Cow Face Detection for Precision Livestock Management using YOLOv8

Umair Ali[1], Wasif Muhammad[1*]
[1]Intelligent Systems Laboratory, Department of Electrical Engineering, University of Gujrat, Pakistan.
*** Correspondence:** Wasif Muhammad: syed.wasif@uog.edu.pk

Citation | Ali. U, Muhammad. W, "Cow Face Detection for Precision Livestock Management using YOLOv8", IJIST, Special Issue. pp 128-132, March 2025
Received | Feb 17, 2025 **Revised |** March 01, 2025 **Accepted |** March 07, 2025 **Published |** March 10, 2025.

Precision livestock management is transforming traditional agricultural practices by boosting productivity, increasing yield, and automating tasks, all while reducing labor requirements and minimizing errors. Conventional methods for animal recognition are often unreliable, which has led to a growing preference for using cameras to identify animals, monitor their health, manage data, and maintain cattle records. However, small-scale farms with limited livestock, such as cows and goats, frequently face overfitting problems in traditional machine learning models due to insufficient training data. Identifying individual cows based on facial features becomes more effective after detecting the cow's face. This study addresses these challenges by fine-tuning YOLOv8, a pretrained model, using a mix of self-captured images and publicly available datasets to detect cow faces in complex environments. Integrating publicly available data and leveraging a pretrained COCO model has significantly improved the model's ability to generalize and accurately detect cow faces. YOLOv8, equipped with the COCO pretrained model, successfully detects nearly all types of cow faces, which can then be used for individual cow classification. This approach enhances cow recognition accuracy, contributing to more efficient farm management applications.

Keywords: Livestock management, Facial Recognition, Cow face detection, Transfer learning, Yolov8

Introduction:

The agriculture industry is evolving rapidly by integrating advanced technologies into traditional methods, resulting in higher efficiency and productivity. One such advancement is the emergence of data-driven approaches in farming, which has given rise to a new field known as Precision Agriculture. These innovative techniques play a key role in livestock management, helping to meet the growing global demand for food [1]. However, countries that rely heavily on livestock face challenges in food storage and agricultural resources, which are vital for human benefit. Additionally, significant issues related to animal health, labor shortages, and overall management persist [2]. As the need for increased animal production grows, it is equally essential to implement smart, efficient systems that can enhance productivity while ensuring long-term sustainability [3]. Traditional livestock management methods fall short in today's world due to their time-consuming nature, high risk of errors, and dependence on manual labor. Research shows that poor labor practices can negatively impact livestock health, increasing the likelihood of disease outbreaks due to insufficient care and ineffective management [4]. Another common issue is disputes among farmers, which often arise from the mixing of animals and the inability to track individual cows accurately [5]. To tackle these problems, various identification and tracking methods have been developed. Broadly, there are three types of cow identification techniques: permanent, semi-permanent, and temporary methods.

Permanent identification methods involve physically marking animals. Although these methods are widely used [6], they are often unreliable, especially when managing large herds. Semi-permanent methods, such as attaching tags or collars, are commonly adopted for tracking cows. However, these methods are prone to issues such as tag loss, wear and tear, and inaccuracies due to the involvement of manual labor [7]. Temporary methods, like RFID (Radio-Frequency Identification) tags, have gained popularity because they enable automated tracking. Despite their advantages, RFID systems often require frequent maintenance, and farmers report concerns about their durability. Moreover, the noise, overhead machinery, and handling stress during installation may disturb the animals. There is an increasing need for advanced solutions that can accurately record the history of cows. Camera-based systems have proven to be effective for this purpose. Unlike traditional methods, computer vision-based systems offer accurate and widely adopted solutions for cow recognition [8]. Since cows naturally recognize each other by their facial features, facial recognition technology can be a practical and efficient way to identify individual cows without the need for physical tags. This approach enhances sustainability and scalability, particularly for large herds [9].

One common challenge in object detection is dataset bias, which arises from limited and non-diverse datasets. When a model is trained on a small, specific dataset, it performs well with similar input but struggles with generalization. The scarcity of images captured under varying conditions—such as different weather, lighting, and camera angles—reduces the model's robustness. For instance, cow recognition can fail due to pose variations, including shifts in camera angles and cow movements, especially involving facial features. Convolutional Neural Networks (CNNs) are commonly used for facial recognition because of their ability to extract complex features. However, CNN models may suffer from overfitting when trained on small datasets [10]. Addressing these challenges is critical for effective livestock management, where variations in lighting, angles, and animal movement are inevitable [11]. Data preprocessing and augmentation can improve the performance of cow recognition models, but achieving generalization still requires diverse and extensive datasets. For higher accuracy, it is essential to first detect cow faces and then classify them based on their facial features. A supervised learning approach, with manual annotations to localize cow faces before classification, can enhance accuracy [12]. In the current study, a transfer learning-based approach is proposed, utilizing YOLOv8, a pre-trained model on the COCO dataset, to accurately detect cow faces. This pre-trained model can be fine-tuned with a smaller, specialized dataset to improve cow face localization. YOLO (You Only Look Once) is a cutting-edge object detection model known for its speed and real-time accuracy, making it well-suited for this task [13].

Objectives:

The objective of this research is to enhance cow face detection under varying background conditions, angles, and lighting by leveraging a pre-trained model. Specifically, the study aims to fine-tune YOLOv8, a state-of-the-art object detection model, using a dataset of annotated frames derived from self-collected videos and publicly available sources. By combining the precision of YOLOv8 with a diverse dataset, the study seeks to improve detection accuracy and ensure reliable cow face recognition in complex environments.

Materials and Methods:

Investigation Site:

The dataset for cow facial recognition was created by recording videos of 37 different cows. To generate diverse, environment-based images, every 10th frame from the videos was extracted, and cow faces were manually annotated using the Roboflow tool. Additionally, annotated images from publicly available web-based datasets were incorporated to enhance generalization and achieve higher detection accuracy. The fine-tuning of the pre-trained YOLOv8 model was performed in a simulation environment using this combined dataset.

Methodology:

The proposed methodology consists of several key steps, as outlined in Figure 1.

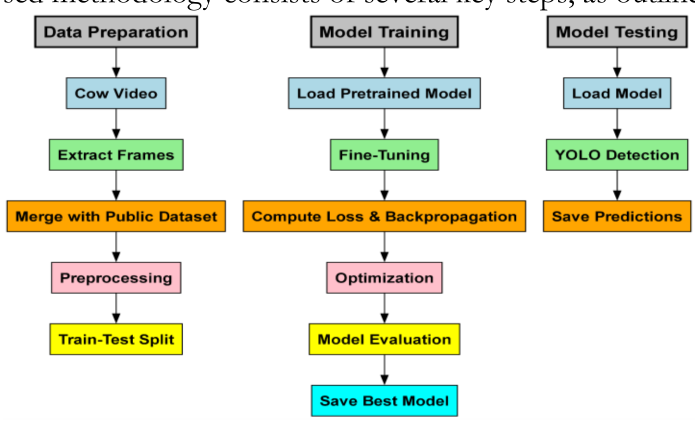

Figure 1. Workflow for Cow Face Detection using YOLOv8: Input videos are processed through frame extraction and annotated for training. The Yolov8 model is fine-tuned and optimized for model evaluation. Model is loaded and tested with web-based images.

Dataset:

Videos of 37 cows were recorded from multiple angles under different lighting and movement conditions. From these videos, 1,078 frames were extracted at regular intervals, with every 10th frame selected to form the initial dataset. To increase the dataset's size and diversity, 1,410 annotated images were added from a publicly available dataset on Roboflow. This brought the total number of images to 2,488 before preprocessing, with three distinct classes. To maintain uniformity, the external dataset classes were consolidated into a single class, labeled *CowFace*, ensuring consistency across all images. The dataset was then preprocessed to prepare it for training. All images were resized to 640×640 pixels to standardize their dimensions while preserving key features. After resizing, the dataset was split into three subsets: 1,800 images for training, 560 for validation, and 86 for testing. To enhance the training set's variability, data augmentation techniques were applied, including horizontal flips, zoom cropping (0–20%), rotations (-15° to +15°), and shear transformations (±10° both horizontally and vertically). These augmentations generated three additional versions of each image, resulting in a total of 5,096 images after preprocessing and augmentation. This process ensured a diverse dataset, optimized for effective model training.

YOLOv8 Architecture:

The basic architecture of the original YOLO version is illustrated in Figure 2.

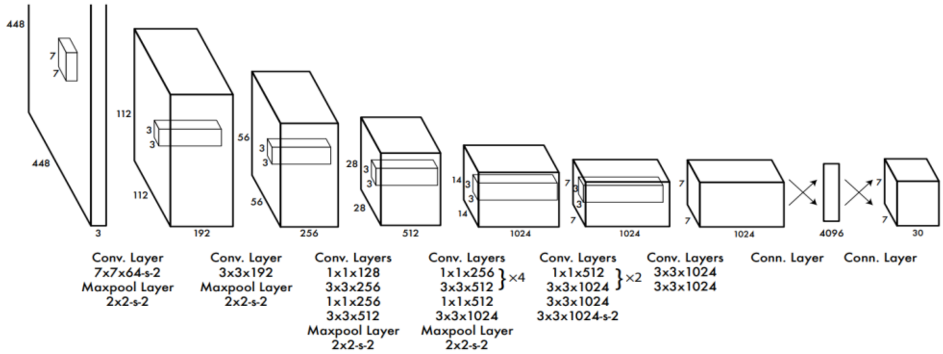

Figure 2. The YOLOv1 architecture [13] processes an input image through sequential convolutional layers for feature extraction, reduces dimensions with max-pooling, and outputs object detection predictions through fully connected layers.

YOLOv1 Architecture:

The original YOLOv1 architecture [13], proposed by Joseph Redmon et al., starts with a 7×7 convolutional layer (stride 2) for large-scale feature extraction, followed by a 2×2 max-pooling layer (stride 2). Next, the network uses alternating 1×1 and 3×3 convolutional layers to balance feature extraction and dimensionality reduction, incorporating a total of 24 convolutional layers. Finally, two fully connected layers are added, with the last layer producing a 7×7×30 output. This output includes bounding boxes, confidence scores, and class probabilities. The design integrates object detection into a single neural network to improve efficiency. Over time, several enhancements led to the development of YOLOv8, a more advanced and stable version that integrates community-driven ideas and modern techniques for better performance and flexibility. While no official paper on YOLOv8 has been published yet, the available architecture can be found in [14] and is illustrated in Figure 3.

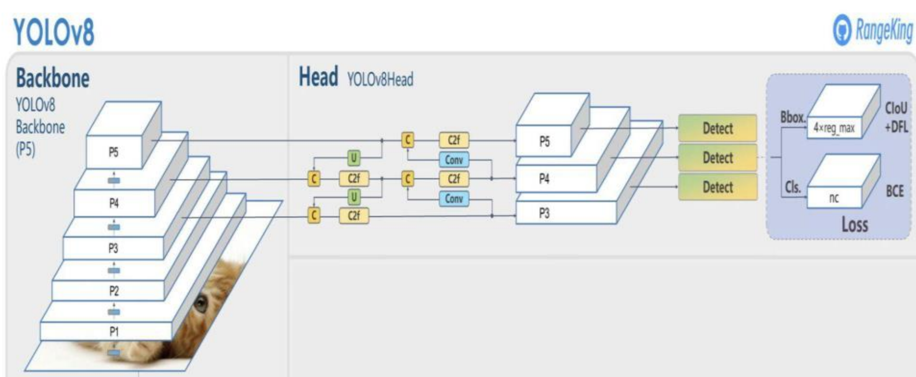

Figure 3. YOLOv8 predicts object centers directly (anchor-free), making it faster and simpler. It uses a new C2f module for better feature extraction and reduces model size with efficient design tweaks. Mosaic data mixing is used early in training but stopped later for accuracy.

YOLOv8:

YOLOv8 brings several key improvements that enhance both its performance and efficiency. It predicts object centers directly, eliminating the need for anchor boxes. This simplifies the model, making it faster and more efficient. Additionally, YOLOv8 includes a

C2f module, which is optimized to extract features more effectively from frames, even in complex environments. The model is scalable and works well on both high-performance GPUs and edge devices.

Implementation of YOLOv8:

YOLOv8 has shown outstanding performance in detecting cow faces across different test datasets and real-world web images. Its high precision is achieved by using transfer learning with COCO weights, followed by fine-tuning on a custom dataset. This dataset, created from video frames and augmented Roboflow images, captures variations in lighting, angles, and environmental conditions, ensuring reliable generalization. Trained over 50 epochs at a 640×640 resolution with a batch size of 16, the model focuses on single-class detection to streamline its accuracy on cow faces. It has already outperformed SSD, Faster R-CNN, and earlier YOLO versions in both speed and accuracy. Thanks to its real-time detection capabilities and efficient computational design, YOLOv8 is an ideal solution for livestock monitoring systems. By balancing data diversity with optimized architecture, it provides scalable and high-performance results. The implementation of YOLOv8 for cow face detection is based on three key mathematical principles: loss optimization, weight updates, and bounding box prediction. The YOLO loss function combines bounding box regression, object confidence, and classification loss to improve detection accuracy.

YOLO Trainer Framework:

The YOLO Trainer Framework fine-tunes YOLOv8 models to detect cow faces, ensuring a streamlined workflow for training, verification, and accurate area estimation of the detected cow face. This framework also automates dataset management, model training, and evaluation while saving the best-performing models for improved accuracy. Additionally, during system training, it records matrices to estimate test images, making it a reliable and effective solution for cow face detection. Table 1 below provides details on the hyperparameters used in the code.

Table 1. The YOLOTrainer fine-tunes YOLOv8 for cow face detection using epochs=50, batch=16, imgsz=640, workers=8, optimizer="SGD", patience=10, and pretrained=True. It organizes the project structure, trains on data.yaml, and saves the best model. Table 1 is showing details of hyperparameters

Hyperparameter	Value	Description
pretrained_model	"yolov8n.pt"	Pretrained YOLO model used for training
data	self.data_path	Path to the dataset configuration file (data.yaml)
epochs	50	Number of training epochs
batch	16	Batch size during training
imgsz	640	Input image size (in pixels)
workers	8	Number of worker threads for data loading
optimizer	"SGD"	Optimization algorithm used for training
patience	10	Number of epochs with no improvement before early stopping
pretrained	True	Whether to use a pretrained model
project	self.results_dir	Directory to store training results
name	"fine_tune_coco"	Experiment name for saving results
pretrained_model	"yolov8n.pt"	Pretrained YOLO model used for training
data	self.data_path	Path to the dataset configuration file (data.yaml)

Results:

YOLOv8-based facial recognition has successfully detected cow faces from test datasets and real-world web images. The model was fine-tuned using a custom dataset that included video frame captures and additional images from a publicly available Roboflow dataset. This dataset was selected for its diversity, featuring variations in angles, lighting conditions, and cow face characteristics. The fine-tuning process involved 50 training epochs, a batch size of 16, and image resizing to 640×640 pixels, aligning with the preprocessing steps. YOLOv8's flexibility in handling single-class detection and its compatibility with augmented datasets made it a perfect fit for this research. Compared to earlier YOLO versions and other object detection frameworks like SSD and Faster R-CNN, YOLOv8 demonstrated superior inference speed and overall performance.

Figure 4. Left: Manually labelled image of cow faces; Right: YOLOv8 prediction of the same image. The model, fine-tuned on a diverse custom dataset, accurately detects cow faces under varying conditions, as shown in the comparison.

The trained model was also tested with various online images. Figure 5 below illustrates the model's functionality, demonstrating its ability to generalize and detect any cow face effectively.

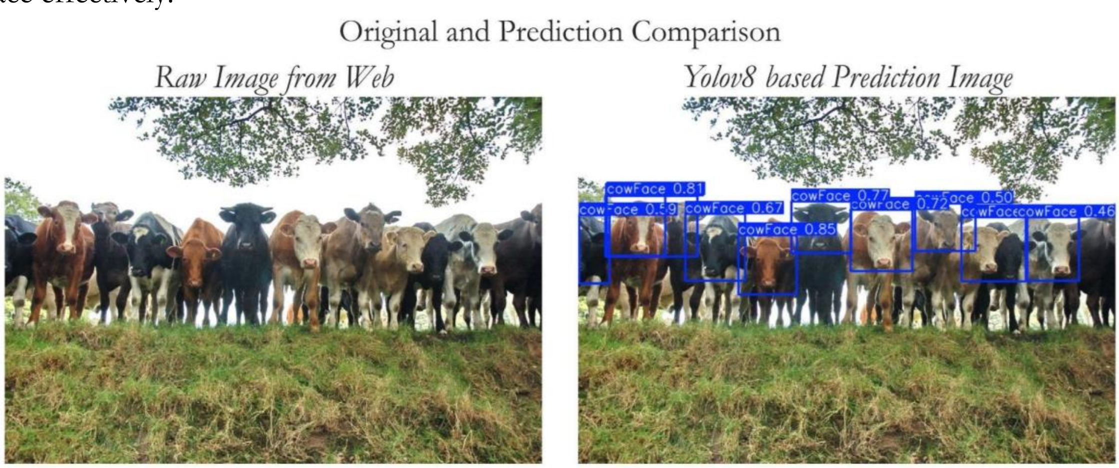

Figure 5 shows a raw image sourced from the web on the left side, representing the original picture. After processing, the fine-tuned model accurately detects and labels the cow face, with the processed image displayed on the right side. The model's ability to generalize on smaller datasets without overfitting enhances its practical utility. Additionally, its streamlined deployment and real-time prediction capabilities make it a powerful tool for livestock monitoring. The combination of a robust architecture, transfer learning using COCO weights, and a fine-tuned dataset ensured YOLOv8's high performance in detecting cow faces, proving its superiority over conventional methods.

Figure 6 illustrates the progression of key performance metrics during training. Precision stabilizes at around 0.976, reflecting a low false-positive rate. Recall converges near

0.955, indicating the model's effectiveness in detecting relevant objects. The mAP50 metric reaches an average of 0.981, highlighting the model's strong detection accuracy at an IoU threshold of 0.5. Together, these metrics demonstrate the YOLOv8 model's high reliability in object detection tasks.

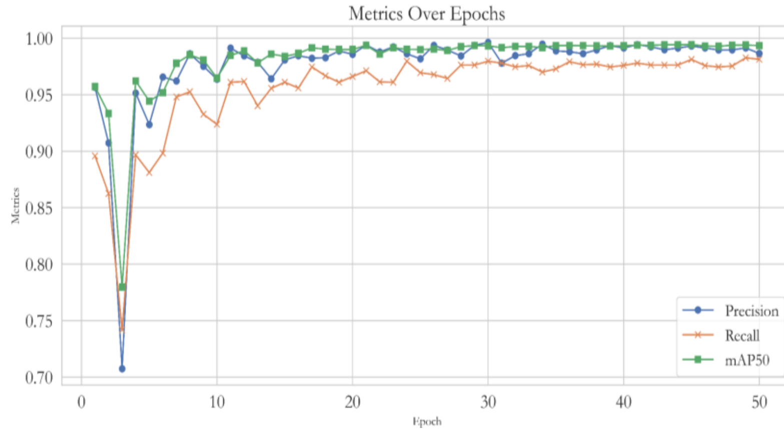

Figure 6. Precision, recall, and mAP50 metrics stabilize at high values, demonstrating consistent, reliable, and well-balanced object detection performance.

Figure 7 illustrates the trends of box loss, classification loss, and distributional focal loss for both training and validation datasets. The gradual decrease in training losses indicates the model's ability to effectively learn and adapt to the dataset. Although validation losses are slightly higher, suggesting minor generalization challenges, the close alignment of the training and validation curves indicates minimal overfitting. This demonstrates that the model successfully captures essential features while maintaining generalization.

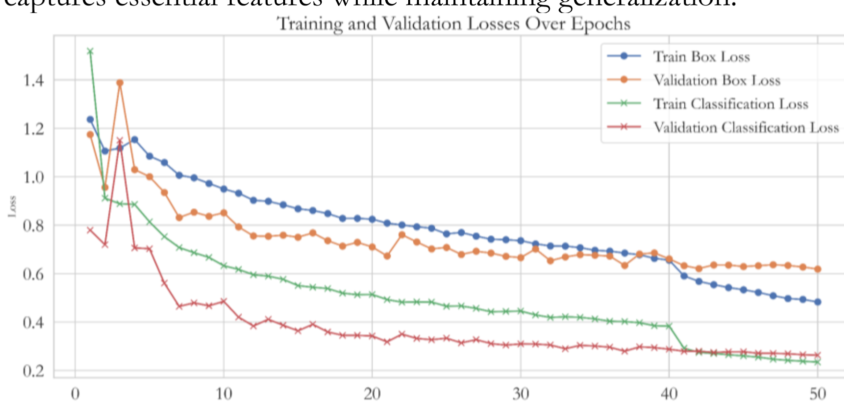

Figure 7 Training and validation losses show a steady decline, indicating effective learning with minimal overfitting, despite slightly higher validation losses

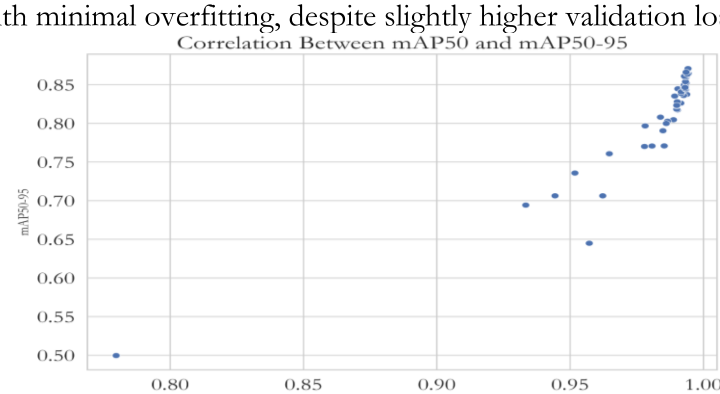

Figure 8. A clear linear relationship between mAP50 and mAP50-95 indicates consistent detection accuracy across different IoU thresholds.

The scatterplot in Figure 8 shows a strong linear correlation (≈ 0.98) between mAP50 and mAP50-95. This indicates that improving detection accuracy at a 0.5 IoU threshold also enhances performance across stricter thresholds. The near-linear pattern highlights the YOLO model's robustness in handling varying levels of localization strictness. Performance metrics are calculated after each training epoch to assess the model's detection ability. Key metrics include:

- **Precision**: The proportion of true positives out of all predicted positives,
- **Recall**: The proportion of true positives out of all actual positives, and
- **mAP50**: Mean average precision at a 0.5 IoU threshold.

The mAP50-95 metric provides a broader evaluation by measuring performance across a range of IoU thresholds (from 0.5 to 0.95). Additionally, validation losses, computed using unseen data, help assess how well the trained model can generalize to new data.

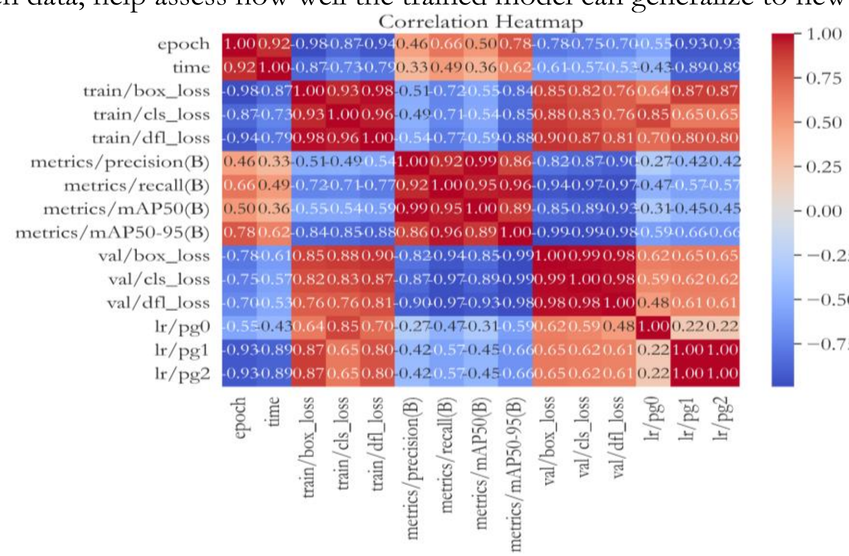

Figure 9. The correlation heatmap highlights key metric interdependencies, showing how reduced losses improve detection accuracy and reflecting the optimization behavior of learning rates.

The heatmap in Figure 9 reveals relationships between key metrics, losses, and learning rates. It shows strong negative correlations between box loss and mAP50/mAP50-95, meaning that lower localization errors lead to higher detection accuracy. Positive correlations between precision and recall demonstrate their interdependence, where enhancing one tends to improve the other. Learning rate parameters (e.g., lr/pg0, lr/pg1, and lr/pg2) exhibit weak direct correlations with precision and recall, indicating that they primarily optimize model weights rather than directly affecting evaluation metrics. Overall, the heatmap provides valuable insights into the optimization dynamics and performance interdependencies of the model. The results confirm the YOLO model's high precision and robustness. The strong correlation between mAP50 and mAP50-95 highlights the model's ability to maintain consistent performance across various IoU thresholds, making it well-suited for applications requiring high detection accuracy. Deviations in validation losses suggest potential areas for further improvement, such as applying advanced augmentation techniques or fine-tuning hyperparameters. These findings emphasize the model's applicability to real-world scenarios where accurate and reliable object detection is critical.

Table 2 presents an analysis of training and validation losses alongside performance metrics, demonstrating the model's effectiveness in object detection. Training box loss steadily decreased across epochs, with validation box loss following a similar downward trend. Slightly higher validation loss values indicate good generalization with minimal overfitting. Classification losses for both training and validation also decreased consistently, reflecting the model's growing accuracy in object classification. Performance metrics further underscore the model's robustness, with precision averaging 0.976, recall at 0.955, and mAP50 reaching an impressive 0.981, which reflects excellent detection accuracy at an IoU threshold of 0.5. Additionally, the mAP50-95 average of 0.812 demonstrates reliable performance across varying IoU thresholds. Gradual reductions in learning rates across parameter groups (lr/pg0, lr/pg1, lr/pg2) contributed to smooth convergence, preventing loss oscillations and ensuring stable training progress.

Discussion:

The findings of this study highlight significant advancements in improving model generalization and real-world applicability through strategic data augmentation and dataset diversity. By fine-tuning YOLOv8 using a combination of self-captured images and publicly available datasets, the model demonstrated enhanced robustness in detecting cow faces under varied lighting, angles, and background conditions. This is a notable improvement compared to earlier object detection frameworks, such as SSD, Faster R-CNN, and YOLOv4, which have been reported to struggle with overfitting when trained on limited datasets. The integration of COCO-pretrained weights further improved YOLOv8's ability to generalize to complex environments, making it more effective and scalable for large-scale livestock monitoring applications. YOLOv8 was chosen over conventional methods like R-CNN and older YOLO versions due to its superior speed and efficiency. Its enhanced backbone structure and anchor-free detection improve object localization, especially in complex environments. Compared to earlier YOLO versions, YOLOv8 offers higher precision, better recall, and faster real-time performance.

Self-captured images alone are insufficient for generalization due to the similarity in environmental conditions. To reduce overfitting and enhance generalization, publicly available data was incorporated into the dataset. Fine-tuning further improved performance by enabling the model to learn dataset-specific features, thereby reducing false positives and improving the accurate detection of cow faces amid complex background elements. YOLOv8 is also more robust than older versions, effectively handling variations in lighting, occlusion, and pose. Fine-tuned models generally perform better in real-world applications compared to generic pretrained models, making YOLOv8 a suitable choice for practical deployment.

In addition to generalization, YOLOv8's comparative performance offers practical implications for precision livestock management. The model's anchor-free detection mechanism and advanced C2f module significantly enhance object localization, resulting in reduced computational overhead, faster inference speed, and improved accuracy. These advantages not only make YOLOv8 suitable for deployment on high-performance GPUs but also adaptable for edge devices, which are increasingly used in smart farms. By achieving high precision, recall, and mAP scores, the proposed approach enhances the accuracy of cow face detection, which can help farmers manage animal health, monitor feeding patterns, and resolve disputes in mixed-herd environments. While the model shows promise, future refinements could further enhance its practical utility and societal impact. Exploring multi-class detection to differentiate between various cattle breeds or integrating temporal analysis to track individual cows over time could offer additional functionalities. Furthermore, incorporating this model into broader Internet of Things (IoT)-based farm management systems could enable real-time livestock monitoring and automated decision-making. Ethical considerations, including animal stress from continuous surveillance and privacy concerns regarding farm data collection, should also be addressed in future implementations. These enhancements would pave the way for a fully autonomous, scalable livestock monitoring system, contributing to increased productivity, reduced labor requirements, and improved animal welfare.

Table 2. shows a detailed breakdown of training and validation losses, highlighting steady improvements in object localization and classification accuracy. Key metrics—precision (0.976), recall (0.955), and mAP50 (0.981)—demonstrate the model's strong detection capability. The gradual learning rate reductions further facilitated smooth convergence and stable training progress.

epoch	train/box_loss	train/cls_loss	metrics/precision (B)	metrics/recall (B)	metrics/ mAP50(B)	metrics/mAP 50-95(B)	val/box_loss	lr/pg0	lr/pg1	lr/pg2
1	1.237	1.519	0.956	0.89581	0.95719	0.64492	1.1744	0.0700	0.003324	0.003324
2	1.106	0.911	0.907	0.86255	0.93337	0.69425	0.9560	0.0399	0.006525	0.006525
3	1.118	0.888	0.707	0.74189	0.77968	0.4997	1.3886	0.0096	0.009595	0.009595
4	1.153	0.885	0.951	0.89676	0.96222	0.70639	1.0297	0.0094	0.009406	0.009406
48	0.497	0.242	0.989	0.97536	0.99387	0.86499	0.6339	0.0006	0.000694	0.000694
49	0.493	0.238	0.991	0.98266	0.99399	0.86363	0.6274	0.0004	0.000496	0.000496
50	0.483	0.235	0.986	0.98136	0.9935	0.86628	0.6191	0.0002	0.000298	0.000298

Conclusion.

The implementation of YOLOv8 for face detection marks a key step toward accurate cow recognition using facial features. Similar to human face recognition, cow face detection holds significant potential for identifying individual animals. By fine-tuning the model, raw background information is filtered out, allowing the system to focus solely on the face, which enhances classification accuracy. This study represents the initial step toward developing a precise cow identification system aimed at tracking and monitoring individual animal performance. A transfer learning-based approach was successfully implemented using YOLOv8 to detect cow faces in diverse farm environments. The model was fine-tuned with a carefully curated dataset that combined frames from video footage with publicly available annotated images. The system demonstrated outstanding performance, achieving a precision of 0.976, a recall of 0.955, and an mAP50 of 0.981, indicating its high accuracy in detecting cow faces. The model's consistent performance improvement across epochs, along with reduced box losses, underscores its robustness. Data augmentation and fine-tuning on a diverse dataset significantly enhanced generalization, while the use of early stopping minimized the risk of overfitting.

Additionally, the strong correlation between mAP50 and mAP50-95 highlighted the model's reliability across varying object localization thresholds. These findings confirm YOLOv8 as an effective tool for real-time face detection in livestock localization and monitoring applications. This approach not only establishes a foundation for accurate cow face detection but also paves the way for future advancements in classification tasks. Future work may focus on refining the model further and integrating it with additional technologies to enhance cow face recognition and improve individual cow identification in practical settings.

Acknowledgement. The authors would like to thank the administration of Intelligent Systems Laboratory, Department of Electrical Engineering, University of Gujrat for providing an opportunity to conduct research and perform experiments with the available laboratory facilities.

Author's Contribution. Umair Ali wrote the initial draft of the article. Wasif Muhammad supervised and drafted the final version of the manuscript.

Conflict of interest. The authors have no conflict of interest for publishing this manuscript in IJIST.

References:

[1] P. L. P. J. Zarco-Tejada, N. Hubbard, "Precision Agriculture: An Opportunity for EU-Farmers – Potential Support with the CAP 2014-2020," Think Tank. Accessed: Mar. 13, 2025. [Online]. Available: https://www.europarl.europa.eu/thinktank/en/document/IPOL-AGRI_NT(2014)529049

[2] S. K. Seelan, S. Laguette, G. M. Casady, and G. A. Seielstad, "Remote sensing applications for precision agriculture: A learning community approach," *Remote Sens. Environ.*, vol. 88, no. 1–2, pp. 157–169, 2003, doi: https://doi.org/10.1016/j.rse.2003.04.007.

[3] P. G. António Monteiro, Sérgio Santos, "Precision Agriculture for Crop and Livestock Farming—Brief Review," *Animals*, vol. 11, no. 8, p. 2345, 2021, doi: https://doi.org/10.3390/ani11082345.

[4] Ali Ismail Awad, "From classical methods to animal biometrics: A review on cattle identification and tracking," *Comput. Electron. Agric.*, vol. 1123, pp. 423–435, 216AD, doi: https://doi.org/10.1016/j.compag.2016.03.014.

[5] L. X. Liu Bo, Liu Yuefeng, Bao Xiang, Wang Yue, Liu Haofeng, "Research on Dairy Cow Identification Methods in Dairy Farm," *Indian J. Anim. Res.*, vol. 52, no. 12, pp. 1733–1739, 2023, doi: 10.18805/IJAR.BF-1660.

[6] A. M. Johnston and D. S. Edwards, "Welfare implications of identification of cattle by ear tags," *Vet. Rec.*, vol. 138, no. 25, pp. 612–614, Jun. 1996, doi: 10.1136/VR.138.25.612.
[7] D D Wardrope, "Problems with the use of ear tags in cattle - PubMed," Vet Rec. Accessed: Mar. 13, 2025. [Online]. Available: https://pubmed.ncbi.nlm.nih.gov/8966978/
[8] V. I. K. and K. M. D. A N Ruchay, K A Dorofeev, V V Kalschikov, "A depth camera-based system for automatic measurement of live cattle body parameters," *IOP Conf. Ser. Earth Environ. Sci.*, vol. 341, 2019, doi: 10.1088/1755-1315/341/1/012148.
[9] M. S. V. Maria Vlad, Roxana Andreea Parvulet, "A Survey of Livestock Identification Systems," Proceedings of the 13th WSEAS International Conference on Automation and Information (ICAI12). Accessed: Mar. 13, 2025. [Online]. Available: https://www.researchgate.net/publication/265227039_A_Survey_of_Livestock_Identification_Systems
[10] P. Thanapol, K. Lavangnananda, P. Bouvry, F. Pinel, and F. Leprevost, "Reducing Overfitting and Improving Generalization in Training Convolutional Neural Network (CNN) under Limited Sample Sizes in Image Recognition," *InCIT 2020 - 5th Int. Conf. Inf. Technol.*, pp. 300–305, Oct. 2020, doi: 10.1109/INCIT50588.2020.9310787.
[11] S. M. A. Imran Khan, Chen Tianhua, Khan Malak Abid Ali, "Object Recognition in Different Lighting Conditions at Various Angles by Deep Learning Method," arXiv preprint arXiv:2210.09618. Accessed: Mar. 13, 2025. [Online]. Available: https://www.researchgate.net/publication/357702329_Object_Recognition_in_Different_Lighting_Conditions_at_Various_Angles_by_Deep_Learning_Method
[12] X. Lei, X. Wen, and Z. Li, "A multi-target cow face detection model in complex scenes," *Vis. Comput.*, vol. 40, no. 12, pp. 9155–9176, Dec. 2024, doi: 10.1007/S00371-024-03301-W/METRICS.
[13] A. Redmon, J., Divvala, S., Girshick, R., Farhadi, "You only look once: Unified, real-time object detection," *2016 IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, doi: 10.1109/CVPR.2016.91.
[14] J. Solawetz and Francesco, "What is YOLOv8? A Complete Guide," Yolov8. Accessed: Mar. 13, 2025. [Online]. Available: https://blog.roboflow.com/what-is-yolov8/