

# DEAL-YOLO: DRONE-BASED EFFICIENT ANIMAL LOCALIZATION USING YOLO

Aditya Prashant Naidu<sup>1\*</sup>, Hem Gosalia<sup>2\*</sup>, Ishaan Gakhar<sup>3\*</sup>, Shaurya Singh Rathore<sup>4\*</sup>  
Krish Didwania<sup>1</sup>, Ujjwal Verma<sup>5</sup>

<sup>1</sup>Dept. of Computer Science and Engineering, <sup>2</sup>Dept. of Mechatronics Engineering

<sup>3</sup>Dept. of Information and Communication Technology, <sup>4</sup>Dept. of Data Science and Computer Applications

<sup>5</sup>Dept. of Electronics and Communication Engineering

Manipal Institute of Technology, Manipal Academy of Higher Education  
Manipal, Karnataka, India

{adityanaidu2004, hem.gosalia3, ishaangakhar04, shauryarathore121}@gmail.com  
krishdidwania0674@gmail.com, ujjwal.verma@manipal.edu \*

## ABSTRACT

Although advances in deep learning and aerial surveillance technology are improving wildlife conservation efforts, complex and erratic environmental conditions still pose a problem, requiring innovative solutions for cost-effective small animal detection. This work introduces DEAL-YOLO, a novel approach that improves small object detection in Unmanned Aerial Vehicle (UAV) images by using multi-objective loss functions like Wise IoU (WIoU) and Normalized Wasserstein Distance (NWD), which prioritize pixels near the centre of the bounding box, ensuring smoother localization and reducing abrupt deviations. Additionally, the model is optimized through efficient feature extraction with Linear Deformable (LD) convolutions, enhancing accuracy while maintaining computational efficiency. The Scaled Sequence Feature Fusion (SSFF) module enhances object detection by effectively capturing inter-scale relationships, improving feature representation, and boosting metrics through optimized multiscale fusion. Comparisons with baseline models reveal high efficacy with up to 69.5% fewer parameters compared to vanilla Yolov8-N, highlighting the robustness of the proposed modifications. Through this approach, our paper aims to facilitate the detection of endangered species, animal population analysis, habitat monitoring, biodiversity research, and various other applications that enrich wildlife conservation efforts. DEAL-YOLO employs a two-stage inference paradigm for object detection, refining selected regions to improve localization and confidence. This approach enhances performance, especially for small instances with low objectness scores.

## 1 INTRODUCTION AND PREVIOUS WORK

Wildlife object detection has proven to be essential for all aspects related to biodiversity conservation. (Chalmers et al., 2021; Delplanque et al., 2022; Peng et al., 2020). Accurate identification and tracking of animal species from aerial imagery allows the evaluation of population trends, habitat changes, and effective protection strategies. Traditional monitoring techniques such as ground surveys and camera trapping, can be hindered by their high costs and potential human biases (et al.). To this end, UAVs present a more efficient alternative, providing cost-effective, high-resolution aerial data with minimal human involvement. Recent advancements in deep learning have significantly enhanced the automation and quality of wildlife detection through Convolutional Neural Networks (CNNs) and object detection models (Axford et al., 2024). However, further improvements are needed to enhance performance on object detection while ensuring computational efficiency, particularly for deployment on UAVs.

---

\*Equal Contribution

Modern object detection models, particularly the You Only Look Once (YOLO) family (Redmon et al., 2016; Bochkovskiy et al., 2020) and Faster R-CNN (Ren et al., 2016), have demonstrated superior accuracy in detecting and classifying objects in complex environments. However, wildlife detection presents unique challenges, particularly in UAV-based imagery. Small animal targets often occupy only a few pixels, making distinguishing them from the background difficult. In addition, occlusions, overlapped species, variations in lighting conditions, and environmental interference further complicate the detection process (Eikelboom et al., 2019). Recent advances in small object detection have introduced various techniques to improve accuracy, yet challenges persist in drone-based wildlife detection. RRNet (Chen et al., 2019) employed AdaResampling for realistic augmentation but struggled with segmentation challenges in natural environments. RFLA (Xu et al., 2022) assigned labels via Gaussian receptive fields but faced limitations with irregularly shaped animals. The Focus & Detect framework (Koyun et al., 2022) enhanced small-object detection through high-resolution cropping but required extensive manual annotations. Cross-layer attention mechanisms (Li et al., 2021) amplified small object features but increased computational costs, while SSPNet (Hong et al., 2022) fused multiscale features but diluted fine details. Wildlife detection models, such as those based on YOLOv5, YOLOv8, and Faster R-CNN (Ocholla et al., 2024), performed well on structured targets like livestock but struggled with camouflage and scale variations. CNN-based approaches for satellite imagery (Bowler et al., 2020) and Faster R-CNN with HRNet (Ma et al., 2022) improved small target recognition but suffered from anchor box limitations and false positives due to vegetation noise. Similarly, YOLOv6L (Cusick et al., 2024) detected static nests but was sensitive to resolution changes. Efficient object detection models have also been explored, with modifications to YOLOv5 (Jung & Choi, 2022) improving efficiency at the cost of fine-grained spatial details. UFPMP-Det (Huang et al., 2022) leveraged attention mechanisms but introduced computational overhead, while Drone-DETR (Kong et al., 2024) relied on large datasets and exhibited slow convergence. Efficient YOLOv7-Drone (Fu et al., 2023) optimized UAV detection but struggled with camouflaged wildlife due to its reliance on accurate mask generation.

Despite these advancements, achieving robust and efficient wildlife detection for previous works is challenging for drone imagery due to limitations in feature resolution, fixed anchor boxes, and difficulty in distinguishing fine details amidst background noise. The main contributions of this work include:

- **Optimization and Restructuring of YOLOv8:** Modifications introduced to YOLOv8 such as efficient convolution modules and an optimized downsampling strategy, significantly reduce computational complexity while preserving high performance.
- **State-of-the-Art Performance at lower computational load:** Achieved superior detection accuracy with up to 69.6% reduction in trainable parameters, effectively optimizing both efficiency and performance, thus showcasing applicability in real-world use cases.
- **Two-Stage Inference Strategy:** Introduced a novel and adaptive two-stage Region of Interest (RoI) based inference approach that enhances detection performance by refining bounding box predictions in ambiguous environments requiring fine-grained differentiation. This results in 4% increase in Precision and 4.2% increase in Recall on average.

## 2 PROPOSED METHODOLOGY

The proposed methodology utilizes a combination of advanced loss functions, architectural modifications, and inference strategies to enhance object detection performance in UAV imagery. In particular, DEAL-YOLO integrates the **Normalized Wasserstein Distance** (Wang et al., 2022b) to model bounding boxes as 2D Gaussian distributions, measuring the similarity between transformed predicted boxes and ground truth labels. By assigning greater importance to pixels near the center, this approach accounts for the smaller size of aerial objects and introduces smoothness to the bounding box deviations, with the Optimal Network theory underpinning the exponential normalization to yield an effective similarity measure. To further mitigate the influence of low-quality examples, the model also incorporates the **Wise IoU** metric (Tong et al., 2023), which minimizes the adverse effects of geometric variations—such as differences in distance and aspect ratio—by penalizing both major and minor misalignments between predicted anchor boxes and target boxes. Its adaptive weighting mechanism is particularly valuable for UAV applications, where altitude variations cause objects to appear at diverse scales, ensuring that smaller objects (often captured at high altitudes) are detected



Table 1: Comparisons with baseline YOLO models on the BuckTales dataset against the proposed approach across various metrics. Suffix 'T' stands for Tiny and 'N' stands for Nano. The '\*\*' represents results with 2-stage inference.

Model	#Params(M)	Precision	Recall	mAP <sub>50</sub>
YOLOv5-N	2.504	46.9	53.1	48.7
YOLOv6-N	4.234	38.7	42.2	42.3
YOLOv8-N	3.006	70.7	41.8	42.8
YOLOv9-T	1.972	59.7	48.8	55.8
YOLOv10-N	2.697	42.0	45.7	46.2
Gold-YOLO	5.610	38.6	75.0	50.7
RT-DETR	41.97	48.1	38.9	35.7
Faster-RCNN	43.060	63.6	75.2	29.7
DEAL-YOLO (Ours)	<b>0.994</b>	75.3	58.2	<b>48.5</b>
DEAL-YOLO (Ours)*	<b>0.994</b>	<b>85.3</b>	<b>87.8</b>	47.6

Table 2: Comparison of SOTA methods against the proposed method on the WAID dataset across various metrics. Not all models have published the mAP<sub>50</sub>, and hence that entry has been left blank. Suffix 'T' stands for Tiny, 'S' stands for Small and 'N' stands for Nano. The '\*\*' represents results with 2-stage inference. '-LD' represents our model with LD convolutions.

Model	#Params(M)	Precision	Recall	mAP <sub>50</sub>
YOLOv7-T	6.000	93.7	92.3	95.23
YOLOv5-S	7.200	<b>96.9</b>	92.9	<b>96.3</b>
ADD-YOLO	1.500	93.0	91.0	95.0
WILD-YOLO	12.380	92.8	91.41	95.0
YOLOv4-S	8.270	38.2	92.7	56.3
MobileNet v2	3.950	40.0	91.5	59.1
YOLOv8-N	3.010	88.5	85.4	89.7
DEAL-YOLO-LD	<b>0.914</b>	91.1	88.9	93.4
DEAL-YOLO	0.994	92.0	88.8	93.3
DEAL-YOLO-LD*	<b>0.914</b>	95.2	94.8	90.5
DEAL-YOLO*	0.994	95.9	<b>95.3</b>	90.8

### 3 EXPERIMENTS AND RESULTS

To validate the proposed methodology, the WAID (Mou et al., 2023) and BuckTales datasets (Naik et al., 2024) have been employed and thorough experimentation has been performed to justify our choice of modules.

As shown in Table 1, we compare various YOLO models—YOLOv6, YOLOv8, YOLOv9, YOLOv10, Gold-YOLO, RT-DETR, Faster R-CNN (Li et al., 2023; Jocher et al., 2023; Wang & Liao, 2024; Wang et al., 2024; 2023; Zhao et al., 2024; Ren et al., 2016)—and DEAL-YOLO. Our model, with 68% fewer parameters, outperforms these baselines by 4.8% on average. This reduction in computational load, coupled with improved accuracy, makes it well-suited for animal detection. Results without 2-stage inference, also in Table 1, highlight the benefits of our approach. The model was trained using the SOAP optimizer (Vyas et al., 2025), ensuring better stability and convergence.

Table 2 further demonstrates our methodology’s competitiveness, achieving **87% fewer** parameters than YOLOv8n while maintaining performance. Compared to YOLOv7-T, YOLOv5-S, ADD-YOLO, WILD-YOLO, YOLOv4-S, and MobileNet v2, DEAL-YOLO LD—leveraging SSFF layers and LD convolutions—excels in detecting small objects, such as animals in the WAID dataset. The 2-stage inference further enhances predictions. Additional experiments and details are mentioned in the appendix.

## REFERENCES

- Daniel Axford, Ferdous Sohel, Mathew A Vanderklift, and Amanda J Hodgson. Collectively advancing deep learning for animal detection in drone imagery: Successes, challenges, and research gaps. *Ecological Informatics*, 83:102842, 2024. ISSN 1574-9541. doi: <https://doi.org/10.1016/j.ecoinf.2024.102842>. URL <https://www.sciencedirect.com/science/article/pii/S1574954124003844>.
- Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. Yolov4: Optimal speed and accuracy of object detection, 2020.
- Ellen Bowler, Peter T. Fretwell, Geoffrey French, and Michal Mackiewicz. Using deep learning to count albatrosses from space: Assessing results in light of ground truth uncertainty. *Remote Sensing*, 12(12), 2020. ISSN 2072-4292. doi: 10.3390/rs12122026. URL <https://www.mdpi.com/2072-4292/12/12/2026>.
- C. Chalmers, P. Fergus, C. Aday Curbelo Montanez, Steven N. Longmore, and Serge A. Wich. Video analysis for the detection of animals using convolutional neural networks and consumer-grade drones. *Journal of Unmanned Vehicle Systems*, 9(2):112–127, 2021. ISSN 2291-3467. doi: <https://doi.org/10.1139/juvs-2020-0018>. URL <https://www.sciencedirect.com/science/article/pii/S2291346721000084>.
- Changrui Chen, Yu Zhang, Qingxuan Lv, Shuo Wei, Xiaorui Wang, Xin Sun, and Junyu Dong. Rrnet: A hybrid detector for object detection in drone-captured images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, Oct 2019.
- Andrew Cusick, Katarzyna Fudala, Piotr Pasza Storożenko, Jędrzej Świeżewski, Joanna Kaleta, W. Chris Oosthuizen, Christian Pfeifer, and Robert Józef Bialik. Using machine learning to count antarctic shag (*leucocarbo bransfieldensis*) nests on images captured by remotely piloted aircraft systems. *Ecological Informatics*, 82:102707, 2024. ISSN 1574-9541. doi: <https://doi.org/10.1016/j.ecoinf.2024.102707>. URL <https://www.sciencedirect.com/science/article/pii/S1574954124002498>.
- Alexandre Delplanque, Samuel Foucher, Philippe Lejeune, Julie Linchant, and Jérôme Théau. Multispecies detection and identification of african mammals in aerial imagery using convolutional neural networks. *Remote Sensing in Ecology and Conservation*, 8(2):166–179, 2022. doi: <https://doi.org/10.1002/rse2.234>. URL <https://zslpublications.onlinelibrary.wiley.com/doi/abs/10.1002/rse2.234>.
- Jan A. J. Eikelboom, W. Daniel Kissling, and C. M. G. Groen. Automated detection of animals in aerial images using deep learning. *Remote Sensing in Ecology and Conservation*, 5(4):456–469, 2019. doi: 10.1002/rse2.123.
- Bruce et al. Large-scale and long-term wildlife research and monitoring using camera traps: a continental synthesis. *Biological Reviews*, n/a(n/a). doi: <https://doi.org/10.1111/brv.13152>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/brv.13152>.
- Xiaofeng Fu, Guoting Wei, Xia Yuan, Yongshun Liang, and Yuming Bo. Efficient yolov7-drone: An enhanced object detection approach for drone aerial imagery. *Drones*, 7(10), 2023. ISSN 2504-446X. doi: 10.3390/drones7100616. URL <https://www.mdpi.com/2504-446X/7/10/616>.
- Mingbo Hong, Shuiwang Li, Yuchao Yang, Feiyu Zhu, Qijun Zhao, and Li Lu. Sspnet: Scale selection pyramid network for tiny person detection from uav images. *IEEE Geoscience and Remote Sensing Letters*, 19:1–5, 2022. doi: 10.1109/LGRS.2021.3103069.
- Yecheng Huang, Jiaxin Chen, and Di Huang. Ufmpm-det: toward accurate and efficient object detection on drone imagery. *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(1):1026–1033, Jun. 2022. doi: 10.1609/aaai.v36i1.19986. URL <https://ojs.aaai.org/index.php/AAAI/article/view/19986>.
- Glenn Jocher, Ayush Chaurasia, and Jing Qiu. Ultralytics yolov8, 2023. URL <https://github.com/ultralytics/ultralytics>.

- Hyun-Ki Jung and Gi-Sang Choi. Improved yolov5: Efficient object detection using drone images under various conditions. *Applied Sciences*, 12(14), 2022. ISSN 2076-3417. doi: 10.3390/app12147255. URL <https://www.mdpi.com/2076-3417/12/14/7255>.
- Ming Kang, Chee-Ming Ting, Fung Fung Ting, and Raphaël C.-W. Phan. Asf-yolo: A novel yolo model with attentional scale sequence fusion for cell instance segmentation. *Image and Vision Computing*, 147:105057, July 2024. ISSN 0262-8856. doi: 10.1016/j.imavis.2024.105057. URL <http://dx.doi.org/10.1016/j.imavis.2024.105057>.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017. URL <https://arxiv.org/abs/1412.6980>.
- Yaning Kong, Xiangfeng Shang, and Shijie Jia. Drone-detr: Efficient small object detection for remote sensing image using enhanced rt-detr model. *Sensors*, 24(17), 2024. ISSN 1424-8220. doi: 10.3390/s24175496. URL <https://www.mdpi.com/1424-8220/24/17/5496>.
- Onur Can Koyun, Reyhan Kevser Keser, İbrahim Batuhan Akkaya, and Behçet Uğur Töreyn. Focus-and-detect: A small object detection framework for aerial images. *Signal Processing: Image Communication*, 104:116675, 2022. ISSN 0923-5965. doi: <https://doi.org/10.1016/j.image.2022.116675>. URL <https://www.sciencedirect.com/science/article/pii/S0923596522000273>.
- Chuyi Li, Lulu Li, Yifei Geng, Hongliang Jiang, Meng Cheng, Bo Zhang, Zaidan Ke, Xiaoming Xu, and Xiangxiang Chu. Yolov6 v3.0: A full-scale reloading, 2023.
- Yangyang Li, Qin Huang, Xuan Pei, Yanqiao Chen, Licheng Jiao, and Ronghua Shang. Cross-layer attention network for small object detection in remote sensing imagery. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 14:2148–2161, 2021. doi: 10.1109/JSTARS.2020.3046482.
- Jiarong Ma, Zhuowei Hu, Quanqin Shao, Yongcai Wang, Yanqiong Zhou, Jiayan Liu, and Shuchao Liu. Detection of large herbivores in uav images: A new method for small target recognition in large-scale images. *Diversity*, 14(8), 2022. ISSN 1424-2818. doi: 10.3390/d14080624. URL <https://www.mdpi.com/1424-2818/14/8/624>.
- Chao Mou, Tengfei Liu, Chengcheng Zhu, and Xiaohui Cui. Waid: A large-scale dataset for wildlife detection with drones. *Applied Sciences*, 13(18), 2023. ISSN 2076-3417. doi: 10.3390/app131810397. URL <https://www.mdpi.com/2076-3417/13/18/10397>.
- Hemal Naik, Junran Yang, Dipin Das, Margaret C Crofoot, Akanksha Rathore, and Vivek Hari Sridhar. Bucktales : A multi-uav dataset for multi-object tracking and re-identification of wild antelopes, 2024. URL <https://arxiv.org/abs/2411.06896>.
- Ian A. Ocholla, Petri Pellikka, Faith Karanja, Ilja Vuorinne, Tuomas Väisänen, Mark Boitt, and Janne Heiskanen. Livestock detection and counting in kenyan rangelands using aerial imagery and deep learning techniques. *Remote Sensing*, 16(16), 2024. ISSN 2072-4292. doi: 10.3390/rs16162929. URL <https://www.mdpi.com/2072-4292/16/16/2929>.
- Jinbang Peng, Dongliang Wang, Xiaohan Liao, Quanqin Shao, Zhigang Sun, Huanyin Yue, and Huping Ye. Wild animal survey using uas imagery and deep learning: modified faster r-cnn for kiang detection in tibetan plateau. *ISPRS Journal of Photogrammetry and Remote Sensing*, 169:364–376, 2020. ISSN 0924-2716. doi: <https://doi.org/10.1016/j.isprsjprs.2020.08.026>. URL <https://www.sciencedirect.com/science/article/pii/S0924271620302409>.
- Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 779–788, 2016.
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pp. 91–99, 2016.
- Zanjia Tong, Yuhang Chen, Zewei Xu, and Rong Yu. Wise-iou: Bounding box regression loss with dynamic focusing mechanism, 2023. URL <https://arxiv.org/abs/2301.10051>.

- Nikhil Vyas, Depen Morwani, Rosie Zhao, Mujin Kwun, Itai Shapira, David Brandfonbrener, Lucas Janson, and Sham Kakade. Soap: Improving and stabilizing shampoo using adam, 2025. URL <https://arxiv.org/abs/2409.11321>.
- Ao Wang, Hui Chen, Lihao Liu, and et al. Yolov10: Real-time end-to-end object detection. *arXiv preprint arXiv:2405.14458*, 2024.
- Chengcheng Wang, Wei He, Ying Nie, Jianyuan Guo, Chuanjian Liu, Kai Han, and Yunhe Wang. Gold-yolo: Efficient object detector via gather-and-distribute mechanism, 2023. URL <https://arxiv.org/abs/2309.11331>.
- Chien-Yao Wang and Hong-Yuan Mark Liao. Yolov9: Learning what you want to learn using programmable gradient information. 2024.
- Chien-Yao Wang, Alexey Bochkovskiy, and Hong-Yuan Mark Liao. Yolov7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors, 2022a. URL <https://arxiv.org/abs/2207.02696>.
- Jinwang Wang, Chang Xu, Wen Yang, and Lei Yu. A normalized gaussian wasserstein distance for tiny object detection, 2022b. URL <https://arxiv.org/abs/2110.13389>.
- Chang Xu, Jinwang Wang, Wen Yang, Huai Yu, Lei Yu, and Gui-Song Xia. Rfla: Gaussian receptive field based label assignment for tiny object detection. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner (eds.), *Computer Vision – ECCV 2022*, pp. 526–543, Cham, 2022. Springer Nature Switzerland. ISBN 978-3-031-20077-9.
- Xin Zhang, Yingze Song, Tingting Song, Degang Yang, Yichen Ye, Jie Zhou, and Liming Zhang. Ldconv: Linear deformable convolution for improving convolutional neural networks. *Image and Vision Computing*, 149:105190, 2024. ISSN 0262-8856. doi: <https://doi.org/10.1016/j.imavis.2024.105190>. URL <https://www.sciencedirect.com/science/article/pii/S0262885624002956>.
- Yian Zhao, Wenyu Lv, Shangliang Xu, Jinman Wei, Guanzhong Wang, Qingqing Dang, Yi Liu, and Jie Chen. Detrs beat yolos on real-time object detection, 2024. URL <https://arxiv.org/abs/2304.08069>.

## A DETAILED METHODOLOGY

The Scaled Sequential Feature Fusion (SSFF) (Kang et al., 2024) block enhances multi-scale feature representation by refining feature maps (P3, P4, P5) sequentially using Gaussian smoothing before fusion. Each feature map  $f(i, j)$  is convolved with a Gaussian kernel  $G_\sigma(x, y)$ , defined as:

$$F_\sigma(i, j) = \sum_u \sum_v f(i - u, j - v) \times G_\sigma(u, v) \quad (1)$$

where  $G_\sigma(x, y)$  is given by:

$$G_\sigma(x, y) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}} \quad (2)$$

This progressively smooths the feature maps with increasing standard deviation  $\sigma$ , ensuring robust feature refinement. The smoothed feature maps are then fused sequentially, allowing finer-scale details from P3 to progressively enhance coarser features in P4 and P5, preserving spatial relationships and improving object detection across scales.

We use the Normalized Wasserstein Distance to achieve smoothness of the bounding box deviations according to the formula

$$NWD(N_a, N_b) = \exp\left(-\frac{\sqrt{W_2^2(N_a, N_b)}}{C}\right), \quad (3)$$

Where  $N_a$  and  $N_b$  represent two Gaussian distributions. The term  $W_2^2(N_a, N_b)$  denotes the squared 2-Wasserstein distance between these distributions, measuring the optimal transport cost between them. The constant  $C$  serves as a normalization constant.

This is combined with the Wise IoU metric to minimize the adverse effects of geometric variation according to

$$\mathcal{L}_{WIoU} = \mathcal{R}_{WIoU} \mathcal{L}_{IoU} \quad (4)$$

$$\mathcal{R}_{WIoU} = \exp\left(\frac{(x - x_{gt})^2 + (y - y_{gt})^2}{(W_g^2 + H_g^2)^*}\right) \quad (5)$$

In the equations,  $\mathcal{L}_{WIoU}$  denotes the weighted IoU loss (computed as the product of the standard IoU loss  $\mathcal{L}_{IoU}$  and the scaling factor  $\mathcal{R}_{WIoU}$ ).  $\mathcal{R}_{WIoU}$  scales the loss based on the squared Euclidean distance between the predicted box centre  $(x, y)$  and the ground truth centre  $(x_{gt}, y_{gt})$ , with  $(W_g^2 + H_g^2)^*$  serving as a normalization term.

LDCConv (Linear Deformable Convolution) introduced by Zhang et al. (2024) is an innovative convolutional operation that facilitates arbitrarily sampled shapes and accommodates a flexible number of parameters, distinguishing it from conventional fixed-grid convolutions. This approach generates initial sampled positions and learns offsets to adjust the receptive field dynamically. As a result, it enables more efficient feature extraction with linear parameter growth. This adaptability allows LDCConv to cater to various target shapes while optimizing computational efficiency.

## B ABLATION STUDY

In this study, the impact of individual components and their combination is presented for each dataset. As seen in Table 3, the results of the SOAP optimizer are seen in Row 2. The results of Row 1 are Vanilla YOLOv8 (Jocher et al., 2023) with the Adam Optimizer (Kingma & Ba, 2017). Inclusion of the SOAP optimizer results in a 7% increase while integrating the SSFF module along with WIoU and NWD Loss demonstrates an increase in performance of 6.625%. Finally, changing P5 to P2 causes a 66.93% reduction in trainable parameters while having a negligible effect on the quantitative metrics. These results showcase the effectiveness and fine-grained choice of our modules and the applicability of our approach in real-world scenarios.

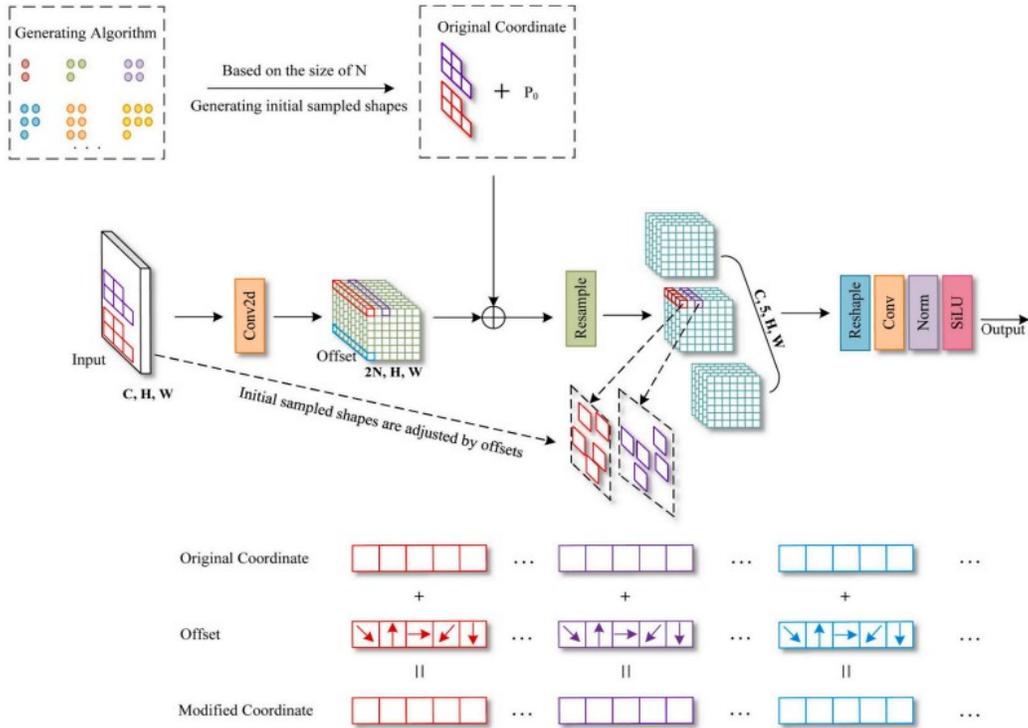


Figure 2: Schematic overview of the structure of LDConv.(Zhang et al., 2024) The initial sampled coordinates are assigned to a convolution of arbitrary size, and the sample shape is adjusted using learnable offsets. This process modifies the original sampled shape at each position through resampling.

Table 3: Ablation study on the impact of individual proposed changes in YOLOv8-N for the BuckTales Dataset. 'WIoU + NWD' represents the integration of WIoU and NWD losses.

SOAP	SSFF	WIoU+NWD	P2 for P5	#Params(M)	Precision	Recall	mAP <sub>50</sub>	AP <sub>50-95</sub>
✓				3.006	92.2	87.7	93.1	61.0
✓	✓			2.490	91.4	89.0	93.6	<b>68.0</b>
✓	✓	✓		2.490	<b>94.4</b>	<b>89.4</b>	<b>94.2</b>	60.9
✓	✓	✓	✓	0.994	92.0	88.8	93.3	60.9
✓	✓	✓	✓	<b>0.914</b>	91.1	88.9	93.4	60.3

The decision to resize images to 640 at inference was based on practical considerations and the need for consistency with the patched dataset. Since the training was conducted on patched images ( $608 \times 513$ ) resized to 1280, it was important to ensure that inference-time resizing did not introduce distortions that could impact model performance.

Resizing test images to 640 provides a reasonable balance between preserving spatial information and maintaining consistency with the training data. Since the patched images used during training are relatively small ( $608 \times 513$ ), a test resolution of 640 minimizes excessive resizing, helping retain object details and prevent artifacts. As observed in Table 5, models evaluated on patched images (resized to 1280) achieve significantly better performance compared to those tested on unpatched images with larger resizing scales.

Additionally, using 640 aligns closely with standard YOLO input sizes, ensuring compatibility with pre-trained backbone architectures while keeping computational requirements manageable. Since

Table 4: Ablation study on the impact of individual proposed changes for WAID Dataset. 'WIoU + NWD' represents the integration of WIoU and NWD losses.

SOAP	SSFF	WIoU+NWD	P2 for P5	#Params(M)	Precision	Recall	mAP <sub>50</sub>	AP <sub>50-95</sub>
				3.006	70.7	41.8	42.8	26.0
✓				3.006	62.4	57.1	56.6	33.2
✓	✓			2.490	65.4	53.7	58.0	33.0
✓	✓	✓		2.490	<b>73.4</b>	<b>60.7</b>	<b>63.1</b>	<b>39.4</b>
✓	✓	✓	✓	<b>0.994</b>	67.4	57.5	63.0	37.6

Table 5: Ablation study on the effect of using the patched/unpatched versions of the BuckTales dataset, as well as different image resizing during inference. 'Patched/1280' means patched dataset and images were resized to 1280 during inference. Suffix 'T' stands for Tiny and 'N' stands for Nano.

BuckTales	Model	#Params(M)	Precision	Recall	mAP <sub>50</sub>	mAP <sub>50-95</sub>
Patched/1280	YOLOv5-N	2.504	58.6	63.0	<b>65.9</b>	35.0
	YOLOv6-N	4.234	54.9	61.6	61.6	34.6
	YOLOv8-N	3.006	57.5	54.5	59.7	34.7
	YOLOv9-T	1.972	57.3	<b>66.0</b>	65.1	<b>40.1</b>
	YOLOv10-N	2.697	<b>63.7</b>	55.1	57.6	33.5
	DEAL-YOLO (Ours)	<b>0.994</b>	57.5	61.0	62.2	38.3
Unpatched/2560	YOLOv5-N	2.504	9.00	10.00	4.58	1.45
	YOLOv6-N	4.234	7.87	8.93	3.68	1.15
	YOLOv8-N	3.006	10.30	3.46	<b>5.73</b>	<b>1.59</b>
	YOLOv9-T	1.972	9.39	11.10	5.21	1.69
	YOLOv10-N	2.697	5.41	<b>14.20</b>	3.36	1.16
	DEAL-YOLO (Ours)	<b>0.994</b>	<b>32.50</b>	7.34	7.82	2.72
Unpatched/3840	YOLOv5-N	2.504	28.7	28.7	28.4	11.4
	YOLOv6-N	4.234	47.8	29.0	28.9	11.6
	YOLOv8-N	3.006	<b>51.6</b>	27.6	30.7	13.8
	YOLOv9-T	1.972	48.0	35.3	36.8	15.4
	YOLOv10-N	2.697	27.2	26.7	25.2	9.8
	DEAL-YOLO (Ours)	<b>0.994</b>	40.0	<b>39.1</b>	<b>38.5</b>	<b>19.0</b>

UAV imagery contains small objects across large backgrounds, aggressive resizing (either upscaling or downscaling) could lead to a loss of fine details or unnecessary blurring.

While the exact impact of different test-time resizing strategies would require further empirical validation, the choice of 640 appears to be a well-reasoned approach that maintains consistency with training, minimizes distortions, and balances computational efficiency without introducing significant domain shifts.

## C QUALITATIVE RESULTS

In this section, we analyze the qualitative performance of our model on the WAID and BuckTales datasets. By visualizing the predictions as shown in ??, we assess how well the model localizes animals and generalizes beyond the provided annotations. The following observations highlight key aspects of the model's effectiveness in real-world scenarios.

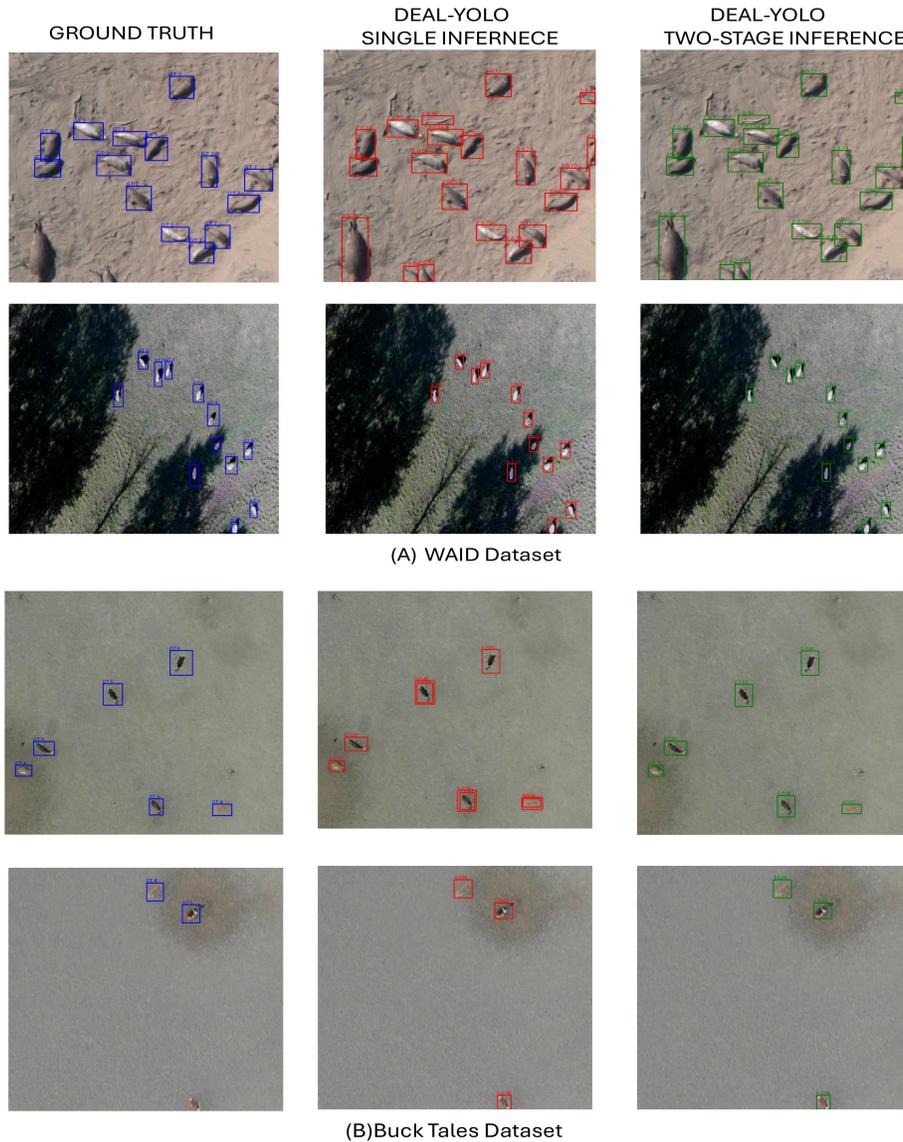


Figure 3: Qualitative results on the WAID and BuckTales datasets. Ground truth annotations are shown in blue, single-stage inference predictions in red, and two-stage inference predictions in green. The left column represents the Ground Truth bounding boxes, the middle column represents DEAL-YOLO with standard inference and the right column represents results of two-stage inference.

The model’s predictions exhibit a closer and more precise alignment with the detected animals than the ground truth annotations, demonstrating superior localization. Notably, the model also identifies animals that are missing from the ground truth labels, highlighting its ability to generalize beyond the provided annotations. The use of two-stage inference further enhances detection performance by boosting confidence scores and effectively resolving overlapping bounding boxes. This approach ensures more precise predictions and better differentiation of multiple animals within the same frame, ultimately improving overall detection accuracy.

During our analysis, we identified a potential issue in both datasets. DEAL-YOLO, when using standard inference, achieves higher confidence scores than vanilla YOLOv8. As shown in Fig. ??, certain instances in both datasets lack highly accurate bounding boxes. Visualizing DEAL-YOLO’s results revealed that two-stage inference produced bounding boxes that were more compact and closely fitted than the provided labels. As shown in Fig. ??, zooming into the ROI further illustrates

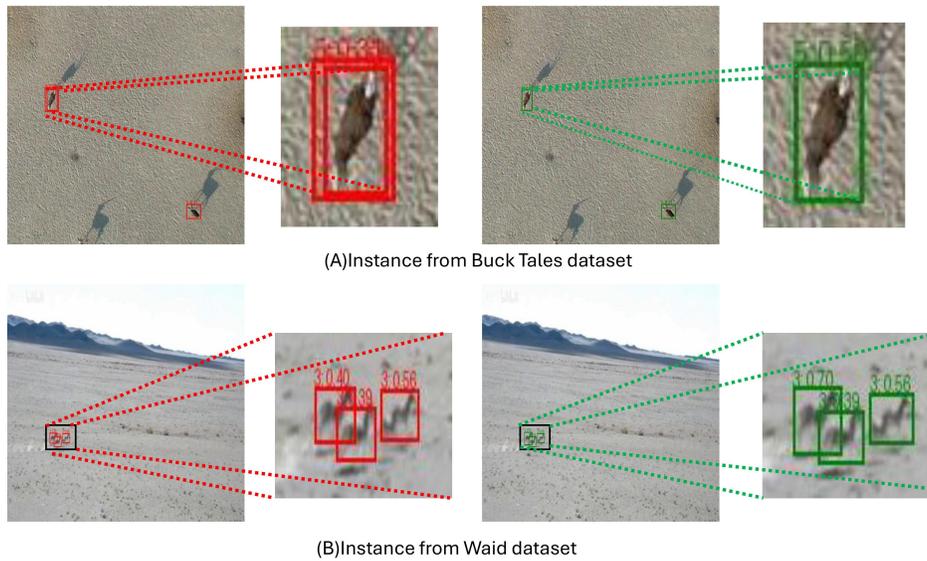


Figure 4: Comparing the ROI of predicted anchor boxes from a single inference (shown in red) versus a two-step inference (shown in green), highlighting the removal of overlapping boxes and the increase in object confidence scores.

the advantages of two-stage inference over single-stage inference. The authors recognize this as an open problem and plan to explore its implications in drone surveillance, performance quantification, and wild animal detection in future work.

## D CONCLUSION AND ACKNOWLEDGEMENT

In this work, we have presented DEAL-YOLO, a novel approach to animal detection that showcases the superior performance of up to 66.93% lesser trainable parameters on BuckTales and comparable metrics to SOTA at 69.59% lesser trainable parameters on WAID. The SSFF module, LD convolutions, and our novel 2-stage inference paradigm demonstrate excellent results across UAV-captured datasets like WAID and BuckTales.

We would like to thank Mars Rover Manipal, an interdisciplinary student project of MAHE, for providing the essential resources and infrastructure that supported our research. We also extend our gratitude to Mohammed Sulaiman for his contributions in facilitating access to additional resources crucial to this work. refer to tables, explain our results in