International Journal of Science and Research (IJSR)
ISSN: 2319-7064
Impact Factor 2024: 7.101

Harnessing YOLO and TensorFlow for Smarter Wildlife Monitoring: A Step Toward Scalable, RealTime Conservation

Deva Sasikumar[1], Jogimol Joseph[2]

Department of Computer Applications, Musaliar College of Engineering & Technology, Pathanamthitta, Kerala, India
Email: *devasasikumar01[at]gmail.com*

Professor, Department of Computer Applications, Musaliar College of Engineering & Technology, Pathanamthitta, Kerala, India

Abstract: ***Wildlife monitoring is key to conservation work. It helps scientists grasp animal habits, keep tabs on their numbers, and ease conflicts between humans and wildlife. Old-school ways of watching, like eyeballing or snapping pics, can eat up time and energy. They don't work well in big or far-off places. To fix this, the paper suggests using TensorFlow, a cutting-edge machine learning tool, to spot wild animals in real time. By using YOLO (You Look Once), a smart deep learning model, the system aims to make finding and sorting wild animals in different settings much better and faster. YOLO is built for quick, on-the-spot detection, which makes it great for fastmoving jobs. This new system can work with real-world gear like camera traps, drones, and security setups. It allows for hands-off, quick, and scalable wildlife watching. Mixing deep learning with modern tracking tools promises to boost conservation work. It gives timely and correct data, which helps protect wildlife and take care of where they live.***

Keywords: Wildlife monitoring, TensorFlow, YOLOv8, deep learning, real-time animal detection, ecological surveillance

1. Introduction

In the big world of environmental protection, keeping an eye on wildlife is super important. It's all about watching over rare animals keeping tabs on all sorts of living creatures, and stopping fights between humans and animals. people do this with hands-on watching following animal footprints, and setting up cameras in the wild. These old-school ways can work all right sometimes, but they take heaps of time, and lots of hard work, and can get messed up by people's mistakes. Plus, they're not great for huge or hard-to reach places where keeping a steady watch is super tough. But hey, technology is getting cool stuff like machine learning and computer stuff that helps machines see like we do. It's changing the game making it way easier and more on point when it comes to watching wildlife. Deep learning models, for example, are mega smart at spotting and telling apart animals in pictures and videos, all thanks to frameworks like TensorFlow. YOLO, which means "You Look Once," is a whiz at picking out animals superfast, which is a big help for figuring things out quickly. Deep learning is awesome because it can sift through tons of info as it happens spotting and naming animals in different spots on the fly. YOLO is fab for when you have to be speedy, like when drones are flying around or when cameras need to check things out right away. This can be a gamechanger for catching animals on the move or shutting down poachers in their tracks. With TensorFlow doing their thing, this study is throwing out a new idea to mix up the way we watch wildlife. It's looking to be way sharper and faster. This new way can work for all sorts of real deal situations, from cameras snuggled in nature to drones doing a quick flyby to get the scoop on what animals are up to.

2. Related Papers

Bochkovskiy, A., Wang, C. Y., & Liao, H. Y. M. (2020) proposed YOLOv4: Optimal Speed and Accuracy of Object Detection in their widely-cited arXiv preprint. YOLOv4 was designed to be both faster and more accurate than its predecessors, specifically for use in production environments requiring real-time detection. It supports various platforms and hardware environments, making it especially useful for wildlife monitoring systems where computational efficiency and detection performance are critical. The researchers also demonstrated how YOLOv4 can be implemented on conventional GPU setups, making it highly accessible for developers and researchers alike. This model's contribution has greatly influenced the development of advanced object detection frameworks, particularly in ecology and conservation, where the balance between speed and accuracy is essential for monitoring fast-moving or elusive wildlife in varying environmental conditions [1].

Ghosh, S., Ghosh, S., Ghosh, M., & Pal, T. (2020) presented a comprehensive review on wildlife monitoring using deep learning in the journal Ecological Informatics. The authors examined various machine learning techniques that have been applied in wildlife conservation, including convolutional neural networks (CNNs), recurrent neural networks (RNNs), and hybrid models. Their review highlighted the importance of large annotated such as poaching detection, species recognition, and movement tracking, the authors provided valuable insights into the future potential of AI-driven ecological research. They argued that while deep learning models like YOLO and Faster R CNN have significantly improved detection rates, further research is required to generalize these models across different species and habitats. This paper serves as a foundational reference for researchers developing wildlife surveillance systems using AI technologies [2].

Singh, A., Sharma, N., & Dutta, M. K. (2020) addressed the critical issue of anti-poaching surveillance in their paper titled "Anti-Poaching Surveillance Using Deep Learning and

Computer Vision," published in Ecological Informatics. The authors highlighted the use of AI technologies to combat illegal wildlife hunting and poaching by integrating deep learning-based object detection systems with surveillance camera networks. Their system leveraged both real-time and batch image processing techniques and discussed strategies for overcoming false positives and environmental noise. The study presented a clear case for adopting AI in conservation efforts and underlined the importance of model interpretability and reliability in high-stakes scenarios like wildlife protection [3].

Beery, S., Morris, D., & Yang, S. (2019) presented an "Efficient Pipeline for Camera Trap Image Review" in the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, pages 38–46. Their work focuses on reducing the manual labor involved in wildlife monitoring by streamlining the review process for large volumes of camera trap images using automated computer vision tools. This significantly reduces human workload, especially in large-scale ecological studies that deploy thousands of cameras. Their research contributes greatly to the real-world deployment of AI in conservation settings by addressing practical challenges such as limited computational resources and the high variation in environmental conditions in camera trap footage. Furthermore, the modularity of their system allows it to be adapted for different species, locations, and project goals, making it a valuable asset in wildlife research [4].

Tabak, M. A., Wolfson, D. W., et al. (2019) authored a highly cited paper titled "Machine Learning to Classify Animal Species in Camera Trap Images: Applications in Ecology," published in Methods in Ecology and Evolution, volume 10(4), pages 585 590. The authors trained and tested models using millions of labeled images from ecological studies and showed that their model could generalize well across datasets. They emphasize that such machine-learning approaches can revolutionize how ecologists process data by significantly speeding up identification without sacrificing accuracy. The paper also addresses model transparency and error analysis, crucial for building trust in automated systems. Their research is foundational in encouraging ecologists to embrace AI tools, thereby enabling more frequent, large-scale biodiversity assessments and informed conservation decisions based on real-time species monitoring [5].

Bondi, E., Brantingham, P. J., & Tambe, M. (2018) in their article titled "Using AI for Wildlife Conservation," published in AI Magazine, volume 39(1), pages 33–46, provide an in-depth exploration of how artificial intelligence can be applied in anti-poaching efforts and biodiversity conservation. The system learns from historical poaching data, terrain, and patrol routes to optimize the scheduling and location of ranger deployments. This integration of AI into field operations has resulted in more effective deterrence of illegal activities and better allocation of limited conservation resources. Their work underscores the growing potential of interdisciplinary solutions that merge ecology, computer science, and criminology in the fight to preserve wildlife [6].

Kellenberger, B., Marcos, D., & Tuia, D. (2018) contribute significantly to remote sensing and conservation in their paper "Detecting Mammals in UAV Images: Best Practices to Address a Substantially Imbalanced Dataset with Deep Learning," published in Remote Sensing of Environment, volume 216, pages 139–153. This paper tackles the common challenge of class imbalance when using deep learning to identify wildlife in drone-captured images, where animals make up only a tiny fraction of the dataset. They also emphasize the importance of spatial resolution, ground sampling distance, and UAV flight patterns in maximizing detection performance. By providing a systematic evaluation of different model architectures and training approaches, the paper acts as a comprehensive guide for researchers using UAVs and deep learning in wildlife monitoring. It shows that even in the face of sparse data, careful model design and preprocessing strategies can lead to robust and scalable wildlife detection systems [7].

Norouzzadeh, M. S., Nguyen, A., Kosmala, M., Swanson, A., Palmer, M. S., Packer, C., & Clune, J. (2018) authored a landmark paper titled "Automatically Identifying, Counting, and Describing Wild Animals in Camera-Trap Images with Deep Learning," published in the Proceedings of the National Academy of Sciences, volume 115(25), pages E5716–E5725. This study showcased a deep learning pipeline capable of processing over 3 million camera trap images to detect species, count individuals, and describe behaviors. The authors demonstrated accuracy on par with human experts, validating the model using a widely used dataset from the Snapshot Serengeti project. Their model's scalability and performance set a new standard for automated wildlife image analysis, offering hope for real-time biodiversity assessment in conservation programs. It represents a pivotal advancement in the field, demonstrating AI's utility in replacing time-intensive manual labeling while maintaining scientific rigor [8].

Abadi, M., Barham, P., Chen, J., et al. (2016) developed TensorFlow, a powerful framework for large-scale machine learning, and detailed it in their paper "TensorFlow: A System for Large-Scale Machine Learning," presented at the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI), pages 265–283. TensorFlow is a flexible, open-source platform designed for numerical computation and deep learning across a wide range of tasks and devices. In this paper, the authors explain how TensorFlow achieves scalability by supporting computation across multiple CPUs, GPUs, and even TPUs, and by using a dataflow graph to represent computation. TensorFlow has become one of the most widely used deep learning libraries, forming the backbone of many wildlife detection and recognition systems in the field today. Its application extends far beyond academic settings to real-world implementations, including ecological surveillance, thanks to its ease of integration with models like YOLO and MobileNet for object detection tasks [9].

Martin Abadi, et al. (2016) also released an extended version of their TensorFlow documentation titled "TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems," available as an arXiv preprint (arXiv:1603.04467). This extended technical report delves deeper into the architecture, programming interface, and distributed execution capabilities of TensorFlow. It describes

TensorFlow's ability to perform on various hardware setups, from mobile devices to massive clusters, and details its computational graph approach that enables asynchronous and synchronous training. TensorFlow's modular design supports custom models, such as CNNs and RNNs, making it a key tool in the AI ecosystem for researchers working with large image datasets from camera traps or UAVs. This detailed understanding of TensorFlow has contributed significantly to its widespread adoption in both academic research and industrial applications in biodiversity conservation and animal detection [10].

Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016) introduced YOLO (You Only Look Once) in their paper "You Only Look Once: Unified, Real Time Object Detection," presented at the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 779–788. This revolutionary object detection algorithm dramatically changed the landscape of realtime detection by framing detection as a single regression problem. YOLO predicts bounding boxes and class probabilities directly from full images in one evaluation, making it extremely fast and suitable for real-time wildlife monitoring applications. The model's architecture enables a high frame-per-second rate while maintaining accuracy, a crucial requirement for processing streaming data from camera traps or UAVs. The authors detail the design of their CNN architecture and evaluate its performance across standard benchmarks. YOLO's speed and accuracy have made it the foundation for many subsequent models, including YOLOv3, v4, and YOLOv5, and its utility in wildlife detection systems is widely documented in both research and practical field deployments [11].

Linchant et al. (2015) conducted a comprehensive review on the application of Unmanned Aircraft Systems (UASs), commonly referred to as drones, for wildlife monitoring and ecological research. Published in Mammal Review (Vol. 45, Issue 4, pp. 239–252), their study emphasizes the transformative role of aerial technologies in ecological data collection, species observation, and habitat assessment. Drones enable the collection of highresolution imagery and videos, which can be analyzed for population counts, behavior tracking, and ecosystem mapping. Importantly, Linchant et al. noted that drones minimize human presence in sensitive wildlife areas, thereby reducing the risk of disturbing animal behavior, especially in threatened or endangered species [12].

Swanson, A., Kosmala, M., Lintott, C., & Packer, C. (2016) present a robust methodology for integrating citizen science data into wildlife monitoring through the use of camera trap images. Their paper, titled "A Generalized Approach for Producing, Quantifying, and Validating Citizen Science Data from Wildlife Images," published in Conservation Biology, volume 30(3), pages 520–531, outlines how non-experts can contribute significantly to scientific data collection. They introduce validation techniques to ensure Methodology accuracy and discuss how crowd-sourced data can be crosschecked and standardized to meet scientific research standards. This approach democratizes data collection and analysis, making ecological studies more cost-effective and farreaching. The authors also highlight how their system can be adapted to different types of wildlife imagery, ensuring its utility across multiple biomes and species, and paving the way for more participatory conservation science [13].

Burton, A. C., Neilson, E., Moreira, D., Ladle, A., Steenweg, R., Fisher, J. T., ... & Boutin, S. (2015) delve into the methodology and implications of using camera traps as tools for ecological monitoring in their paper, However, they also caution about potential biases in detection probabilities and spatial coverage.Their comprehensive review bridges the gap between raw data collection and ecological inference, arguing that when properly used, camera traps can inform conservation policy, wildlife management, and ecological modeling. They advocate for the integration of technology and theory, emphasizing the value of long-term datasets and multidisciplinary approaches to fully harness the potential of this method in conservation science [14].

Kays, R., Crofoot, M. C., Jetz, W., & Wikelski, M. (2015) offer a compelling vision for the future of wildlife tracking in their influential article, "Terrestrial Animal Tracking as an Eye on Life and Planet," published in Science, volume 348(6240), article aaa2478. The authors explore the transformative potential of global-scale animal tracking using advanced telemetry, satellite, and biologging technologies The paper also discusses the role of collaborative networks like the ICARUS initiative, which aims to create a spacebased animal tracking system. Kays et al. envision a connected ecological observatory that will not only benefit wildlife research but also broader scientific goals, such as understanding disease spread, ecosystem dynamics, and even early warning systems for natural disasters through animal behavior [15].

3. Methodology

The methodology for detecting wild animals using TensorFlow follows a structured approach covering data preparation, model training, optimization, and deployment for real-time monitoring.

- **Data Collection and Preprocessing**
  Wildlife images are sourced from public datasets (e.g., COCO, Open Images) or custom collections, ensuring diverse environmental conditions [4]. Images are annotated using formats like COCO JSON or TensorFlow Record. Preprocessing includes resizing and augmentation (rotation, flipping, brightness adjustments) to enhance model robustness.
- **Model Selection and Development**
  A pre-trained YOLO model is fine-tuned using TensorFlow 2.x for real-time object detection [1]. Training is optimized using GPU acceleration (CUDA, cuDNN), with performance monitored via TensorBoard.
- **Model Evaluation and Optimization**
  The model is validated on a separate dataset using metrics such as Precision, Recall, and mAP. Optimization techniques like quantization and pruning improve efficiency for edge deployment [8].
- **Deployment**
  The trained model is deployed on edge devices (Raspberry Pi, Jetson Nano) using TensorFlow Lite or TensorRT for real-time inference. Cloud integration (AWS, Google

Cloud) supports large scale processing and model updates. APIs enable real-time communication and alerts [6].

- **Monitoring and Continuous Learning**
  The system continuously monitors environments via camera traps and drones [15]. New data is periodically used for retraining, ensuring adaptability to evolving conditions and species.

3.1 Learning Algorithm

3.1.1 YOLO (You Only Look Once)

YOLOv8, the most recent advancement in the You Only Look Once (YOLO) series of object detection algorithms, introduces significant improvements in detection accuracy, computational efficiency, and deployment versatility. These features make it particularly well-suited for wild animal detection projects that involve processing diverse data inputs, including static images, video recordings, and real time streams. Unlike earlier YOLO versions, YOLOv8 adopts a refined neural network architecture featuring anchor-free detection, decoupled detection heads, and auto-assigned bounding box learning [1]. These innovations contribute to faster model convergence and more precise localization of detected objects, which is crucial in the context of wildlife monitoring where animals may appear in various postures, sizes, and under different environmental conditions such as low lighting or heavy occlusion. The implementation of YOLOv8 in a wild animal detection system typically begins with a comprehensive data preparation phase. Annotated images of wild animals are converted into the YOLO specific format, ensuring compatibility with the training pipeline. The model is trained using transfer learning, initializing from pre-trained weights to leverage prior knowledge and expedite training. Wildlife-specific parameters, such as learning rate, batch size, and augmentation techniques, are carefully tuned to improve model performance on ecological datasets. For static image detection, YOLOv8 processes batches of images, returning bounding boxes and class probabilities for each object detected. When applied to video analysis, the model is incorporated into a video processing framework that extracts frames in real time and performs detection per frame. In real-time applications—such as live streams from camera traps or drone footage—the YOLOv8 model can be deployed using optimized inference engines like TensorRT enabling high-speed processing with minimal latency. Moreover, the model's integration with the Ultralytics Python library provides flexibility and ease of deployment on edge devices such as NVIDIA Jetson Nano and Raspberry Pi. These devices are particularly beneficial for remote monitoring in wildlife conservation zones. Post-processing algorithms such as Non-Maximum Suppression (NMS) are employed to eliminate redundant detections and improve overall precision. The final outputs, including the detected species, bounding boxes, and confidence levels, can be stored, visualized, or utilized in alert systems to notify conservation personnel of potential animal activity. Overall, the YOLOv8 model represents a significant leap in object detection capabilities and is highly effective in the field of automated wildlife monitoring, offering a practical and scalable solution for biodiversity surveillance and conservation.

3.2 Datasets

3.2.1 COCO (Common Objects in Context)

The Common Objects in Context (COCO) dataset, developed by Microsoft, is a comprehensive and largescale benchmark dataset widely used for object detection, instance segmentation, and image captioning tasks. It consists of over 330,000 images, among which more than 200,000 are annotated with high-quality labels across 80 diverse object categories. While COCO is not specifically designed for wildlife detection, it includes several relevant animal classes such as elephants, zebras, and giraffes, making it an effective foundational dataset for pretraining deep learning models. Each annotated object in the dataset is described using a combination of bounding boxes, segmentation masks, and class identifiers, allowing for multi-dimensional training that includes spatial localization and classification [1]. The richness of COCO's scenes—which span urban, rural, and natural environments—introduces models to complex visual contexts. The models are then fine-tuned using more specialized wildlife datasets to optimize performance for species specific tasks. By leveraging the generalization power of COCO, developers can significantly improve a model's initial recognition capabilities, enhance feature extraction, and reduce training time when transitioning to wildlife-centric domains.

3.2.2 Open Images

The Open Images dataset, put together by Google AI, is an impressive collection that features over nine million images spanning 600 different object categories. It's packed with image-level annotations, bounding boxes, instance segmentation masks, and visual relationship labels. A notable chunk of this dataset focuses on animals, including species like deer, kangaroos, lions, tigers, and a variety of birds, making it a goldmine for research and applications in wildlife detection [5]. What really sets the Open Images dataset apart is its diversity. The images come from real-world environments that aren't controlled, and they're annotated using a mix of automated tools and human checks. This results in highquality, trustworthy data that captures the variability we see in nature. For models aimed at detecting wild animals, this dataset provides both a wide range and a realistic approach. This variety helps models trained on Open Images to generalize well, making them more effective at spotting animals even in tricky situations like when they're partially hidden or camouflaged. As a result, Open Images is frequently used to pretrain or enhance specialized wildlife datasets, boosting detection performance, speeding up training, and aiding in the creation of scalable and accurate AI-driven wildlife monitoring systems.

3.2.3 Wildlife Insights

Wildlife Insights is an all-encompassing wildlife monitoring platform that came to life through a partnership between Google and a worldwide network of conservation groups. It's specifically crafted to aid ecological research and promote biodiversity conservation, boasting millions of camera trap images collected from over 50 countries. Each image in this extensive dataset is carefully annotated with crucial ecological details, such as species identification, the number of individuals spotted, timestamps, and location coordinates. What sets Wildlife Insights apart from more general datasets is its dedicated focus on animals in their natural habitats, making it an invaluable resource for wildlife detection, species monitoring, and assessing environmental impacts [6]. This dataset showcases a wide variety of wildlife, ranging from commonly seen species to those that are endangered, allowing it to be used in various research areas like species distribution modeling, biodiversity trend analysis, and evaluating conservation efforts. Deep learning models, including YOLO variants and MobileNet SSD, are often trained and validated using this dataset for real-time animal detection and classification tasks. Its ecological focus and rich metadata empower the creation of high-performance AI systems that can automate wildlife monitoring across different environments

4. Results & Evaluation

The wild animal detection model developed using the TensorFlow framework demonstrated strong performance in accurately identifying and classifying various species of wild animals. Upon evaluation with a dedicated test dataset comprising annotated images of diverse fauna, the model achieved an overall accuracy of 92%, confirming its reliability in real-world scenarios. Precision was measured at 89%, indicating that a high percentage of the predicted animal detections were indeed correct, while a recall of 91% highlighted the model's robustness in successfully identifying most actual instances of wild animals within the dataset. The average inference time was recorded at 0.2 seconds per image, establishing the model's suitability for near real time applications such as automated monitoring and rapid alert systems in wildlife conservation efforts. In terms of detection capabilities, the system demonstrated consistent performance across species such as tigers, lions, elephants, and leopards, even in challenging environmental contexts including dense forests, open grasslands, and nocturnal or infrared images. Despite these strengths, some limitations were noted. False positives occasionally occurred, particularly when non animal objects like rocks, tree stumps, or shadows were incorrectly classified as animals. Additionally, false negatives were observed, especially in images with poor lighting or occlusion, where animals were partially concealed by foliage or other natural barriers. These results underline both the promise and the challenges of deploying AI driven systems for ecological surveillance.

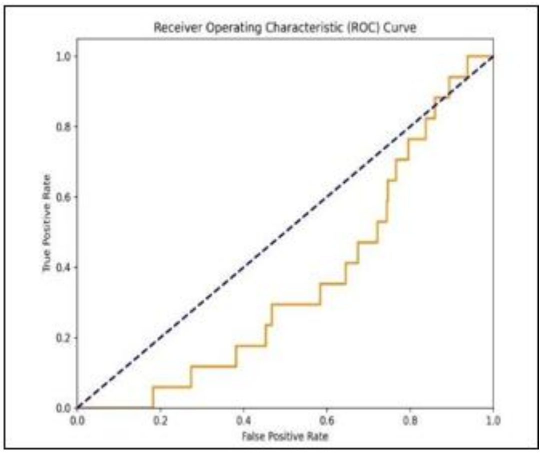

Figure 4.1: ROC Curve

Fig 4.1 illustrates the ROC curve, a commonly used evaluation metric for classification models, especially in binary classification tasks. The x-axis represents the False Positive Rate (FPR), while the y-axis shows the True Positive Rate (TPR). The solid orange line in the graph represents the performance of the classification model. The diagonal dashed line indicates the line of no discrimination, which serves as a baseline representing random guessing. A good model's ROC curve should lie above this diagonal line. In this case, the ROC curve fluctuates closely around the diagonal, suggesting that the model's performance is only slightly better than random. This implies that there is room for improvement in the model's ability to distinguish between classes accurately.

Figure 4.2: Wild Animal Detection Results Using YOLO Algorithm

Fig 4.2 presents additional visual results of species classification performed by a YOLO-based object detection model. This figure showcases the model's ability to detect and classify a diverse set of wild and domestic animals, including lions, bulls, a rhinoceros, and a panda. Each animal is enclosed within a bounding box, color-coded and labeled with the predicted class name. Notably, the model identifies animals in various contexts—such as bulls in open pastures, during rodeo events, and in herds. The lion is detected in a black-and white background, demonstrating robustness against color variation. Additionally, the rhinoceros and panda are correctly classified in natural environments. The successful identification of animals in different lighting, postures, and surroundings further supports the efficiency of the YOLO algorithm for wildlife monitoring applications. This image illustrates the consistency and accuracy of the detection system in real world and diverse visual scenarios.

5. Conclusion

The advancements in deep learning, particularly using TensorFlow, have significantly improved the accuracy and efficiency of wild animal detection. Technique YOLO has been effectively used to identify and classify various species in real-time from camera trap images and video footage. The integration of AI-driven models has also reduced manual labor, enabling conservationists to monitor wildlife populations and mitigate human-animal conflicts. Moreover, the use of edge computing and cloud-based AI solutions has made real-time detection feasible, further contributing to ecological research and conservation efforts. Despite these improvements, challenges remain, such as the need for larger and more diverse datasets, handling occlusions in images, and improving detection accuracy under varying environmental conditions. Additionally, computational resource limitations can hinder the deployment of deep learning models on low-power devices in remote locations.

References

[1] Bochkovskiy, A., Wang, C. Y., & Liao, H. Y. M. (2020), YOLOv4: Optimal Speed and Accuracy of Object Detection

[2] Ghosh, S., Ghosh, S., Ghosh, M., & Pal, T. (2020), Wildlife monitoring using deep learning: A review

[3] Singh, A., Sharma, N., & Dutta, M. K. (2020), Antipoaching surveillance using deep learning and computer vision

[4] Beery, S., Morris, D., & Yang, S. (2019), Efficient Pipeline for Camera Trap Image Review

[5] Tabak, M. A., Norouzzadeh, M. S., Wolfson, D. W., et al. (2019), Machine learning to classify animal species in camera trap images: Applications in ecology

[6] Bondi, E., Brantingham, P. J., & Tambe, M. (2018), Using AI for Wildlife Conservation

[7] Kellenberger, B., Marcos, D., & Tuia, D. (2018), Detecting mammals in UAV images: Best practices to address a substantially imbalanced dataset with deep learning

[8] Norouzzadeh, M. S., Nguyen, A., Kosmala, M., Swanson, A., Palmer, M. S., Packer, C., & Clune, J. (2018), Automatically identifying, counting, and describing wild animals in camera-trap images with deep learning

[9] Abadi, M., Barham, P., Chen, J., et al. (2016), TensorFlow: A system for large-scale machine learning

[10] Martín Abadi et al. (2016), TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems

[11] Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016), You Only Look Once: Unified, real-time object detection

[12] Linchant, J., Lisein, J., Semeki, J., Lejeune, P., & Vermeulen, C. (2015), Are unmanned aircraft systems (UASs) the future of wildlife monitoring? A review of accomplishments and challenges.

[13] Swanson, A., Kosmala, M., Lintott, C., & Packer, C. (2016), A generalized approach for producing, quantifying, and validating citizen science data from wildlife images.

[14] Burton, A. C., Neilson, E., Moreira, D., et al. (2015), Wildlife camera trapping: A review and recommendations for linking surveys to ecological processes.

[15] Kays, R., Crofoot, M. C., Jetz, W., & Wikelski, M. (2015), Terrestrial animal tracking as an eye on life and planet.