

Democratizing Multi-Species Object Detection in Drone Imagery for Cost-Effective Population Monitoring of Endangered Animals

Sankaran, Sowmya (School: Albuquerque Academy)

Animal populations worldwide are rapidly declining, and a low-cost, efficient technology that can accurately count endangered species could be vital for monitoring population changes over many years. This research had three main contributions: training a foundation model on drone images, creating a workflow that can be generalized for a wide range of species and models, and a low-cost hardware solution for parks worldwide. Developing a novel foundation model required following the entire machine learning workflow, from data collection, preparation, and augmentation to training, fine-tuning, and evaluating the models. Thousands of images taken using a Bwine F7 drone and large, openly available drone image datasets, with a wide variety of species such as cranes, zebras, dogs, and impalas, were used to train and fine-tune several machine learning models with a baseline YOLOv8 architecture. We demonstrate our workflow by training 30 different models, with the largest having 43.7 million parameters and 365 layers. We also used hyperparameter tuning and data augmentation techniques to improve accuracy. The workflow uses Pytorch in Google Colab on NVIDIA V100 and A100 GPUs, and a comprehensive testing process helped identify the best models and evaluate the workflow for easy and efficient use. The foundation model created using our workflow has high accuracy (up to 98.2%) on multi-species drone images, while the state-of-the-art YOLOv8 architecture had only 4% accuracy. Finally, a hardware solution was implemented on the Jetson Orin Nano for use in real-time species detection for park rangers. This study presents a workflow that is a low-cost, low-power, fast, and effective way for non-machine learning experts to create and deploy robust models in low-resource environments.