Environmental Computational Science
and Earth Observation Laboratory

Master thesis subject

Exploring 3D Rendering for Automated Wildlife Localisation with Deep Learning

Context

Wildlife censuses play a vital role in monitoring endangered species, such as rhinos. Such censuses are increasingly conducted by means of animal detection in aerial images, captured with drones. However, the large number of images that are needed bears a significant burden, due to the requirement for tediously drawn bounding boxes around the oftentimes minuscule and rare animals in the images. In this work, we attempt to solve this problem by means of automated animal detection through state-of-the-art deep learning models, trained on synthetically rendered animals.

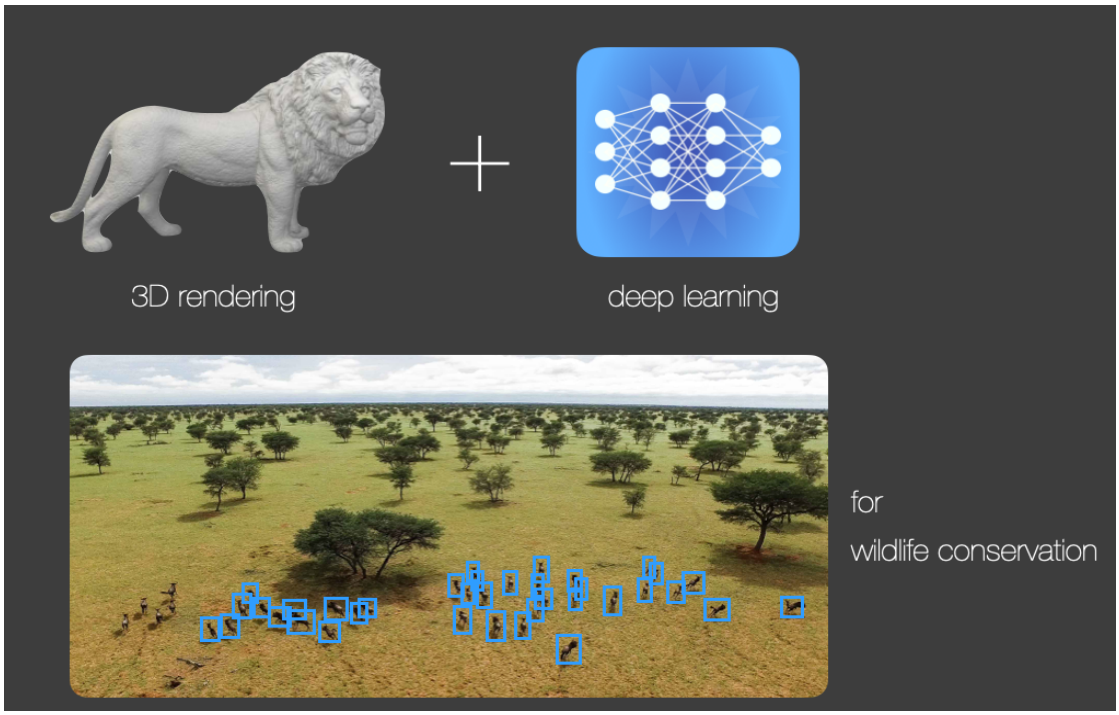

ECEO

EPFL ENAC IIE
Devis Tuia, Prof.
Rue de l'Industrie 17
Case Postale 440
CH – 1951 Sion

Téléphone : +41 21 69 382 83 (secr).
E-mail : devis.tuia@epfl.ch

With the advent of deep learning, in particular Convolutional Neural Networks (CNNs), machine vision tasks like object detection and classification have seen tremendous progress and rise in various fields, such as person identification or autonomous driving. In the context of wildlife conservation, CNNs have successfully been employed to detect large mammals, such as zebras or rhino, in images acquired from an aerial point of view, e.g. with drones. However, CNNs require a large amount of training data, which requires extensive amounts of work to create—basically, many hundreds of images need to be manually screened and bounding box annotations be drawn for thousands of animals therein.

In this project, the student will instead explore the possibilities of training CNNs for wildlife detection on artificially rendered data. In an initial step, this requires setting up a deep learning infrastructure and obtaining a baseline performance estimate on real, labelled images (provided). In a second step, the student will explore ways of rendering and positioning animals in the aerial image and training a CNN to detect it. This model is then to be applied on the real images. If successful, the end-result will be an animal detector that works sufficiently well on real images but requires no or very little manually annotated training images.

Objectives

- Familiarise and successfully set up a CNN-based object detector for animal localisation in UAV images
- Investigate the possibilities of using 3D rendering engines for realistic wildlife visualisations and detector training
- Provide an animal detector that is trained on synthetically rendered examples and can detect animals in real aerial images

Requirements and practical info

- Background in machine/deep learning is welcome.
- Programming skills in Python.
- The thesis will be supervised from the Sion campus.
- Access to parallel computing resources is provided.

Literature

- Kellenberger, B., Marcos, D. and Tuia, D. "Detecting mammals in UAV images: Best practices to address a substantially imbalanced dataset with deep learning." Remote Sensing of Environment 216 (2018): 139-153.
- Lin, T.-Yi, et al. "Focal loss for dense object detection." IEEE International Conference on Computer Vision, 2017.
- Zuffi, S., et al. "3D Menagerie: Modeling the 3D Shape and Pose of Animals." IEEE Conference on Computer Vision and Pattern Recognition, 2017.

Contact

Prof. Devis Tuia, devis.tuia@epfl.ch