Classifying Drones and Birds from Images using Deep Neural Networks

Aaron Zhong, River Hill High School, aaronzhong3513@gmail.com

A Summer intern project at the Johns Hopkins University Applied Physics Laboratory

Abstract

With the increasing use of drones in both urban and rural environments, and the need to distinguish them from wildlife, particularly birds, object detection techniques using images have gained significant attention in recent years. Applications in security, wildlife monitoring, and airspace management have underscored the importance of accurately identifying and differentiating drones from birds. This study investigates the state-of-the-art image recognition technology to enable real-time classification of drones and birds. The primary objective of this research is to leverage machine learning and image processing techniques to develop automated systems capable of accurately identifying these aerial objects. Such advancements have the potential to significantly benefit industries including law enforcement, aviation safety, and environmental conservation. Drones often fly in close proximity to birds which poses unique challenges differentiating them.

One of the state-of-the-art object detection algorithms, YOLO (You Only Look Once) was chosen for its exceptional speed and accuracy. By training YOLO on the annotated datasets, the system was able to learn and distinguish the unique features of drones, such as their shapes and flight patterns, from those of birds. Extensive testing in a variety of environmental and temporal conditions have been conducted. These tests evaluated the model's ability to accurately classify drones and birds under realistic scenarios, ensuring its robustness and adaptability. The object detection model demonstrated promising results in distinguishing between drones and birds. Under controlled conditions, the YOLO-based system achieved an impressive accuracy rate of approximately 88% in correctly identifying drones and birds based on hundreds to thousands of test images and multiple videos. While accuracy declined slightly in more challenging scenarios, such as poor lighting or irregular movement, it remained robust at around 80%. The system excelled at handling overlapping flight patterns, effectively differentiating drones from birds of

similar size. While the current object detection model effectively differentiates between birds and drones, there are several opportunities to enhance its accuracy and efficiency further. One key improvement involves expanding the training dataset to include a more diverse range of bird and drone species, as well as varying environmental conditions, such as different lighting, weather, and backgrounds. By increasing the diversity of the training data, the model's adaptability can be significantly improved, enabling it to perform reliably in a wider array of real-world scenarios. These enhancements would make the system particularly valuable for applications such as security, where accurately detecting unauthorized drones is critical.