SALTO White Paper

Secure Active Learning for Territorial Observations

ISSeP
UCLouvain
Oscars s.a.

October 2023

The SALTO project is part of the Defence-related Research Action (DEFRA) program, managed by the Royal Higher Institute for Defence (RHID) in collaboration with the Belgian Federal Science Policy Office (BELSPO). The SALTO consortium consists of three key partners: ISSeP, an expert in Earth observation, the PiLAB research group at UCLouvain, renowned for their focus on image processing and artificial intelligence; and Oscars s.a., a leading private firm that crafts IT solutions for data management.

Active Learning and synthetic models to better train models and human analysts

The influx of Earth observation images from satellites is expanding each year, outpacing the capacity of available human resources for detailed analysis. This situation is particularly critical for Belgian Defence Imagery Analysts who require prompt extraction of relevant information from newly received images. The SALTO research project seeks to address this issue by developing an automated analysis tool for highlighting main features on satellite images so that imagery analysts, whether civilian or military, are provided only with the information that are most relevant to their missions. The resulting tool will possess the capability to discern various types of aircraft and vehicles visible on a satellite image and to detect significant changes in airport areas, such as the position and the amount of these objects.

This tool includes the entire processing chain, starting from the loading of new data to generating the final report that highlights elements of interest to the imagery analyst. The automatic detection of these features relies on artificial intelligence, using an active learning model where the analyst monitors the detection process to enhance the reliability of the outcomes. This tool specializes in monitoring airports and identifying aircraft present there.

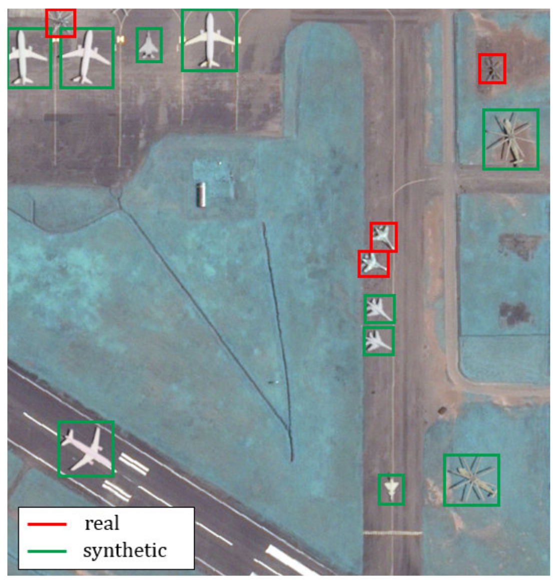

Figure 1 Real (red) and synthetic (green) aircraft on a Pléiades image.

To train the model effectively, various sets of civilian data were used. This involved working with images where the locations of aircraft were already known. Initially, an open-source image library was utilized. Afterwards, around fifty airport images obtained from Pléiades satellites were carefully annotated by hand. This created a collection of approximately 1,800 aircraft and 200 helicopters that served as training objects for the model.

Training on synthetic satellite data

The training of the model is limited by the amount of data available and by the type of aircrafts annotated on the open-source databases. One way to improve the quality and the quantity of the training datasets is to produce our own data.

Generation of synthetic satellite data

We developed a procedure to produce "fake" synthetic images using the Blender software. This free and open-source tool drivable with Python scripts allows to create a 3D scene on a real Pléiades raster background and to add in this scene models of aircrafts or any other type of objects (Figure 2). The colour of the objects is randomly chosen to increase the variety of object appearance. Shadows and "ray-tracing" rendering are calculated based on the Pléiades image metadata indicating the position of the Sun in the sky during the acquisition. The 3D model resolution is degraded to match the Pléiades image resolution and improve the quality of the synthetic image. At the end of the process, a transparent layer with aircraft "rasterized" objects is merged with the existing raster. Hence, synthetic images with aircrafts rarely present on real images are used to test the model and to train it better.

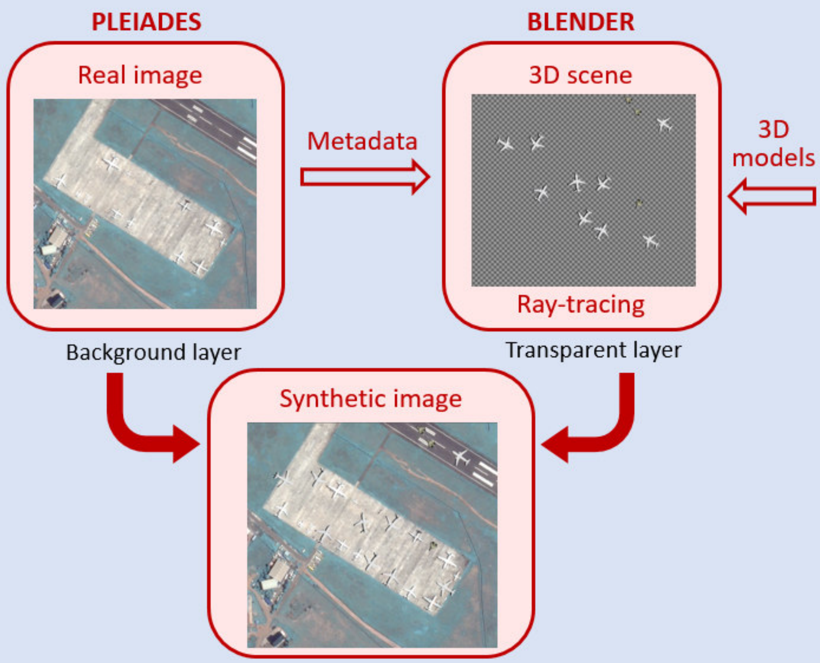

Figure 2 Scheme of the synthetic image generation

Notably, using synthetic images for training the model resulted in a noticeable 20% improvement in the model's precision, as measured by the mAP (mean Average Precision) metric (Figure 3). The generation of synthetic images also provide the flexibility to apprehend new original situations and let the model evolve accordingly.

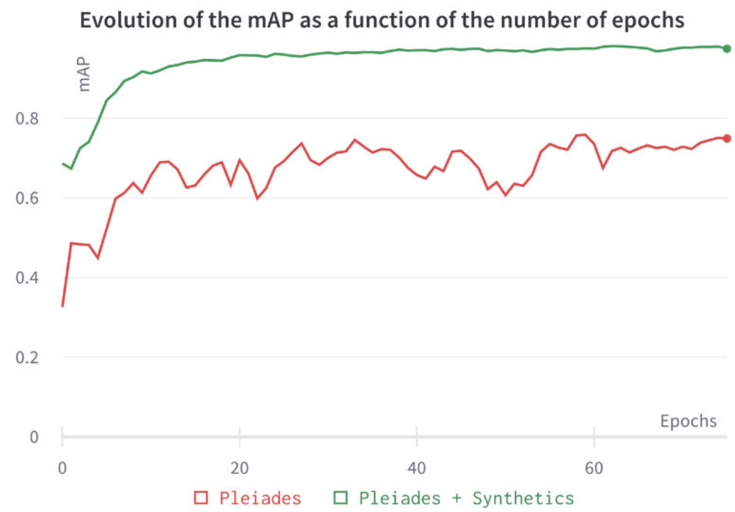

Figure 3 Evolution as a function of epochs (iterations) of the mean Average Precision (mAP) of the detection model after a training including Pléiades images only (red) and including Pléiades and synthetic images (green).

YOLO based detection and active learning

The detection model used in this project is YOLOv5, a convolutional neural network developed by Ultralytics[1]. It first extracts features in the image which are then used to localize and classify objects in the form of a bounding box as shown in Figure 4. For each detection, the model predicts the centre position, width and height of the bounding box, along with the object class and subclass, while simultaneously assigning a confidence score.

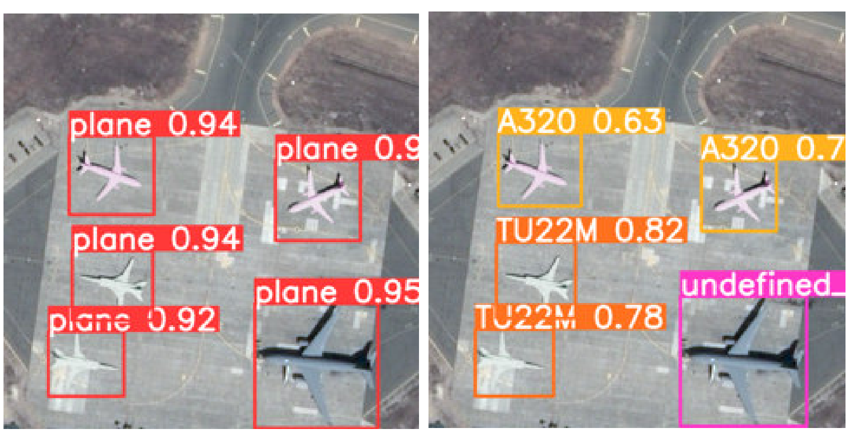

Figure 4 Example of the results of the detection model returning the bounding boxes of the detected objects, their classification (left) and subclassification (right), and a confidence score (maximal confidence = 1.00).

The model training is carried out in two steps. In the initial phase, the model is trained using an open-source dataset containing many annotated satellite images to acquire the basic knowledge for object detection. The parameters learned during this training serve as a starting point for the second training phase, which involves fine-tuning the model using datasets composed of the images captured by Pleiades satellites and synthetic

[1] https://github.com/ultralytics/yolov5

images. A training with and without active learning as well as different query strategies of active learning was tested.

Active learning

The principal objective of active learning is to select the most relevant images for model training in order to reduce the size of the dataset and thereby decrease the costs associated with annotation processes. As shown in Figure 5, active learning operates as an iterative workflow. The most interesting images are selected, annotated by experts and placed in the training set. This set is then used to train the model. These steps are repeated until the model reaches a predefined performance or the entire annotation budget is used.

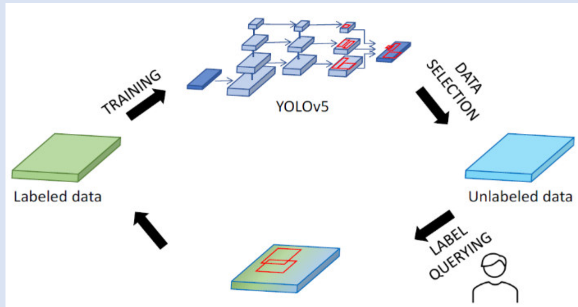

Figure 5 Active Learning scheme

There are two major strategies for querying the samples to label: uncertainty-based and diversity-based methods. Uncertainty-based query strategies select instances about which the model is the most uncertain while diversity-based strategies aim to select the most diverse samples based on the features extracted by the model. Other methods take advantage of both methods by combining them.

When tested, no particular strategy achieves significantly higher mAP compared to others. Thus, a mixed method was chosen as query strategy. This ensure that analysts are presented with the most ambiguous images, which likely represent the most challenging cases, while also showcasing a diverse range of images and avoiding redundancy.

Furthermore, this detection model has been integrated into a platform that encompasses all the training data. The platform enables annotation of the data and production of synthetic images. It generates reports that provide a comprehensive list of objects of interest and their respective characteristics according to the analyst's needs.

Beyond the primary scope of the project, the platform can be used as a training tool for the imagery analysts, in particular thanks to the generation of different scenarios through the production of synthetic data. Initially developed for the Belgian Defence Imagery Analysts and focusing on airports, the SALTO product would be useful in other diverse contexts for civilian users, such as emergency services, public authorities or humanitarian organisations.