusion, and lighting conditions across thousands of frames. This annotation burden renders traditional supervised segmentation approaches (e.g., Mask R-CNN (He et al., 2017) typically requiring thousands of manually annotated masks) impractical for livestock monitoring applications.

SAM 2.1's zero-shot capability directly addresses this limitation by eliminating the need for domain-specific training data. The model's pre-training on the SA-V dataset (50,800 videos, 642,600 masklets covering diverse real-world scenarios) (Ravi et al., 2024) enables robust generalization to agricultural contexts without cattle-specific annotations. Furthermore, SAM 2.1's box-prompting interface seamlessly integrates with our YOLOv11s detection pipeline, converting detected bounding boxes into precise segmentation masks through a unified two-stage architecture.

Key advantages of SAM 2.1 for cattle segmentation include the zero-shot generalization, where direct deployment without cattle-specific mask annotations addresses the data scarcity challenge; box-prompt integration, where Native support for bounding box inputs enables seamless coupling with YOLOv11s detections; superior accuracy, where achieves a 76.4% J&F score on the SA-V benchmark, outperforming SAM 1.0 by 8.2 percentage points (Ravi et al., 2024); as well as computational efficiency, where SAM 2.1 (Tiny) delivers 47.2 FPS—nearly 4× faster than SAM 1.0 while maintaining comparable accuracy (75.0% vs. 68.2% J&F) (Ravi et al., 2024).

Following object detection, detections classified as *lying* are immediately discarded to eliminate unsuitable postures. Only detections classified as *standing* with a confidence score $p_i \geq \theta_{conf}$ are forwarded to the segmentation pipeline. In this work, the Detection Confidence Threshold was set to $\theta_{conf} = 0.7$, and a sensitivity study of this parameter is presented in Section 3.3.3. The bounding boxes generated by YOLOv11s for *standing* cattle (format: $[x_1, y_1, x_2, y_2]$ in pixel coordinates) served as spatial prompts to guide SAM 2.1's attention mechanism toward individual cattle regions. For each detection i satisfying class(d_i) = standing and $p_i \geq \theta_{conf}$, the corresponding bounding box coordinates were directly input to SAM 2.1's box-prompt interface:

$$\text{Prompt}_i = \{(x_{1,i}, y_{1,i}), (x_{2,i}, y_{2,i})\} \tag{1}$$

where $(x_{1,i}, y_{1,i})$ and $(x_{2,i}, y_{2,i})$ represent the top-left and bottom-right corners of the ith detected *standing* cattle bounding box. This box-conditioned prompting strategy significantly improves segmentation precision by eliminating ambiguous background regions and focusing computational resources on cattle body boundaries. The model outputs binary segmentation masks $M_i \in \{0,1\}^{H \times W}$ delineating pixel-level cattle contours for all standing cattle instances. We deployed the SAM 2.1 (Large) configuration to prioritize segmentation quality over inference speed, as our pipeline processes pre-recorded videos rather than real-time streams.

2.5. Hierarchical instance filtering strategy

A hierarchical instance filtering strategy was developed to remove cattle detections unsuitable for morphological analysis. The pipeline integrates posture evaluation, spatial context constraints, and shape-based quality control into a single coherent process.

2.5.1. YOLO detection confidence filtering

Specifically, YOLOv11s is first fine-tuned on a manually annotated subset, and then applied to automatically infer cattle locations in the remaining UAV images, producing bounding boxes and standing/lying posture predictions for all visible individuals.

For each detected instance i, the model outputs a bounding box b_i, a posture label $y_i \in \{standing, lying\}$, and an associated confidence score $p_i \in [0,1]$, representing the posterior probability estimated by the network. Only detections satisfying

$$p_i \geq \theta_{conf}$$

are retained for subsequent processing, where θ_{conf} is a predefined Detection Confidence Threshold parameter. Detections with $p_i < \theta_{conf}$ are discarded to suppress false positives and unreliable posture classifications.

2.5.2. ROI size filtering

The second stage filters cattle using a spatial Region Of Interest (ROI) to mitigate radial distortion and oblique perspective deformation introduced by UAV wide-angle lenses. All spatial operations are performed in the normalized image coordinate system, where both bounding-box coordinates and ROI boundaries are expressed in the range $[0, 1]$ relative to image width and height.

This stage consists of two sequential steps: (i) defining a central ROI in normalized coordinates and (ii) retaining only those cattle instances whose bounding-box centroids fall inside the ROI.

Let the ROI be denoted as

$$\text{ROI}_{norm} = (x_{min}, y_{min}, x_{max}, y_{max}), \tag{2}$$

where all parameters are normalized with respect to the image width and height. A sensitivity study on the parameter ROI has been presented in Section 3.3.1.

For each detected cattle instance i with bounding box coordinates $(x_{min}^{b_i}, y_{min}^{b_i}, x_{max}^{b_i}, y_{max}^{b_i})$ in normalized form, the bounding-box centroid is computed as

$$c_x = \frac{x_{min}^{b_i} + x_{max}^{b_i}}{2}, \quad c_y = \frac{y_{min}^{b_i} + y_{max}^{b_i}}{2}. \tag{3}$$

An instance is retained if and only if its centroid lies within the ROI, i.e.,

$$x_{min} \leq c_x \leq x_{max} \ \wedge \ y_{min} \leq c_y \leq y_{max}. \tag{4}$$

In this work, the ROI Size parameter is denoted as S_{ROI}. Fig. 3 illustrates an example where the ROI was defined as $(x_{min}, y_{min}, x_{max}, y_{max}) = (0.15, 0.15, 0.85, 0.85)$, corresponding to the $S_{ROI} = 70\%$ nadir region.

By operating entirely in normalized coordinates, this spatial filtering strategy is invariant to absolute image resolution and UAV flight altitude, ensuring consistent filtering behavior across all acquisition sessions.

2.5.3. W/L range filtering

The final and main stage applies W/L ratio filtering to exclude cattle instances affected by various geometric distortions. Minimum-area rotated rectangles were fitted to each segmentation contour to estimate body dimensions. From these rectangles, cattle width and length were obtained. The Width-to-Length ratio (W/L) was selected as the primary indicator for body-condition assessment due to its robustness to UAV altitude and scale variations:

$$\text{W/L} = \frac{\text{Width}}{\text{Length}} \tag{5}$$

The W/L ratio computed from each segmentation mask is used as a geometric proxy to assess pose suitability:

$$(\text{W/L})_{\min} < (\text{W/L}) \leq (\text{W/L})_{\max}, \tag{6}$$

where $(\text{W/L})_{\min}$ and $(\text{W/L})_{\max}$ are empirically derived thresholds. W/L values below $(\text{W/L})_{\min}$ typically result from excessive leg extension or stretched standing postures, which artificially increase the apparent body length and produce abnormally small W/L ratios, whereas values above $(\text{W/L})_{\max}$ are generally caused by geometric distortion, occlusion effects, head turning, instance merging, inaccurate bounding-box localization, or occasional misclassified *lying* individuals. This filtering step aims to remove standing cattle whose apparent body proportions are severely distorted and therefore unsuitable for reliable morphometric extraction. In this work, the W/L Range parameter is denoted as $R_{W/L}$.

Through this three-stage filtering pipeline (posture, spatial ROI, and W/L ratio), only high-quality standing cattle instances with minimal geometric distortion are retained for subsequent morphometric extraction.

Fig. 3. Example image illustrating ROI and cattle centroid filtering. The *red dashed box* denotes the ROI boundary ($S_{ROI} = 70\%$). *Green boxes* indicate detected standing cattle whose centroids lie within the ROI. *Orange circles with white crosses* mark bounding-box centroids. Cattle without green bounding boxes were detected by YOLOv11s but excluded by the ROI and centroid filtering criteria.

2.6. Feature extraction and growth curve modeling

For each segmentation mask M_i that passed the hierarchical filtering strategy, morphological features were extracted to characterize cattle body geometry. The largest connected component was selected using 8-connectivity, and a standard morphological closing operation was applied to fill small holes in the mask.

For all cattle retained after ROI and centroid filtering, the width-to-length (W/L) ratios are computed at all 12 time points (approximately 10-day intervals). These ratios serve as the final and primary filtering criterion for selecting well-posed cattle. At the same time, they provide a stable measure of body compactness: lower values indicate more elongated, lean individuals, whereas higher values correspond to broader bodies typically associated with better physiological condition.

Finally, at each time point, the mean W/L ratio is computed over all cattle of that time point satisfying the condition established in Eq. (6), representing the overall herd growth trajectory at that time point.

The temporal evolution of the mean W/L ratio was modeled using a sigmoidal (S-shaped) logistic function:

$$y(t) = \frac{L}{1 + e^{-k(t-t_0)}} + b \tag{7}$$

In this logistic formulation, $y(t)$ denotes the mean W/L ratio at time t. The parameter L defines the upper asymptote and corresponds to the maximum attainable increase in mean W/L ratio, whereas b represents the lower asymptote associated with the baseline ratio in early growth. The coefficient k (day^{-1}) controls the rate of morphological change, and t_0 specifies the inflection point, indicating the time at which growth transitions from acceleration to deceleration. The variable t denotes the elapsed number of days from the start of monitoring.

The logistic function was selected because cattle growth patterns follow S-shaped trajectories, with an initial lag phase, a rapid linear growth phase, and a terminal plateau phase as the animals near market weight or mature size. The parameters carry biological interpretability: t_0 approximates the timing of maximum growth velocity (often corresponding to the critical developmental stage in feedlot cattle), k quantifies population-level growth rate, and L and b bracket the physiological range of morphological indices.

Non-linear least squares optimization (Levenberg–Marquardt algorithm) was employed to fit Eq. (7) to the herd-averaged growth trajectory. Initial parameters were estimated from the observed mean W/L ratios and their temporal derivatives. The optimization procedure minimized the residual sum of squares:

$$\chi^2 = \sum_{i=1}^{12} (y_{\text{observed},i} - y_{\text{predicted},i})^2 \tag{8}$$

where $y_{\text{observed},i}$ represents the mean W/L ratio at time point i, and $y_{\text{predicted},i}$ is the corresponding value predicted by the fitted logistic model.

Goodness-of-fit was assessed using the Coefficient of Determination (R^2) and RMSE, ensuring the reliability of the fitted growth model for the herd-level trajectory.

Based on the fitted growth curve, a suitable selling period can be estimated. It is the time interval when the herd-averaged growth rate (dy/dt) falls below a species-specific threshold after the increasing stage, coupled with the achievement of a target body condition score estimated from the mean W/L ratio. Mathematically:

$$\left.\frac{dy}{dt}\right|_{t=t^*} = k \cdot \frac{L \cdot e^{-k(t^*-t_0)}}{(1 + e^{-k(t^*-t_0)})^2} < \theta_{\text{threshold}} \tag{9}$$

where t^* is the recommended selling time, and $\theta_{\text{threshold}}$ is a species/market-specific threshold (determined from the literature or farm management targets).

2.7. Experimental setup

All training experiments were conducted on an NVIDIA RTX Ada 5000 GPU (32 GB VRAM) using the Ultralytics YOLOv11s framework (Jocher and Qiu, 2024) and PyTorch 2.5.1.

The complete dataset used in this study comprises 904 UAV images acquired on multiple observation dates over a 112-day fattening period. From this dataset, 370 images were randomly selected and manually annotated for YOLO fine-tuning in the cattle detection task. All images were processed through the full pipeline, including cattle detection, posture filtering, instance segmentation, and morphological measurement.

For YOLOv11s, fine-tuning was performed over 200 epochs with a batch size of 16 and an input resolution of 832 × 832. The learning rate was initialized at 3×10^{-4} and scheduled using cosine decay, with optimization performed using AdamW (weight decay = 1×10^{-3}). These configurations were used for all YOLOv11s experiments unless stated otherwise.

In addition to model training, the inference efficiency of the complete processing pipeline was evaluated to assess its practical feasibility. All inference experiments were performed using single-image processing without batch acceleration. The end-to-end pipeline, including YOLOv11-based cattle detection followed by SAM 2.1 instance segmentation, required an average inference time of 0.953 s per image, corresponding to a throughput of approximately 1.05 images per second.

These measurements reflect pure model inference time and exclude UAV flight, image acquisition, and disk I/O overhead. Given the temporal resolution of UAV-based monitoring in commercial feedlot scenarios, this processing speed is sufficient to support large-scale offline analysis and near-real-time decision support.

3. Experimental results

This section reports the experimental evaluation of the proposed UAV-based cattle growth monitoring framework. It includes quantitative assessments of cattle detection and instance segmentation performance, followed by population-level growth modeling, growth dynamics characterization, and estimation of an economically appropriate slaughter endpoint.

3.1. Cattle detection evaluation

The YOLOv11s detector was trained for 200 epochs. The model checkpoint selected at epoch 188 achieved the best overall performance on the validation set and was subsequently evaluated on the held-out test set. On the test data, the detector achieved a precision of 95.56% and a recall of 98.38% at an IoU threshold of 0.5. The corresponding mAP@0.5 reached 98.18%, while mAP@0.75 was 97.82%, and the aggregated mAP@0.5:0.95 was 94.85%. In addition, the mean inference time per image was 0.378 s, corresponding to an average processing speed of 2.65 FPS, demonstrating reliable detection accuracy and practical inference efficiency under real commercial feedlot conditions.

Training and validation loss curves, together with the evolution of detection metrics across epochs, are provided in the Supplementary Materials (Figs. S1–S3).

3.2. Instance segmentation results

Because no manually annotated ground-truth masks were available, segmentation performance was assessed qualitatively. The evaluation focused on boundary accuracy, silhouette completeness, and robustness under varying illumination, shadows, and inter-animal occlusion.

Without any fine-tuning, SAM 2.1 produced high-quality instance segmentation masks that closely adhered to individual cattle contours. Representative examples are shown in Fig. 4. These results demonstrate the model's capability to accurately delineate individual cattle, even under dense feedlot conditions and partial occlusions.

3.3. Growth dynamics modeling and evaluation

This subsection presents cattle growth modeling and the discovery of the developmental phases.

3.3.1. Herd-level vs individual-level inference

Although the proposed pipeline enables instance-level detection and segmentation of individual cattle, reliable longitudinal tracking of the same animal across repeated UAV missions is not feasible under commercial feedlot conditions due to occlusions, high animal density, and the lack of individual visual identifiers. Therefore, the subsequent growth analysis is conducted at the herd (population) level rather than at the individual-animal level.

In this study, instance-level morphological measurements are aggregated to characterize population-level growth dynamics, which provides a robust and scalable representation of developmental trends while remaining consistent with the practical constraints of aerial monitoring in large commercial feedlots.

To mitigate dependencies arising from the high spatial overlap between consecutive UAV images, only detections located within the central region of each image were used for statistical analysis. Based on a sensitivity study, we determined that a central window covering approximately 70% of the original image size effectively reduces the probability of repeatedly sampling the same animal across overlapping frames acquired on the same date.

3.3.2. Morphological feature extraction

To ensure that only geometrically reliable cattle instances are used for morphometric analysis, we first examine typical failure cases produced during automatic detection and segmentation.

Fig. 5 illustrates representative failure cases with abnormally large W/L ratios. The examples correspond to different error sources, including mild wide-angle camera distortion, occlusion by neighboring cattle, excessive head turning, merged instance segmentation, inaccurate bounding-box localization, and occasional misclassified *lying* individuals missed by the posture classifier. These cases visually demonstrate the necessity of the proposed W/L-based filtering strategy for excluding distorted or unreliable cattle instances.

Complementary to these high-ratio failures, Fig. 6 presents typical examples of rejected instances with W/L ratios below the lower threshold. In these cases, excessive leg extension or stretched postures artificially increase the apparent body length, leading to abnormally small W/L values despite the animals being in an upright posture. Such geometric distortions cannot be reliably corrected at the segmentation stage and would introduce significant bias into morphometric measurements if retained, thereby justifying their exclusion.

After excluding these two categories of unreliable instances, only cattle segmentations satisfying all posture, ROI, and W/L filtering criteria are retained for morphometric analysis. Representative cattle instances that satisfy all filtering stages are illustrated in Fig. 7, showing the segmentation produced by SAM 2.1.

Applying this complete filtering pipeline across all UAV acquisitions resulted in a refined dataset of geometrically consistent cattle instances suitable for population-level growth modeling.

Across the 12 time points, we conducted a series of controlled experiments to evaluate the sensitivity of valid instance extraction to the parameters of Detection Confidence Threshold (θ_{conf}), ROI Size (S_{ROI}), and W/L Range ($R_{W/L}$). Tables 3–5 report the number of selected cattle instances before filtering and after varying each of these three parameters, respectively. In Tables 3–5, the bolded rows indicate the parameter settings that were empirically identified as providing the best trade-off between instance quantity and quality, and were therefore adopted as the default configuration in the subsequent analysis ($\theta_{conf} = 0.7$, $S_{ROI} = 70\%$, and $R_{W/L} \in [0.250, 0.333]$). Instead of relying on a single configuration as in our preliminary analysis, these three factors were systematically varied, and the resulting number of retained

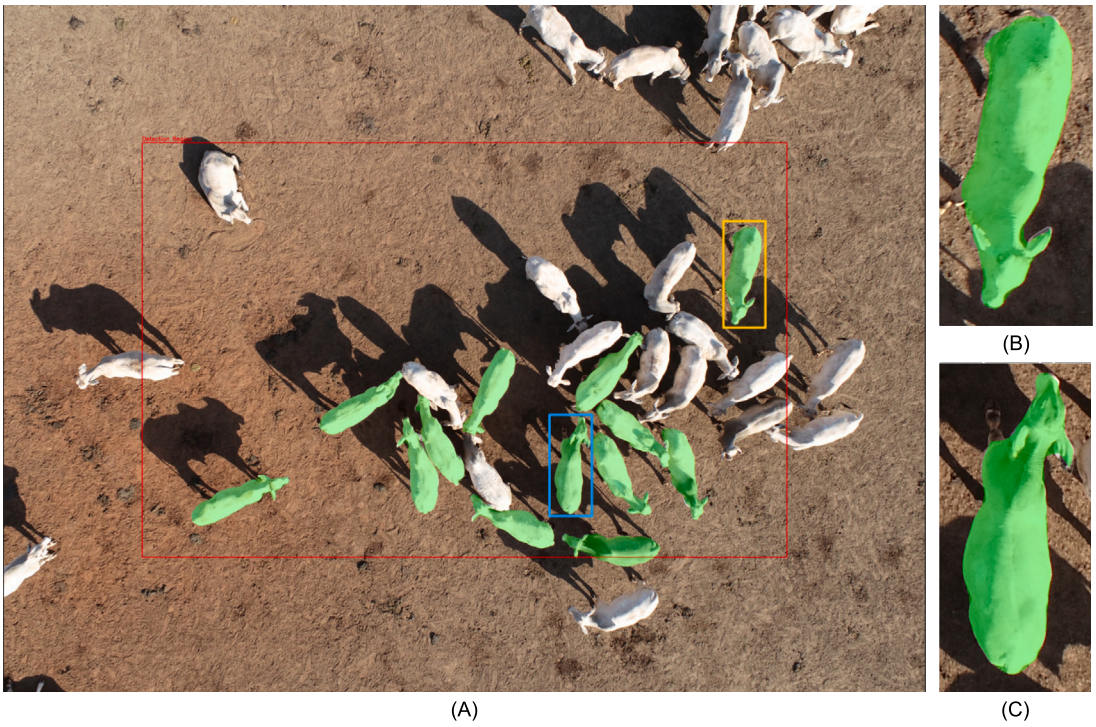

Fig. 4. Instance segmentation results using the SAM 2.1 model. (A) Original UAV image. (B) Segmentation result corresponding to the yellow box in (A). (C) Segmentation result corresponding to the blue box in (A).

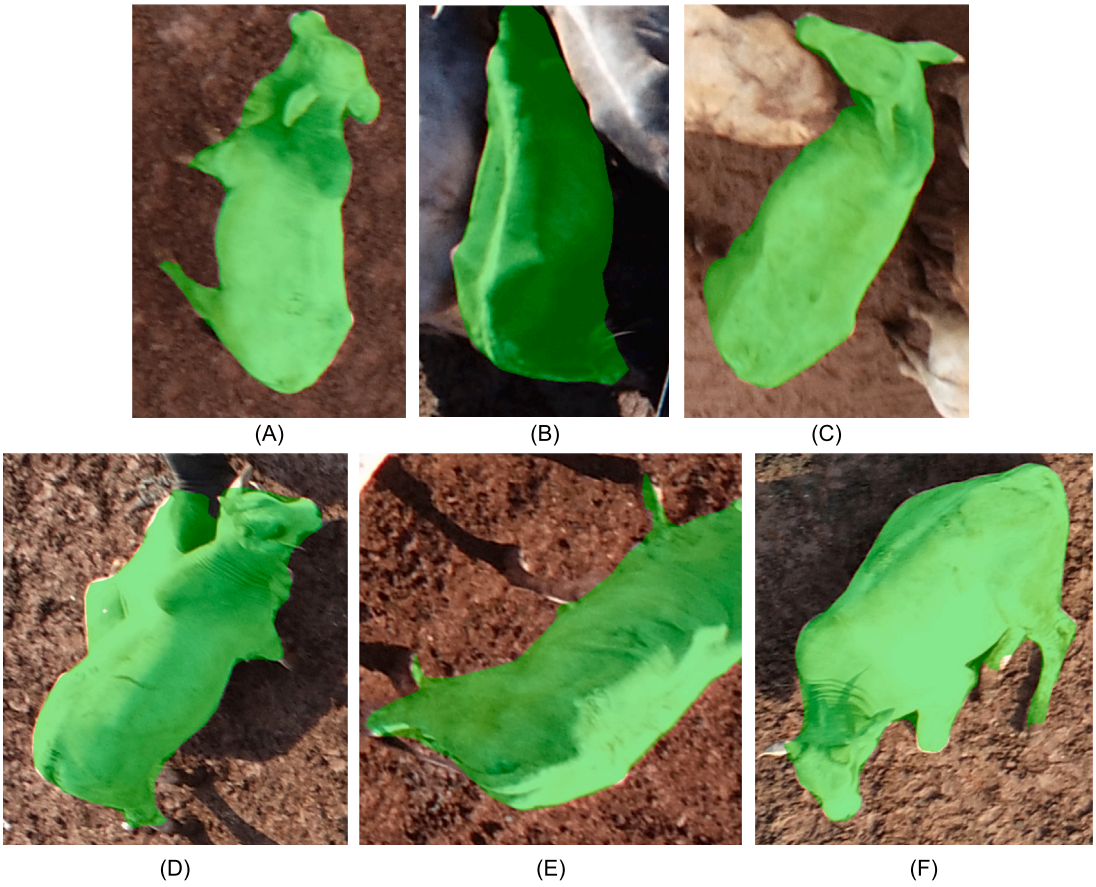

Fig. 5. Examples of rejected cattle instance segmentation with W/L ratios above the upper threshold (0.333). The green overlay indicates the SAM-generated masks. Panels (A)–(F) present invalid instances with W/L ratios of 0.533, 0.415, 0.453, 0.524, 0.519, and 0.707, respectively, corresponding to (A) mild wide-angle camera distortion, (B) occlusion by neighboring cattle, (C) large head turning angle, (D) merged instance segmentation, (E) inaccurate bounding-box localization, and (F) a misclassified *lying* individual missed by the posture classifier.

Fig. 6. Examples of rejected cattle instance segmentations with W/L ratios below the lower threshold (0.250). The green overlay indicates the SAM-generated masks. (A)–(C) show invalid instances with W/L ratios of 0.222, 0.222, and 0.209, respectively, all caused by excessive leg extension, which artificially increases the apparent body length and results in abnormally low W/L ratios.

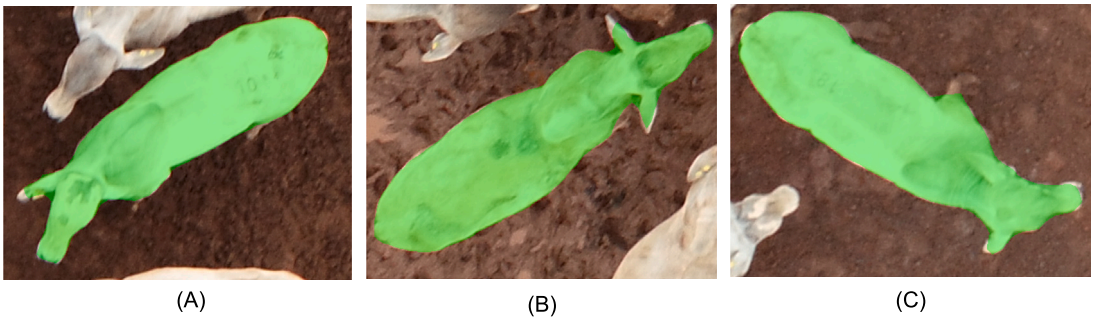

Fig. 7. Examples of retained cattle instance segmentations after posture, ROI, and W/L ratio filtering. The green overlay indicates the SAM-generated mask for each detected animal. (A)–(C) show representative valid instances with W/L ratios of 0.295, 0.285, and 0.289, respectively, which fall within the predefined W/L acceptance range $R_{W/L} = [0.250, 0.333]$ and therefore satisfy the geometric criteria for subsequent analysis.

Table 3
Standing cattle instance counts under varying Detection Confidence Thresholds (θ_{conf}), fixing $S_{ROI} = 70\%$ and $R_{W/L} = [0.250, 0.333]$.

Parameter	Day1	Day11	Day22	Day33	Day43	Day53	Day74	Day81	Day88	Day95	Day102	Day110
Before filtering	1494	443	419	417	423	388	316	508	503	275	254	226
$\theta_{conf} = 0.5$	474	119	94	93	88	92	103	118	101	63	56	56
$\theta_{conf} = 0.6$	474	119	94	93	88	92	102	118	100	63	56	56
$\theta_{conf} = 0.7$	**474**	**119**	**94**	**93**	**88**	**92**	**102**	**118**	**100**	**63**	**56**	**56**
$\theta_{conf} = 0.8$	474	119	94	93	88	92	102	118	99	63	56	56

Table 4
Standing cattle instance counts under varying ROI Size (S_{ROI}), fixing $\theta_{conf} = 0.7$ and $R_{W/L} = [0.250, 0.333]$.

Parameter	Day1	Day11	Day22	Day33	Day43	Day53	Day74	Day81	Day88	Day95	Day102	Day110
Before filtering	1494	443	419	417	423	388	316	508	503	275	254	226
$S_{ROI} = 50\%$	294	73	43	51	61	52	65	69	74	34	26	38
$S_{ROI} = 60\%$	376	97	59	70	74	74	83	97	86	48	45	47
$S_{ROI} = 70\%$	**474**	**119**	**94**	**93**	**88**	**92**	**102**	**118**	**100**	**63**	**56**	**56**
$S_{ROI} = 80\%$	554	144	137	121	106	100	131	144	121	81	61	64

standing cattle instances was quantified at each time point. This variability required a population-averaged analytical approach rather than individual tracking, as consistent identification of all animals at all time points was not feasible under field conditions.

For each time point, the mean W/L ratio was calculated from the set of filtered and validated cattle instances described in Section 2.5. The W/L ratio was chosen to yield an ascending S-shaped trajectory suitable for logistic growth modeling, reflecting the progressive increase in body width relative to length during the fattening process.

The controlled experiments summarized in Tables 3–5 show that the number of retained instances is relatively insensitive to Detection Confidence Thresholds above 0.5, but varies with ROI Size and the W/L Range, confirming that geometric filtering criteria primarily determine the effective sample size.

3.3.3. Population-level growth curve fitting

To characterize cattle morphological development over the fattening period, population-level growth trajectories were modeled using the mean W/L ratio aggregated from filtered standing cattle instances at each acquisition time point. A logistic growth function (Eq. (7)) was fitted using non-linear least squares optimization (Levenberg–Marquardt algorithm).

Table 5
Standing cattle instance counts under varying W/L Range ($R_{W/L}$), fixing $\theta_{conf} = 0.7$ and $S_{ROI} = 70\%$.

Parameter	Day1	Day11	Day22	Day33	Day43	Day53	Day74	Day81	Day88	Day95	Day102	Day110
Before filtering	1494	443	419	417	423	388	316	508	503	275	254	226
$R_{W/L} = [0.250, 0.33]$	**474**	**119**	**94**	**93**	**88**	**92**	**102**	**118**	**100**	**63**	**56**	**56**
$R_{W/L} = [0.26, 0.33]$	455	114	91	88	86	90	101	110	98	63	55	56
$R_{W/L} = [0.27, 0.33]$	423	99	79	82	80	84	93	106	90	61	54	53
$R_{W/L} = [0.28, 0.33]$	350	87	62	71	66	77	79	91	85	56	53	52

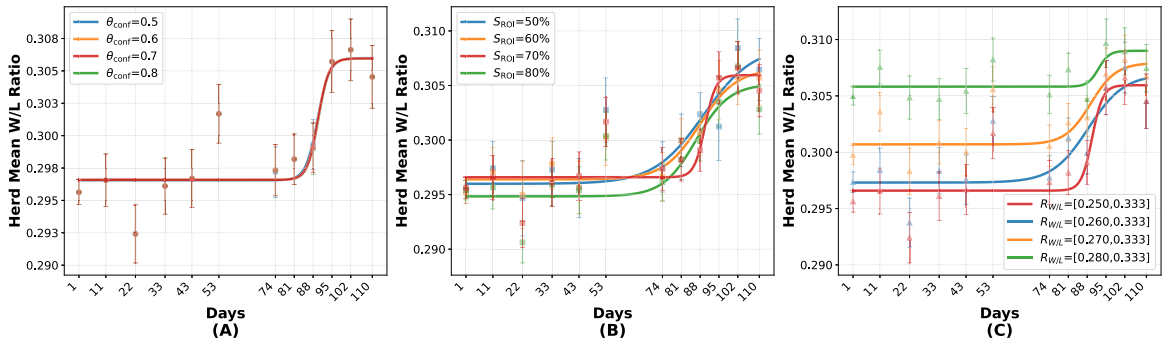

Fig. 8. Population-level W/L growth trajectories under different filtering configurations. (A) Effect of Detection Confidence Threshold ($\theta_{conf} = 0.5$ to 0.8); (B) Effect of ROI Size ($S_{ROI} = 50\%$ to 80%); (C) Effect of W/L Range ($R_{W/L} = [0.250, 0.333]$ to $[0.280, 0.333]$). In all panels, markers denote observed herd-level mean W/L ratios at each time point (distinct shapes correspond to different parameter settings), vertical error bars represent the SEM ($SD/\sqrt{n}$) of the fitted logistic model, and solid curves indicate the corresponding logistic fits. The red curves highlight the selected optimal configuration ($\theta_{conf} = 0.7$, $S_{ROI} = 70\%$, $R_{W/L} = [0.250, 0.333]$).

Rather than relying on a single parameter configuration, a series of sensitivity analyses were conducted to examine how Detection Confidence Threshold (θ_{conf}), ROI Size (S_{ROI}), and W/L Range ($R_{W/L}$) influence the inferred growth trajectories. For each parameter setting, the mean W/L values at all time points were independently fitted with a logistic model, enabling a direct comparison of growth curve shapes and goodness-of-fit metrics.

Fig. 8 presents representative growth trajectories under different parameter variations, where each subpanel corresponds to one controlled parameter: Detection Confidence Threshold (A), ROI Size (B), and W/L Range (C). In all cases, the observed data points represent the empirical mean W/L values at each time point, while the error bars indicate the corresponding Standard Error of the Mean (SEM = $SD/\sqrt{n}$) at each time point. The experimental results shown in Fig. 8 demonstrate that the fitted S-shaped logistic growth function is robust to variations in Detection Confidence Threshold, ROI Size, and W/L Range parameters.

To further assess whether parameter variations influence not only the fitted growth curves but also the inferred growth dynamics, the first derivative of each fitted logistic model was computed and compared across parameter configurations (Fig. 9). This derivative-based analysis provides a direct characterization of temporal growth velocity and highlights the timing and sharpness of peak morphological development. Across all tested configurations, the resulting growth-rate trajectories consistently exhibited a unimodal pattern with a well-defined peak near the inflection point, indicating that the timing of maximum morphological change is robust to reasonable variations in detection confidence, spatial filtering, and W/L acceptance ranges.

For all tested parameter configurations, the fitted logistic models yielded statistically significant results ($p < 0.01$; Tables 6–8), indicating that the observed temporal evolution of the population-level W/L ratio cannot be explained by random variation alone. Here, the reported p-values are used to confirm the statistical validity of the growth model fits rather than to perform hypothesis testing between different parameter settings.

Across the Detection Confidence Thresholds ranging from 0.5 to 0.8, the resulting growth curves exhibited highly consistent S-shaped trajectories, with minimal variation in both inflection timing and asymptotic behavior. As summarized in Table 6, the coefficient of determination ($R^2 \approx 0.77$) and RMSE values remained nearly invariant, indicating that the population-level W/L trajectory is largely insensitive to moderate changes in detection confidence once low-quality detections are removed.

Table 6
Goodness-of-fit statistics of logistic growth models under different Detection Confidence Threshold, fixing $S_{ROI} = 70\%$ and $R_{W/L} = [0.250, 0.333]$.

θ_{conf}	R^2	RMSE	p-value
$\theta_{conf} = 0.5$	0.772	0.002038	1.72×10^{-3}
$\theta_{conf} = 0.6$	0.771	0.002041	1.77×10^{-3}
$\theta_{conf} = 0.7$	**0.771**	**0.002041**	1.77×10^{-3}
$\theta_{conf} = 0.8$	0.771	0.002042	1.84×10^{-3}

In contrast, variations in ROI Size and W/L Range produced more pronounced effects on curve stability and fitting quality (Tables 7 and 8). Overly restrictive ROIs (e.g., $S_{ROI} = 60\%$) retained fewer valid instances (Table 7), leading to increased temporal sparsity and noisier trajectories, particularly toward the upper portion of the growth curve, despite yielding a relatively higher R^2. Conversely, excessively permissive ROIs (e.g., $S_{ROI} = 80\%$) admitted a larger number of peripheral instances affected by wide-angle distortion and perspective bias, which could not be fully mitigated by subsequent filtering and resulted in reduced fitting quality and lower R^2 values. A moderate ROI size of $S_{ROI} = 70\%$ provided a balanced trade-off between instance quantity and geometric reliability, yielding smoother temporal trajectories and more stable curve fitting. Similar trends were consistently observed across different W/L acceptance ranges, where progressively tightening the range reduced the number of retained instances, especially at the last observation time points, leading to degraded curve smoothness and explanatory power. The qualitative consistency of these patterns across parameter configurations demonstrates the robustness of the proposed filtering strategy.

Based on the goodness-of-fit statistics presented in Tables 6–8, the parameter configuration $\theta_{conf} = 0.7$, $S_{ROI} = 70\%$, $R_{W/L} = [0.250, 0.333]$ was identified as a suitable setting, which provides a favorable balance

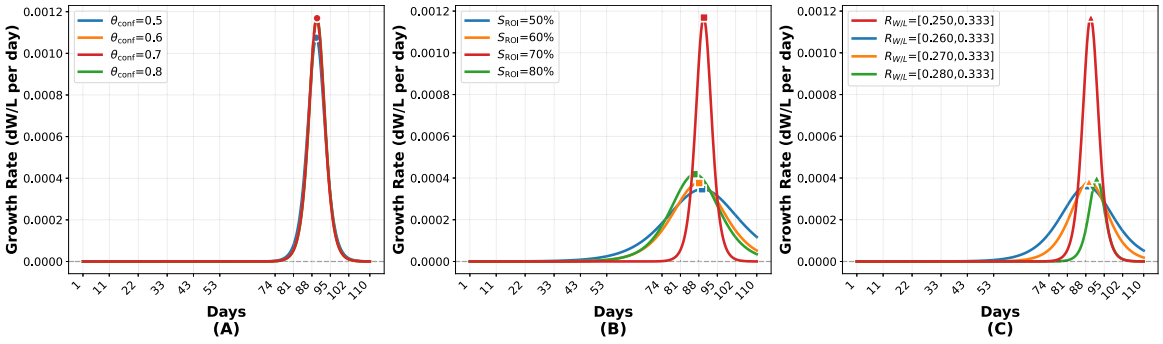

Fig. 9. Temporal evolution of population-level growth rates derived from the first derivative of the fitted logistic models under different instance filtering configurations. (A) Growth-rate trajectories obtained using different detection confidence thresholds. (B) Growth-rate trajectories under varying region of interest (ROI) sizes. (C) Growth-rate trajectories corresponding to different W/L ratio filtering ranges. In each subpanel, curves represent the mean growth-rate trajectories inferred from one parameter configuration, while the red curve highlights the selected optimal setting ($\theta_{conf} = 0.7$, $S_{ROI} = 70\%$, $R_{W/L} = [0.250, 0.333]$).

Table 7
Goodness-of-fit statistics of logistic growth models under different ROI Size configurations, fixing $\theta_{conf} = 0.7$ and $R_{W/L} = [0.250, 0.333]$.

ROI Size	$\mathbf{R}^2$	RMSE	*p*-value
S_{ROI} = 50%	0.703	0.002315	1.27×10^{-3}
S_{ROI} = 60%	0.865	0.001406	4.17×10^{-4}
S_{ROI} = 70%	**0.771**	**0.002041**	1.77×10^{-3}
S_{ROI} = 80%	0.719	0.002262	1.69×10^{-3}

Table 8
Goodness-of-fit statistics of logistic growth models under different W/L filtering ranges, fixing $\theta_{conf} = 0.7$, $S_{ROI} = 70\%$.

W/L Range	$\mathbf{R}^2$	RMSE	*p*-value
$R_{W/L} = [0.250, 0.333]$	**0.771**	**0.002041**	1.77×10^{-3}
$R_{W/L} = [0.260, 0.333]$	0.663	0.002244	3.36×10^{-3}
$R_{W/L} = [0.270, 0.333]$	0.629	0.001895	1.08×10^{-2}
$R_{W/L} = [0.280, 0.333]$	0.426	0.001313	1.03×10^{-1}

between fitting quality (training $R^2 = 0.771$) and temporal consistency of the inferred growth trajectories (Figs. 8–9). Therefore, this configuration was formally selected as the default setting for subsequent analyze.

Under the selected parameter configuration, the fitted logistic model achieved an R^2 of 0.771 and an RMSE of 0.002 in W/L units (Table 9). The estimated lower asymptote ($b = 0.296$) represents baseline body proportions at the onset of monitoring, while the final asymptote ($b + L = 0.307$) corresponds to the morphological plateau associated with market-ready cattle. The inflection point ($t_0 = 87.950$ days) indicates the period of maximum morphological change, separating the slow initial growth phase from the subsequent plateau.

An uncertainty analysis of the fitted logistic growth model parameters shows that the asymptotic increment (L) and the growth rate parameter (k) present substantially higher coefficients of variation than the lower asymptote (b) and the inflection point (t_0) (Table 10). This behavior is mainly attributable to the reduced number of valid cattle instances available at the last observation time points, where sample sizes decrease due to posture variability, thereby limiting the information content for parameters that are primarily constrained by observations near the upper tail of the growth trajectory. In contrast, the lower asymptote and the inflection point remain well constrained, reflecting the denser data coverage during earlier portions of the observation period.

Despite the increased uncertainty associated with L and k, the inflection timing and the final asymptotic level remain stable, supporting the robustness of the inferred population-level growth dynamics.

Overall, the resulting growth curves consistently display a canonical sigmoidal pattern, reflecting an initial adaptation phase, a rapid morphological development period, and a saturation phase toward the end of the feedlot cycle. These findings confirm that population-averaged W/L trajectories derived from UAV imagery provide a stable and interpretable proxy for monitoring cattle growth dynamics under commercial feedlot conditions.

Although the above goodness-of-fit statistics demonstrate stable model performance under different parameter configurations, they are computed using the full set of available instances. To further assess whether the inferred population-level growth trajectory is sensitive to variations in instance availability at individual time points, we conducted a subsampling robustness analysis.

For the selected parameter configuration ($\theta_{conf} = 0.7$, $S_{ROI} = 70\%$, $R_{W/L} = [0.250, 0.333]$), repeated random subsampling was performed at two retention levels: 90% and 80% of available instances. Within each time point, instances were randomly retained while preserving temporal coverage, and the logistic growth model was refitted independently. This procedure was repeated 100 times for each retention level, yielding distributions of fitted growth trajectories and goodness-of-fit statistics.

Fig. 10 illustrates the resulting population-level W/L growth trajectories under subsampling. The solid curves represent the mean fitted trajectories across 100 iterations for each retention level, while the shaded regions indicate the corresponding standard deviation envelopes. The close overlap between the baseline trajectory, generated using 100% data samples, and the subsampled trajectories demonstrates that the inferred growth dynamics are stable even when 10%–20% of instances are randomly excluded at each time point.

Table 11 summarizes the goodness-of-fit statistics under different subsampling levels. The coefficient of determination (R^2) remains consistently high across all configurations (0.722–0.771), with only a modest decrease as the retention fraction decreases. Similarly, RMSE values exhibit limited variation, indicating that moderate reductions in instance availability do not substantially degrade the quality of the fitted growth model.

Overall, the subsampling robustness analysis confirms that the population-level W/L growth trajectory is stable under realistic variations in instance availability. The limited degradation in goodness-of-fit metrics under 20% data removal indicates that the inferred growth dynamics are not driven by a small subset of instances. This robustness is particularly relevant in practical UAV-based monitoring scenarios, where posture variability and occlusion can lead to uneven sample sizes across acquisition time points.

3.3.4. Growth dynamics and phase analysis

Based on the optimal parameter configuration identified in the previous section ($\theta_{conf} = 0.7$, $S_{ROI} = 70\%$, $R_{W/L} = [0.250, 0.333]$), growth dynamics were further examined by computing the first derivative of

Table 9
Logistic growth parameters for mean W/L ratio.

Parameter	Symbol	Fitted Value	Interpretation
Lower asymptote	b	0.296	Baseline W/L ratio representing young cattle.
Upper asymptote increase	L	0.010	Indicates the magnitude of morphological change during growth.
Final asymptote	$L + b$	0.307	Corresponds to mature W/L ratio at market-ready stage.
Growth rate	k	0.149 day^{-1}	Reflects the speed of morphological transition.
Inflection point	t_0	87.950 days	Represents the time of maximum growth velocity.
Goodness-of-fit metrics:			
Coefficient of Determination	R^2	0.771	Indicates a satisfactory model fit.
Root Mean Square Error	RMSE	0.002	Reflects low prediction error.

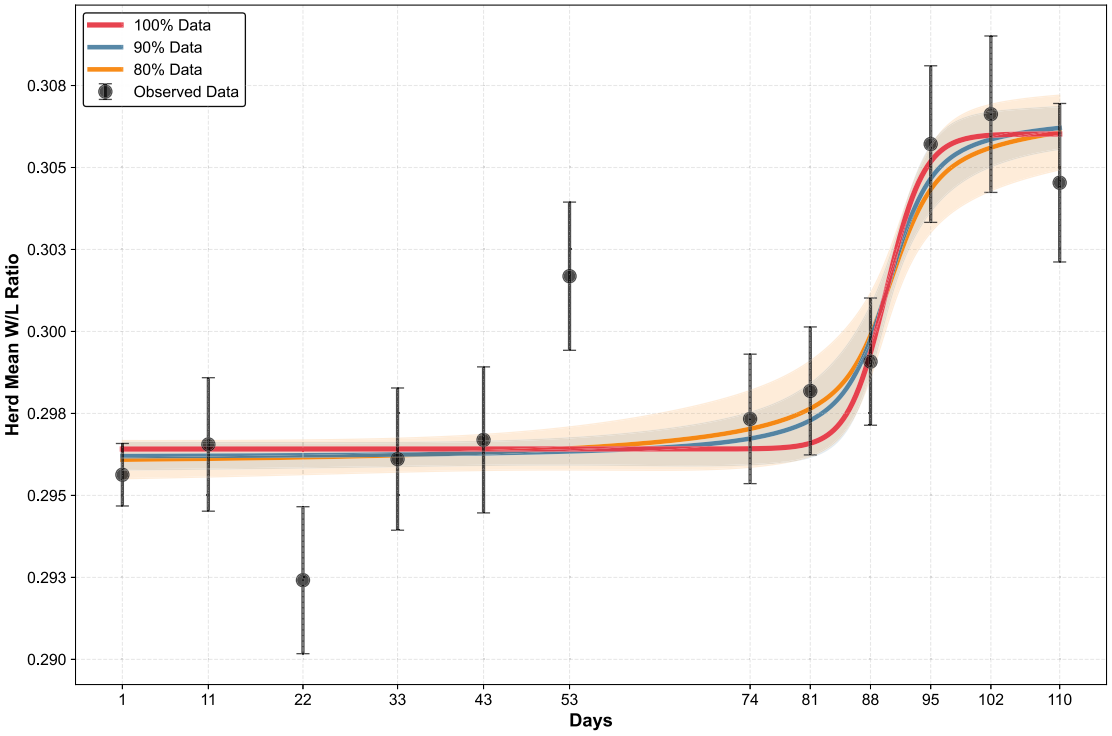

Fig. 10. Subsampling robustness analysis of population-level growth trajectories. Black markers with error bars represent observed herd-level mean W/L ratios ± SEM at each time point. Solid curves show the fitted logistic growth models using 100% (blue), 90% (purple), and 80% (orange) of retained instances. For the 90% and 80% subsampled datasets, each curve represents the mean trajectory across 100 independent subsampling iterations, and the shaded regions indicate ±1 standard deviation envelopes. The strong overlap among curves indicates robustness of the inferred growth dynamics to moderate reductions in sample size.

Table 10
Uncertainty analysis of the fitted logistic growth model parameters under the selected parameter configuration ($\theta_{conf} = 0.7$, $S_{ROI} = 70\%$, $W/L = [0.250, 0.333]$).

Parameter	Estimate	Std. Error	95% CI	CV (%)
b (lower asymptote)	0.296	0.00063	[0.295, 0.298]	0.21
L (asymptotic increment)	0.010	0.00247	[0.007, 0.013]	24.42
k (growth rate)	0.149	0.09048	[−0.029, 0.326]	60.90
t_0 (inflection point, days)	87.950	5.08000	[77.990, 97.910]	5.78

Table 11
Goodness-of-fit statistics under subsampling conditions. Values for 90% and 80% retention represent mean ± standard deviation across 100 random subsampling iterations.

Data Fraction	$\mathbf{R}^2$ (mean ± SD)	RMSE (mean ± SD)
100% (Baseline)	0.771 ± 0.000	0.00204 ± 0.00000
90% Retention	0.747 ± 0.015	0.00225 ± 0.00031
80% Retention	0.722 ± 0.021	0.00234 ± 0.00038

the fitted logistic growth function (Eq. (9)). In this study, the species-specific threshold value calculated by Eq. (9) is $\theta_{threshold} = 98.7$. Again, this value is just an approximation and it may vary from case to case. This derivative represents the instantaneous rate of change in the population-level mean W/L ratio and provides a direct quantitative measure of morphological growth velocity.

For the selected optimal configuration (highlighted in red), three biologically interpretable growth phases can be identified. (i) During the initial lag phase (approximately Days 0–80), growth rates remain close to zero, indicating limited observable changes in body proportions during early adaptation. (ii) The rapid growth phase (approximately Days 80–100) is characterized by a sharp increase and subsequent decline in growth velocity, with the peak occurring near the inflection point ($t_0 \approx 88$ days). This phase corresponds to the main fattening stage in feedlot production, during which cattle exhibit the most pronounced improvements in body condition. (iii) In the terminal plateau phase (after Day 100), growth rates decrease toward zero, reflecting diminishing marginal morphological gains as animals approach their asymptotic body proportions.

It should be noted that the exact timing information presented here is specific to the analysis of the current dataset and may vary across different datasets. Nevertheless, we expect the qualitative phases of cattle development to be preserved.

4. Discussion

This study proposes a UAV-based framework for population-level cattle growth monitoring by integrating automatic detection, instance segmentation, geometric filtering, and logistic growth modeling. Beyond demonstrating the technical feasibility of extracting morphological indicators from aerial imagery, the results provide insights into the robustness, interpretability, and practical relevance of population-averaged growth dynamics under commercial feedlot conditions.

4.1. Robustness of population-level growth inference

A central concern when applying automated vision-based pipelines in real-world production environments is the sensitivity of downstream biological inferences to parameter selection and data filtering strategies. In this study, systematic perturbations of detection confidence thresholds, region-of-interest (ROI) definitions, and W/L acceptance ranges produced only minor variations in fitted growth curves and growth-rate profiles. Notably, the timing of the inflection point and the unimodal structure of the growth-rate trajectories remained consistent across all tested configurations.

This stability suggests that the inferred growth dynamics are not artifacts of a particular parameter setting, but instead reflect genuine population-level biological patterns. While stricter filtering criteria reduced sample size and increased late-stage variability, especially under narrow W/L ranges or small ROIs, the overall sigmoidal growth structure and the transition from rapid development to saturation were preserved. Such robustness is particularly important in UAV-based monitoring scenarios, where acquisition conditions, animal density, and posture distributions cannot be fully controlled.

The results further indicate that geometric filtering criteria, especially ROI definition and W/L range selection, exert a stronger influence on effective sample size than detection confidence once low-quality detections are removed. This observation highlights the importance of explicitly reporting filtering strategies in future studies to ensure reproducibility and fair comparison across datasets and operational settings.

4.2. Interpretation of population-averaged morphological growth

Due to the lack of persistent individual identifiers and the high degree of occlusion in dense feedlot environments, reliable longitudinal tracking of individual animals remains infeasible using overhead UAV imagery alone. Consequently, growth analysis in this study was conducted at the herd level by aggregating instance-level morphological measurements across time points.

While this population-averaged approach does not capture individual growth heterogeneity, it provides a stable and scalable representation of developmental trends that is well aligned with the operational realities of commercial feedlots. The observed sigmoidal trajectory of the mean W/L ratio is consistent with established biological expectations of cattle fattening, characterized by an initial adaptation phase, a period of rapid morphological development, and a terminal plateau with diminishing marginal gains.

Importantly, the consistency of the inferred inflection point across parameter configurations supports its interpretation as a meaningful biological transition rather than a statistical artifact. This reinforces the suitability of population-level W/L trajectories as a proxy for monitoring collective growth dynamics when individual-level measurements are impractical.

4.3. Implications for economic decision-making

By interpreting the first derivative of the fitted logistic model as an operational indicator of growth efficiency, this study translates biological growth dynamics into a practical decision-support framework for feedlot management. The decline of the population-level growth rate below a predefined threshold marks the transition from biologically efficient growth to a terminal phase characterized by diminishing returns on feed investment.

Rather than prescribing a fixed slaughter date, the proposed approach identifies an economically appropriate marketing window that balances residual growth potential with operational flexibility. This data-driven strategy contrasts with conventional management practices that rely on fixed feeding durations independent of herd-specific growth trajectories and environmental conditions.

Although the specific growth-rate threshold used in this study is system-dependent and may vary across production environments, the general framework is extensible. With appropriate calibration, it can be adapted to different breeds, feeding regimes, and market contexts, providing a quantitative basis for adaptive marketing decisions grounded in observed biological dynamics.

4.4. Limitations and future directions

Several limitations of the present study warrant consideration. First, the analysis is inherently population-averaged and does not account for individual-level variability in growth trajectories. Future work integrating animal identification technologies or multi-view tracking may enable individual longitudinal analysis and a more nuanced understanding of growth heterogeneity.

Second, morphological indicators derived from UAV imagery represent indirect proxies for body condition and weight gain. While the W/L ratio demonstrated favorable growth dynamics in this study, integrating additional features such as body area, volume estimation, or multimodal data sources could further improve biological interpretability.

Finally, the proposed framework was evaluated under a single commercial feedlot setting. Extending validation across different management systems, climatic conditions, and production stages will be essential to assess the generalizability of the approach.

Despite these limitations, the present study demonstrates that UAV-based morphological monitoring combined with robust growth modeling provides a stable, interpretable, and operationally meaningful tool for assessing cattle growth dynamics and supporting evidence-based management decisions in modern feedlot systems.

5. Conclusion

This study presented a comprehensive UAV-based computer vision framework for fully automated, non-invasive, and longitudinal monitoring of beef cattle growth under commercial feedlot conditions, rather than a direct weight-estimation model. By combining high-resolution aerial imagery with YOLOv11s for robust cattle detection and posture-aware filtering and SAM 2.1 for instance segmentation, the system achieved reliable individual identification and morphological measurement throughout a 112-day production cycle. Both models demonstrated strong performance, YOLOv11s with high precision (95.56%) and recall (98.36%), and SAM 2.1 with consistently accurate segmentation masks under varying illumination, occlusion, and posture conditions, without requiring additional fine-tuning.

To obtain biologically reliable measurements, an automated instance selection pipeline was employed to retain only laterally oriented animals with minimal body curvature. From these filtered instances, key morphological traits such as body length, width, and the W/L ratio were extracted and used as population-level indicators of growth status.

Modeling the temporal evolution of the W/L ratio using a logistic growth function revealed distinct growth phases and provided a stable representation of herd-level developmental dynamics. Based on this modeling framework, an economically appropriate slaughter window was identified, corresponding to the transition from biologically efficient growth to a terminal low-efficiency phase. This window balances marginal morphological gains against continued feeding costs and offers operational flexibility for marketing decisions under commercial feedlot conditions. While the fitted model accurately captured population-level trends, the present study does not explicitly quantify economic returns; therefore, the direct financial impact of the proposed decision rule warrants further investigation.

In general, the proposed framework demonstrates the feasibility and practical value of UAV-based image-driven approaches as scalable tools for precision livestock management. By providing high-throughput and non-invasive assessments of cattle growth, the system supports evidence-based decisions related to growth monitoring, feed efficiency evaluation, and marketing strategy optimization in commercial beef production systems.

Future work will extend the proposed framework to additional breeds and production systems to assess its generalizability across genetic backgrounds and management conditions. In addition, integration of real-time kinematic (RTK) positioning and economic performance indicators may further enhance measurement precision and decision support capabilities.

CRediT authorship contribution statement

Jianglong Yan: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation. **Everton C. Tetila:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Liang Zhao:** Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization. **Rian C. Gonçalves:** Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation. **Letícia F. Castanheiro:** Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation. **Lucas P. Valem:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation. **Jayme G.A. Barbedo:** Writing – review & editing, Writing – original draft, Validation, Supervision, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We thank the management and staff of Fazenda Campanário for providing access to the feedlot facilities and confined cattle. We also express our appreciation to the farm's veterinary team for their valuable collaboration and technical support during data collection and animal handling. We acknowledge the computational resources provided by the RTX Ada 5000 GPU infrastructure. We also thank the University of São Paulo (PRPI Ordinance No. 1032, New Faculty Support Program). This work was partially supported by the São Paulo Research Foundation (FAPESP) under Grants #2023/03870-8, #2022/09319-9, #2024/15430-5, #2025/05985-2, and #2025/10602-5 and by the Foundation for Support to the Development of Education, Science and Technology of the State of Mato Grosso do Sul (FUNDECT) under Grants #83/029.295/2024 and #83/044.460/2024, and by the Brazilian National Council for Scientific and Technological Development (CNPq) under Grants #406417/2022-9, 403307/2025-2, and #306788/2023-3.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.compag.2026.111559.

Data availability

Data will be made available on request.

References

Afridi, H., Ullah, M., Nordbø, Ø., Hoff, S.C., Furre, S., Larsgard, A.G., Cheikh, F.A., 2024. Analyzing data modalities for cattle weight estimation using deep learning models. J. Imaging 10, 72.

Aquilani, C., Confessore, A., Bozzi, R., Sirtori, F., Pugliese, C., 2022. Precision livestock farming technologies in pasture-based livestock systems. Animal 16, 100429.

Bai, L., Guo, C., Song, J., 2025. Cattle weight estimation model through readily photos. Eng. Appl. Artif. Intell. 143, 109976.

Besler, B.C., Mojabi, P., Lasemiimeni, Z., Murphy, J.E., Wang, Z., Baker, R., Pearson, J.M., Fear, E.C., 2024. Scoping review of precision technologies for cattle monitoring. Smart Agric. Technol. 10, 100621.

Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S., 2020. End-to-end object detection with transformers. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16. Springer, pp. 213–229.

Costa Jr, C., Tedeschi, L.O., Gonzalez-Quintero, R., Arango, J., Burkart, S., Grosjean, G., Dittmer, K.M., Wollenberg, E., Becoña, G., Micol, L., et al., 2025. South America's pasture intensification can increase beef production, reduce emissions by 30% and mitigate warming from methane by 2050. Sci. Rep. 15, 35734.

Curti, P.F.d.F., Silva, W.H.d., Sá Filho, M.F.d., Basso, A.C., Nogueira, E., Taira, A.R., Martins, S.M.B., Abe, A.U., Braga, J.R.M., Pantoja, M.H., et al., 2023. Applications of livestock monitoring devices and machine learning in livestock reproduction and production. Anim. Reprod. 20, e20230077.

Firdaus, F., Atmoko, B.A., Baliarti, E., Widi, T.S. M., Maharani, D., Panjono, P., 2023. The meta-analysis of beef cattle body weight prediction using body measurement approach with breed, sex, and age categories. J. Adv. Vet. Anim. Res. 10, 630.

Giannone, C., Sahraeibelverdy, M., Lamanna, M., Cavallini, D., Formigoni, A., Tassinari, P., Torreggiani, D., Bovo, M., 2025. Automated dairy cow identification and feeding behaviour analysis using a computer vision model based on YOLOv8. Smart Agric. Technol. 12, 101304.

He, K., Gkioxari, G., Dollár, P., Girshick, R., 2017. Mask R-CNN. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2961–2969.

Hou, Z., Huang, L., Zhang, Q., Miao, Y., 2023. Body weight estimation of beef cattle with 3D deep learning model: PointNet++. Comput. Electron. Agric. 213, 108184.

Jocher, G., Chaurasia, A., Qiu, J., 2023. Ultralytics yolov8. https://github.com/ultralytics/ultralytics. (Accessed 10 January 2023).

Jocher, G., Qiu, J., 2024. Yolo11: ultralytics yolo11. https://github.com/ultralytics/ultralytics. (Accessed 27 September 2024).

Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al., 2023. Segment anything. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4015–4026.

Lathuillière, M.J., Flach, R., Wang-Erlandsson, L., Ribeiro, V., zu Ermgassen, E.K.H.J., Souza, C.M., 2025. International reliance on Brazil's water through soy and beef supply chains. Commun. Earth & Environ. 6, 688.

Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L., 2014. Microsoft COCO: Common objects in context. In: Computer Vision–ECCV 2014: 13th European Conference. Zurich, Switzerland September 6–12, 2014, Proceedings, Part V 13, Springer, pp. 740–755.

Los, S., Mücher, C.A., Kramer, H., Franke, G.J., Kamphuis, C., 2023. Estimating body dimensions and weight of cattle on pasture with 3D models from UAV imagery. Smart Agric. Technol. 4, 100167.

Ravi, N., Gabeur, V., Hu, Y.T., Hu, R., Ryali, C., Ma, T., Khedr, H., Rädle, R., Rolland, C., Gustafson, L., et al., 2024. Sam 2: segment anything in images and videos. arXiv preprint arXiv:2408.00714.

Ren, S., He, K., Girshick, R., Sun, J., 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In: Advances in Neural Information Processing Systems, 28.

Wang, A., Chen, H., Liu, L., Chen, K., Lin, Z., Han, J., Ding, G., 2024. Yolov10: real-time end-to-end object detection. arXiv preprint arXiv:2405.14458.

Wang, Y., Mücher, S., Wang, W., Kooistra, L., 2025. Body weight estimation of cattle in standing and lying postures using point clouds derived from UAV-based LiDAR. Drones 9, 84.

Weber, V.A.M., Weber, F.d.L., Oliveira, A.d.S., Astolfi, G., Menezes, G.V., Porto, J.V.d.A., Rezende, F.P.C., de Moraes, P.H., Matsubara, E.T., Mateus, R.G., et al., 2020. Cattle weight estimation using active contour models and regression trees bagging. Comput. Electron. Agric. 179, 105804.

Xiong, Y., Condotta, I.C.F.S., Musgrave, J.A., Brown-Brandl, T.M., Mulliniks, J.T., 2023. Estimating body weight and body condition score of mature beef cows using depth images. Transl. Animal Sci. 7, txad085.

Xu, B., Mao, Y., Wang, W., Chen, G., 2024. Intelligent weight prediction of cows based on semantic segmentation and back propagation neural network. Front. Artif. Intell. 7, 1299169.

Zhao, Y., Lv, W., Xu, S., Wei, J., Wang, G., Dang, Q., Liu, Y., Chen, J., 2024. DETRs beat YOLOs on real-time object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16965–16974.