

CONTINUAL YOLO-BASED DETECTION FOR LONG-TERM MONITORING OF THE GREATER ONE-HORNED RHINO

Nelson R Varte^{1*}, Kaustubh Bhattacharyya², Navajit Saikia¹

¹ Department of Computer Applications, Assam Engineering College, Jalukbari, Guwahati, India

² Department of Electronics and Communication Engineering, Assam Don Bosco University, Azara, Guwahati, India

Emails:kaustubh.d.electronics@gmail.com²; navajit.ete@aec.ac.in¹

***Corresponding Author:**

Email: nelson.varte@gmail.com

Abstract

YOLO-based object detectors support fast wildlife monitoring across diverse habitats. Static training limits these detectors in the field because ecological conditions shift over time. Continual learning provides an incremental update process that preserves earlier knowledge while absorbing new information.

This article reviews recent progress in continual learning for YOLO detection, with emphasis on experience replay, self-distillation, and human-in-the-loop supervision. These approaches protect past knowledge, reduce annotation demands, and deliver more stable predictions during long-term deployments.

A focused case study on Greater One-Horned Rhino monitoring shows how adaptive learning pipelines strengthen detection reliability during seasonal changes, altered terrain, and new camera placements. The review outlines methods with strong potential for long-term conservation work and highlights future directions for resilient wildlife monitoring systems.

Keywords: Continual Object Detection, Wildlife monitoring, YOLO, replay memory, self-distillation, human-in-loop, catastrophic forgetting

1. Introduction

Accurate monitoring of endangered wildlife relies on reliable and scalable computer vision tools. YOLO detectors support real-time identification and deliver strong performance across varied environments [1][2][3]. Traditional YOLO training uses static datasets. This limits performance when field conditions shift, new viewpoints arise, or animal appearances vary. Continual learning addresses this issue by updating models through incremental steps while retaining earlier detections [1][2][4]. Experience replay and distillation-based methods produce stable updates and preserve key knowledge [5][6]. These advances support detectors used in conservation scenarios with frequent environmental changes.

Ecological monitoring introduces additional obstacles. Data scarcity for rare species, heavy class imbalance, and persistent domain shifts complicate long-term deployments [4][5][7].

Conservation teams also rely on feedback loops involving domain experts. These loops reduce errors in uncertain cases and improve long-term accuracy.

This article reviews technical progress in continual object detection within the YOLO family. It examines key methods and connects them to the specific demands of Greater One-Horned Rhino monitoring.

2. Background

2.1 Deep Learning and the YOLO Framework

Object detection identifies objects and localizes them within images. YOLO approaches this with a single forward pass, which produces strong speed and accuracy [8][9][10]. Successive YOLO versions improved detection of small objects, performance in crowded scenes, and efficiency during deployment [8][9][10][11].

These properties support wildlife monitoring tasks that involve large volumes of image data from camera traps, drones, or fixed surveillance systems [9][12].

2.2 Continual Learning

Conventional training procedures assume stable datasets. Field deployments present new classes, lighting variations, and fluctuating backgrounds. Continual learning introduces training routines that integrate new data in stages while retaining earlier outputs [13][11][7].

Common strategies include:

- experience replay, where past samples remain in a buffer [7]
- distillation losses, which transfer knowledge from an earlier model to a new model [11][7]
- memory-augmented methods that balance retention and adaptation [7]

2.3 Wildlife Monitoring Requirements

AI supports wildlife work through camera trap analysis, drone-based surveys, and habitat assessment [12][17]. Many conservation projects operate under difficult conditions, limited staffing, and irregular data flow. Endangered animals appear infrequently, which results in sparse labels and severe dataset imbalance [18][19].

Continual learning aligns with these constraints because it supports periodic updates instead of retraining from scratch.

2.4 Integration of Continual Learning and Ecological Monitoring

Recent work links continual learning with ecological monitoring [13][7]. These efforts aim to maintain detection quality throughout seasonal transitions, environmental disturbances, and long-term conservation programs. YOLO-based detectors paired with continual learning offer a practical route for adaptive wildlife monitoring.

3. Technical Advances in Continual YOLO Object Detection

3.1 Overview

Continual learning research focuses on stable updates under long-term sequences of new data [1]. Effective methods reduce catastrophic forgetting while improving accuracy on new tasks.

3.2 Experience Replay

Experience replay stores selected samples from earlier tasks in a memory buffer [2][7]. Mixing these samples with new data during fine-tuning slows knowledge loss.

Balanced replay strategies, including label propagation and sample consolidation, improve stability in long training sequences and support rare-species monitoring [10]. These strategies help maintain consistent predictions across changing field conditions.

3.3 Self-Distillation and Regularization

Self-distillation transfers outputs from an earlier model to a new model through soft targets [1]. YOLO-specific research refined distillation for object detection by masking unreliable predictions and adjusting loss terms for regression tasks. These improvements stabilized updates across multiple stages and produced strong results on standard benchmarks.

3.4 Human-in-the-Loop

Human supervisors validate uncertain detections and supply high-quality labels for difficult cases [5]. Selective annotation reduces workload and strengthens long-term detection accuracy in ecological datasets. Studies report significant reductions in labeling time when experts focus on ambiguous samples [6].

3.5 Performance Across Benchmarks

Replay, distillation, and human-in-the-loop workflows outperform naive fine-tuning approaches on VOC, COCO, and wildlife datasets [1][4]. These methods support long-term ecological deployments that require consistent results under evolving conditions.

4. Current Challenges and Open Problems

4.1 Catastrophic Forgetting

Models lose earlier detection abilities when exposed to extended sequences of new cases. Replay and distillation slow this decline but do not eliminate it [20][4].

4.2 Distribution Shifts

Wildlife imagery varies with weather, vegetation, and sensor placement. Sudden changes in background conditions degrade detector performance until new data gets integrated [20].

4.3 Small, Overlapping, or Rare Objects

YOLO struggles with small or occluded animals. Rare species worsen this issue because few samples appear in the replay buffer [2].

4.4 Annotation Gaps

Wildlife datasets include missing labels. Missing labels mislead detectors during incremental updates. Human-in-the-loop review helps, though a complete solution remains unavailable [4][6].

4.5 Resource Constraints

Replay buffers require memory and processing resources. Field deployments often use limited hardware [9].

4.6 Monitoring and Evaluation

Long-term wildlife systems need continuous quality checks, automated alerts, and anomaly detection [20].

4.7 Case Comparisons

Whale monitoring with human-in-the-loop achieved 95 percent mAP and reduced annotation effort by 75 percent [19].

Roadside wildlife detection improved recall by 10 percent when self-training and replay were applied [4].

Other work, such as YOLO-SAG and DEAL-YOLO, improved small-object detection and reduced model size for remote deployments [21][22].

Performance Comparison of YOLO-Based Continual Object Detection Methods Across Wildlife Conservation Case Studies is shown in Table 1 below.

Dataset/Species	Method	mAP	F1 Score	Recall Improvement	Annotation Effort Reduction
Whale Monitoring [19]	YOLO + Human-in-the-Loop	95%	0.95	12%	75%
Endangered Species Road Monitoring [4]	YOLO + Self-Training + Replay	92%	0.90	10%	60%

Table 1. Performance Comparison

5. Ecological Case Study: Greater One-Horned Rhino Monitoring

5.1 Conservation Importance

Greater One-Horned Rhinos remain concentrated in Assam's floodplains. Conservation teams handle poaching pressures, recurring floods, and habitat disturbances [19][23][25]. These challenges increase the need for reliable detection tools that support long-term ecological monitoring and rapid response efforts.

Recent advances in rhino-focused YOLO systems provide strong baselines for automated surveillance and contribute directly to these conservation requirements [23][25].

5.2 Standard YOLO Deployment

Recent research demonstrates that YOLO-based detectors perform reliably for Greater One-Horned Rhino monitoring in controlled and semi-structured environments. Earlier work established a strong baseline for real-time rhino detection, achieving mAP values above 98 percent on curated field datasets [23]. Subsequent advances extended this foundation through deployment-oriented studies.

A YOLO-based road-safety wildlife detection system delivered robust performance under variable lighting and background conditions, confirming the suitability of one-stage detectors for resource-constrained field applications [24]. More recent work introduced a YOLO instance segmentation framework with morphological filtering and semantic false-positive suppression, which improved precision in complex vegetation and cluttered backgrounds common in rhino habitats [25].

Together, these deployments show that YOLO architectures provide strong, adaptable baselines for rhino monitoring before continual learning strategies are applied.

5.3 Integrating Continual Learning

Field work in [4] used a self-training scheme that updated detectors through weakly labeled data streams. The system gained accuracy over time and limited annotation load.

Work in [5] used a human-in-the-loop strategy in aerial surveys. Experts corrected uncertain samples, which guided incremental updates.

For rhino monitoring, an adaptive pipeline would support:

- ongoing updates from new camera trap or drone imagery
- selective expert review for difficult frames
- retention of earlier knowledge through replay
- improved resilience under seasonal shifts

Future work should test these methods in rhino habitats with cloud-edge infrastructure and periodic expert review.

6. Future Directions

6.1 Continual Learning Pipelines

Efficient pipelines must integrate replay, distillation, and targeted expert review. Research should examine lightweight algorithms suited for edge devices in remote areas.

6.2 Efficient Annotation

Semi-supervised learning and selective sampling reduce workload. These methods help process sightings of rare behaviors or unusual habitat conditions.

6.3 Domain Adaptation

Long-term deployments require generalization across multiple sites. Domain-adaptive and federated approaches offer promising directions for shared conservation efforts.

6.4 System Monitoring

Field systems need automated evaluation to flag silent failures or domain drift. Reliable alert mechanisms strengthen conservation response.

6.5 Shared Datasets

Open, regularly updated wildlife datasets support benchmarking and reduce duplication of effort. Strong collaboration between ecologists and engineers strengthens progress.

7. Conclusion

YOLO detectors support fast and accurate wildlife monitoring. Static training limits these detectors in dynamic environments. Continual learning strengthens long-term performance by integrating new data while preserving earlier knowledge.

This review examined replay, distillation, and human-in-the-loop methods, with evidence from whale monitoring and threatened wildlife road monitoring. The case study on Greater One-Horned Rhinos shows clear benefits for adaptive conservation programs.

Future work should focus on lightweight continual learning pipelines, improved annotation tools, and robust evaluation systems. These directions support scalable and resilient wildlife monitoring solutions suited to long-term ecological challenges.

References:

- [1] R. De Monte, D. Dalle Pezze, and G. A. Susto, “Teach YOLO to Remember: A Self-Distillation Approach for Continual Object Detection,” *arXiv:2503.04688*, 2025.
- [2] J.-L. Shieh, Q. M. U. Haq, M. A. Haq, S. Karam, P. Chondro, D.-Q. Gao, and S.-J. Ruan, “Continual Learning Strategy in One-Stage Object Detection Framework Based on Experience Replay for Autonomous Driving Vehicle,” *Sensors*, vol. 20, no. 23, p. 6777, 2020.
- [3] J. Redmon and A. Farhadi, “YOLOv3: An Incremental Improvement,” *arXiv:1804.02767*, 2018.
- [4] K. Li, M. Shan, S. B. Perez, K. Luo, and S. Worrall, “Endangered Alert: A Field-Validated Self-Training Scheme for Detecting and Protecting Threatened Wildlife on Roads and Roadsides,” *arXiv:2412.12222*, 2024.
- [5] J. Boulent, B. Charry, M. M. Kennedy, E. Tissier, R. Fan, M. Marcoux, C. A. Watt, and A. Gagné-Turcotte, “Scaling Whale Monitoring Using Deep Learning: A Human-in-the-Loop Solution for Analyzing Aerial Datasets,” *Frontiers in Marine Science*, vol. 10, p. 1099479, 2023.
- [6] S. E. Huebner, M. S. Palmer, and C. Packer, “Human Supervision Is Key to Achieving Accurate AI-Assisted Wildlife Identifications in Camera Trap Images,” *Citizen Science: Theory and Practice*, vol. 9, no. 1, 2024.
- [7] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, and S. Wermter, “Continual Lifelong Learning With Neural Networks: A Review,” *Neural Networks*, vol. 113, pp. 54–71, 2019.
- [8] Y. Xu, Y. Wu, and X. Wang, “Continual Object Detection: A Review,” *arXiv:2205.15445*, 2022.
- [9] E. Fini, S. Lathuilière, E. Sangineto, M. Nabi, and E. Ricci, “Online Continual Learning Under Extreme Memory Constraints,” in *Proc. ECCV*, 2020, pp. 704–720.
- [10] R. De Monte, D. Dalle Pezze, M. Ceccon, F. Pasti, F. Paissan, E. Farella, G. A. Susto, and N. Bellotto, “Replay Consolidation With Label Propagation for Continual Object Detection,” *arXiv:2409.05650*, 2024.

[11] M. De Lange *et al.*, "A Continual Learning Survey: Defying Forgetting in Classification Tasks," in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 7, pp. 3366-3385, 1 July 2022, doi: 10.1109/TPAMI.2021.3057446.

[12] Beery, S., Van Horn, G., Perona, P. (2018). Recognition in Terra Incognita. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds) Computer Vision – ECCV 2018. ECCV 2018. Lecture Notes in Computer Science(), vol 11220. Springer, Cham. https://doi.org/10.1007/978-3-030-01270-0_28

[13] S. K. Beery *et al.*, "The iWildCam 2020 Competition Dataset: Evaluating Transfer Learning for Wildlife Classification," *arXiv:2004.10340*, 2020.

[14] C.-Y. Wang, A. Bochkovskiy, and H.-Y. M. Liao, "YOLOv7: Trainable Bag-of-Freebies Sets New State-of-the-Art for Real-Time Object Detectors," *arXiv:2207.02696*, 2022.

[15] M. Masana, X. Liu, B. Twardowski, M. Menta, A. D. Bagdanov and J. van de Weijer, "Class-Incremental Learning: Survey and Performance Evaluation on Image Classification," in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 5, pp. 5513-5533, 1 May 2023, doi: 10.1109/TPAMI.2022.3213473.

[16] Schneider, Stefan & Taylor, Graham & Linquist, Stefan & Kremer, Stefan. (2018). "Past, Present and Future Approaches Using Computer Vision for Animal Re-Identification from Camera Trap Data.", *Methods in Ecology and Evolution*, vol. 10, doi: 10.1111/2041-210X.13133.

[17] M.S. Norouzzadeh,A. Nguyen,M. Kosmala,A. Swanson,M.S. Palmer,C. Packer, & J. Clune, Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning, *Proc. Natl. Acad. Sci. U.S.A.* 115 (25) E5716-E5725, <https://doi.org/10.1073/pnas.1719367115> (2018).

[18] Fraunhofer IIS, "Continual Learning in Object Detection," *Fraunhofer Technical Report*, 2023.

[19] S. Vasavi, V. Sripathi, and C. M. Simma, "Visualization of Humpback Whale Tracking on Edge Device Using Space-Borne Remote Sensing Data," *Egyptian Journal of Remote Sensing and Space Sciences*, vol. 27, no. 4, pp. 705–715, 2024.

[20] S. Farquhar and Y. Gal, "Towards Robust Evaluations of Continual Learning," *arXiv:1805.09733*, 2019.

[21] L. Chen, G. Li, S. Zhang, W. Mao, and M. Zhang, "YOLO-SAG: An Improved Wildlife Object Detection Algorithm Based on YOLOv8n," *Ecological Informatics*, vol. 83, 102791, 2024.

[22] A. P. Naidu, H. Gosalia, I. Gakhar, S. S. Rathore, K. Didwania, and U. Verma, "DEAL-YOLO: Drone-Based Efficient Animal Localization Using YOLO," *arXiv:2503.04698*, 2025.

[23] N. R. Varte, K. Bhattacharyya, and N. Saikia, "Advancing Environmental Monitoring: YOLO Algorithm for Real-Time Detection of Greater One-Horned Rhinos," *International Journal of Environmental Science*, 2025.

[24] N. R. Varte *et al.*, "Animal Detection and Alert System for Road Safety," 2025 IEEE Guwahati Subsection Conference (GCON), Itanagar, India, 2025, pp. 1-6, doi: 10.1109/GCON65540.2025.11173390.

[25] N. R. Varte, K. Bhattacharyya, and N. Saikia, "Advancing Greater One-Horned Rhino Conservation: A YOLO Instance Segmentation Framework with Morphological and Semantic False Positive Suppression," *Journal of Environmental Protection and Ecology*, vol. 26, no. 6, pp. 2427–2438, 2025.