Super accurate low latency object detection on a surveillance UAV

Maarten Vandersteegen, Kristof Van Beeck, Toon Goedemé
EAVISE, KU Leuven - Campus De Nayer

{maarten.vandersteegen,kristof.vanbeeck,toon.goedeme}@kuleuven.be

Abstract

Drones have proven to be useful in many industry segments such as security and surveillance, where e.g. on-board real-time object tracking is a necessity for autonomous flying guards. Tracking and following suspicious objects is therefore required in real-time on limited hardware. With an object detector in the loop, low latency becomes extremely important. In this paper, we propose a solution to make object detection for UAVs both fast and super accurate. We propose a multi-dataset learning strategy yielding top eye-sky object detection accuracy. Our model generalizes well on unseen data and can cope with different flying heights, optically zoomed-in shots and different viewing angles. We apply optimization steps such that we achieve minimal latency on embedded on-board hardware by fusing layers, quantizing calculations to 16-bit floats and 8-bit integers, with negligible loss in accuracy. We validate on NVIDIA's Jetson TX2 and Jetson Xavier platforms where we achieve a speed-wise performance boost of more than 10 $\times$.

1 Introduction

While UAVs are demonstrating their potential to offer support for numerous tasks in different industry sectors, there is a rising need for automating their control. On-board vision sensors combined with detection and tracking algorithms are a promising source of input for controlling sky-grade robots. In security and surveillance sectors for example, there is a need for detecting and tracking objects on the ground from hovering vehicles with immediate feedback to steering controls. If one for instance uses the output position of an object detector as input for a PID controller, the latency of the detector should be minimal. In these scenarios, processing video feeds from an on-board camera remotely is hardly doable due to streaming latencies and non-deterministic network throughput, such that on-board processing is highly recommended. Existing algorithms already provide powerful object detection, but often perform poorly for either ground level objects seen from higher altitudes or, if trained on a UAV dataset, perform poorly on frontal scenes. Next to that, state-of-the-art neural nets require heavy processing power, making them unusable for on-board inference.

In this work we propose an accurate low-latency object detector for on-board processing on a surveillance UAV. The algorithm detects common objects such as persons or cars and works well on unseen images from both frontal scenes and birds-eye views. To the best of our knowledge, we are first in achieving this due to the joint dataset training strategy where we train on a super set composed of a large generic common object dataset (MS COCO [25]) and a smaller drone dataset (Visdrone2018 [21]). We prove that this method is superior to training on a single dataset (MS COCO) and even to transfer learning from one dataset (MS COCO) to the other (Visdrone2018), leaving the latter approach behind by a margin of 3.5% mAP on Visdrone2018 and even 50% mAP on MS COCO.

The target platforms are NVIDIA Jetson TX2 and NVIDIA Xavier. While previous works run embedded object detectors/trackers on NVIDIA Jetson platforms [1, 2, 3], and stick to simple implementations, we go a number of steps further in optimizing our network for the target platform. We fuse layers and quantize calculations from floats to half floats (TX2) and even to 8-bit integers (Xavier). Our best model achieves a latency of 12ms which is more than 10 $\times$ faster than our baseline detector on the same hardware.

We made our extensively optimized implementation, based on TensorRT, publicly available ¹.

2 Related work

As a first step, we start with the selection of the most suited object detector algorithm for our UAV application, used as a baseline. Common network architectures for object detection include the (Fast(er)) R-CNN family [12] which go back to the very beginning of the deep-learned object detection area and solve the problem in a two stage manner. Attempts to solve object detection fast in one go result in SSD [13] and YOLO [14]. Many more improvements were made to the fundamental proposals from these detectors like the addition of Feature Pyramids [15], the selection of better backbones like ResNet [16], ResNeXT [17] and MobileNet [9]. These improvements result in networks as RetinaNet [18], Mask-RCNN [19], YOLOV3 [20] and several other Faster-RCNN and SSD derivatives. For our application we require an accurate object detector

¹code: <https://gitlab.com/EAVISE/jetnet>

Figure 1: We train a model on MS COCO + Visdrone2018 and port the trained model to TensorRT to compile it to an inference engine which is executed on a TX2 or Xavier mounted on a UAV. During compilation, layer fusion and weight quantization is applied. A calibration set is used to restore the accuracy when using 8-bit integers.

that can cope with both very small and large objects due to different flying heights and zoom levels. The network architecture still needs to be fast due to the low latency criteria and the limited hardware resources of our target platform. We selected YOLOV3 due to the excellent processing speed, the residual backbone for efficient training and the feature pyramid to support the high dynamic range in object sizes.

Nowadays, many large scale datasets are available for training neural networks. For the sky-grade object detection and tracking challenges for common objects in the wild, the most relevant available datasets are Visdrone2018 [21], SDD [22], UAV123 [23] and UAVDT [24]. Training an object detector to detect common objects in the wild is a difficult task because it requires a large amount of high quality training data. Due to the limited viewpoints and natural variation of most datasets, a network trained on dataset D is likely to perform well on D’s test set but not necessarily on the test set of dataset E. This so-called dataset bias [27] causes difficulties in designing object detectors for real-life applications. To overcome this problem we propose to train on a combination of a general and a task-specific dataset: MS COCO [25] and Visdrone2018 [21]. MS COCO is one of the few exceptions without a strong dataset bias and offers therefore the ideal training set to combine with. Visdrone2018 adds support for bird-eye views and steep viewing angles and is chosen because of its similar relevant class labels. In object detection, to the best of our knowledge, no other work explores joint training with multiple datasets.

Several works target embedded object detection or tracking on similar embedded hardware. Tijtgat et al. [1] compare YOLOV2 and tiny YOLOV2 among a few other classic object detection techniques both in terms of speed and accuracy on a Jetson TX2 platform for a UAV warning system. The same hardware is used by Blanco-Filgueira et al [3]. They propose an embedded real-time multi-object tracker based on a foreground-background detector and the GOTURN tracker. The predecessors of the TX2, the TX1 and TK1, are evaluated in [2] where they combine a tiny version of Faster-RCNN with a KCF tracker for single-object tracking on a drone. As stated before, we go further in optimizing

our network for the target platform. In addition, we make efforts to compress our trained network for faster inference.

Different techniques are explored in literature like computation quantization from floats to half floats and 8-bit integers [4, 26, 5], network pruning [6, 7, 8], architectural modifications [9, 10] and network distillation [11]. NVIDIA’s TensorRT toolkit [26] offers up to $4\times$ FLOPS reduction when quantizing the calculations from floats to 8-bit integers thanks to mixed-precision GPUs. This however requires special training or calibration to cope with coarse quantization errors. A state-of-the-art network pruning technique can suppress the number of calculations with an additional factor of up to $3\times$ [7] which is the case for simple networks like AlexNet and VGG. With more complex residual architectures, the state-of-the-art reports a FLOPS reduction of up to 42% without losing accuracy [6]. Pruning however requires numerous cycles of fine-tuning and increases the total training time with a large factor. One-time retrain pruning techniques [8] also exist, but typically result in much lower compression rates. MobileNet and SqueezeNet [9, 10] are conv-layer designs also targeted to reduce FLOPS and/or weights. These modifications however, make it harder to train a network which usually results in lower precision. Network distillation is also used as a compression technique to transfer the knowledge of a large network to a smaller network. Mehta et al. [11] report better performance with distillation compared to training a small network from scratch. Although the distilled network’s accuracy is 10% mAP lower than its teacher, the speedup compared to the teacher is almost $6\times$.

In our application, we strive for maximal reduction in FLOPS and minimal loss in accuracy. We therefore do not use MobileNet or network distillation. Network quantization seems the most beneficial.

3 Low latency object detection

In this section we discuss our methodology summarized in figure 1 and detail each of the steps: how we design (3.1) and train (3.2) an algorithm based on YOLOV3 and how we optimize it (3.3) for the target

Figure 2: Examples from the MS COCO (top) and Visdrone2018 (bottom) training set, with their modified annotations.

platform. We conclude this section by describing the hardware setup (3.4).

3.1 Network design

As stated earlier, we use YOLOV3 [20] as our base architecture. Standard YOLOV3 is designed to detect 80 object classes and supports a number of different square input resolutions (320x320, 416x416 or 608x608).

We make the following modifications:

1. Select a custom set of supported object classes: *person*, *car*, *bicycle*, *motorbike*, *bus* and *truck*
2. Modify the input resolution to 608×352 (16/9 aspect ratio matching our camera)
3. Recalculate the anchor boxes to best fit our training set

The object class labels chosen are relevant to our application and all appear in the two datasets we use for training. We reduce the number of output filters of the feature pyramid to support the six classes mentioned above.

3.2 Training

The detector is trained on a combination of two datasets: MS COCO [25] and Visdrone2018 [21]. MS COCO contains 120k training images with annotations

for 80 different common object classes, mainly focusing on ground level scenes. Visdrone2018 has a training set of 6.5k images and features 11 common object classes, seen from the sky from different angles and altitudes. Also, both datasets contain images with dense cluttered objects which are bounded in rectangular regions, marked with the *iscrowd* flag (MS COCO) or with the *ignored regions* class label (Visdrone2018).

For MS COCO, we remove all annotations not included in the supported class labels, listed in section 3.1, leaving all resulting unannotated images in our training set as additional negatives [27]. Next, leftover annotations marked with the *iscrowd* flag get a special *ignore* label. The loss function is modified to not penalize predictions inside these *ignore* annotation areas as false positives during training.

For Visdrone2018, we first remap all *people* and *pedestrian* labels to *person* and all *van* labels to *car* to match the supported class labels. Visdrone's labels *ignored regions*, *others* and arguably *tricycle* and *awning-tricycle* are remapped to the special *ignore* label. Figure 2 illustrates samples from the modified training data.

To show the added value of combining two datasets, we train three different models: a model only trained on Visdrone2018, a transfer-learned model initialised with MS COCO weights and fine-tuned on Visdrone2018 and finally a model trained on both datasets. For optimization purposes which are discussed in section 3.3.2, we also train a ReLU variant of the latter model, where all leaky ReLUs are replaced with standard ReLUs.

We use our own fork of the darknet framework² to train our models. Models trained or fine-tuned only on Visdrone2018 are trained for 100k iterations. Models trained on both datasets are trained for 500k iterations. Both are trained with a minibatch of 64 and a learning rate of 0.004 on four NVIDIA Tesla V100 GPUs. Training for 500k iterations takes up to 7 days. Before training, all models are initialized with ImageNet classifier weights. For the ReLU model variant, we first train an ImageNet darknet53 classifier using ReLUs to get representative extraction weights for initialization of the ReLU model. During training, the network input width $w \in 32\mathbb{N} : [416, 960]$ is randomly scaled every 10 iterations together with the network input height $h \in 32\mathbb{N} : [256, 544]$ while trying to fit the 16/9 aspect ratio as close as possible. Scaling the input resolution affects the scale of all layers.

3.3 Optimization

It is previously demonstrated that network pruning of residual networks like YOLOV3 results in low compression rates [8], even with complex pruning algorithms [6]. On top of that, a long iterative pruning process is needed to get a worthwhile compression rate,

²Training framework: <https://gitlab.com/EAVISE/darknet>

specifically on large datasets. We therefore limit additional network optimization to layer fusion and calculation quantization only.

3.3.1 Porting to TensorRT

Our framework is based on TensorRT [26], which we will use to accelerate our model for inference. The first step is to port our trained model to TensorRT, where it will be compiled to a fast inference engine for our specific hardware target (figure 1). During compilation, layer fusion and calculation quantization are applied which are discussed in sections 3.3.2 and 3.3.3.

TensorRT is a very powerful inference engine but offers only a limited set of supported layer types. YOLOV3 contains two layer types that are not supported: Leaky ReLU and the Upsample layer. The conventional way of dealing with unsupported layers in TensorRT is creating a plugin layer with a custom implementation. We do this for both the Leaky ReLU and the Upsample layer.

3.3.2 Layer fusion

When compiling the network to an inference engine, TensorRT first tries to fuse as many layers as possible. Typical convolution - batch norm - activation blocks can be fused to a single layer, greatly improving inference speed. However, plugin layers do not allow to be fused with other layers. Since our Leaky ReLU is a plugin, it will not be fused with its proceeding batch norm/convolution layer. YOLOV3 contains 72 leaky ReLUs which may lead to a significant performance drop compared to using classic ReLUs. As an alternative we therefore create a second implementation replacing the leaky ReLUs with ReLUs. This however requires retraining of the network and might result in lower accuracy [28].

3.3.3 Quantization

Next to layer fusion, one can choose to quantize layer calculations to half floats and 8-bit integers. All native layers can be quantised. Plugin layers however will still be executed using floats regardless of the quantization level which forces TensorRT to insert type conversions before and after every plugin layer. As stated in section 3.3.2, our Leaky ReLU plugin is repeated 72 times which introduces 144 tensor type conversions. This becomes even more problematic than the missed fusion opportunities. The alternative of replacing all leaky ReLUs with ReLUs again solves this problem, but as stated before might not be as accurate. We propose a second alternative to implement a Leaky ReLU with four natively supported layers as depicted in Figure 3. This implementation is slower than a plugin implementation when using floats, but allows to be quantized which may result in better performance in half float

Figure 3: Leaky ReLU with negative slope α implemented with two scale layers, a ReLU layer and an element-wise adder. The first scale layer is likely to be fused into the previous layer.

and 8-bit integer mode compared to its plugin alternative. Furthermore, this leaky ReLU can benefit slightly from layer fusion since the first scale layer is likely to be fused to the previous layer (still leaving three unfused layers).

No additional precautions need to be made when converting the model to half float (16-bit floats) precision. For 8-bit integers however, special training [5] or weight calibration [26] is required. In this work we focus on calibration only, as depicted in figure 1, since it is much faster. We use TensorRT’s built-in entropy calibrator [26] which calculates upper and lower boundaries of feature map activations for each layer. These boundaries are determined by minimizing the information loss between the quantized activation values and the original activation values (based on the Kullback-Leibler divergence) when sending numerous images from the training set through the network. Once the boundaries are calculated, the lower boundary will represent -128 while the upper boundary will represent 127. All values outside the boundaries are clipped. The image set used in calibration is called the calibration set. We use a calibration set of 1000 images, randomly sampled from our training set.

3.4 Hardware setup

Our main target platform is the NVIDIA Jetson TX2. The TX2 has an onboard Pascal GPU with 256 CUDA cores, which supports float and half float calculations. We use TensorRT V4.0 on this platform. We also investigate performance on the Jetson Xavier, NVIDIA’s latest platform, featuring two TPUs and an onboard Volta GPU with 512 CUDA cores which supports floats, half floats and 8-bit integer calculations. Our Xavier is running TensorRT V5.0 and we will only be using its GPU. All experiments are done with maximum compute power enabled on both platforms.

4 Results

4.1 Model accuracy

Our trained models are evaluated on the validation sets of MS COCO and Visdrone2018. The evaluation

Table 1: Accuracy results of our models (%mAP), validated on MS COCO and Visdrone2018 validation sets respectively.

model	MS COCO	Visdrone2018
coco_608	67.83%	24.94%
visdrone_608	4.36%	41.45%
visdrone_coco_init_608	10.40%	41.72%
joint_608	60.56%	45.27%
joint_960	63.48%	56.78%
joint_relu_608	62.06%	44.66%
joint_relu_960	63.91%	56.09%

metric used is mean average-precision for an IoU of 0.5. Table 1 presents the validation results. The name of each model represents the dataset its trained on while the trailing number indicates the input width. The default YOLOV3 model, *coco_608*, is a little better than other models on MS COCO, but performs poorly on Visdrone2018. The contrast is even larger for models *visdrone_608* and *visdrone_coco_init_608* which do quite well on Visdrone2018 but fail completely on MS COCO. This might indeed indicate that MS COCO has a weaker dataset bias compared to Visdrone2018. Model *visdrone_coco_init_608* differs from *visdrone_608* in that it is transfer-learned from MS COCO to Visdrone2018 by initializing with weights from *coco_608* prior to training on Visdrone2018. We propose *joint_608* (trained jointly on both datasets) as the best alternative since it performs well on both validation sets and even outperforms *visdrone_608* and *visdrone_coco_init_608* by at least 3.5% on Visdrone2018 and by more than 50% on MS COCO. This proves that training on a joint set is superior compared to training on a single dataset and even to transfer learning from one dataset to the other.

We also tested *joint_960* which uses the same weights as *joint_608* but has an input resolution of 960×544 . This model outperforms *joint_608* by a large margin of 11% for Visdrone2018, which can be explained by the fact that Visdrone2018 contains very small instances which are better detected in a higher resolution. We only benchmark the latency of *joint_960* on the Xavier platform since it is too heavy for the TX2. As mentioned in section 3.2, we also train a standard ReLU variant of our *joint* model, coined *joint_relu*. As expected, these models perform slightly worse (on Visdrone2018) compared to there leaky ReLU variants, but the difference is less than 1%. Figure 4 gives a qualitative overview of the detection performance on unseen data.

4.2 Optimization

Since we are interested in low-latency rather than throughput, we only benchmark our models with batch size equal to 1. Figure 5 presents our optimization re-

sults in accuracy (mAP) versus latency (ms) charts. All measurements are GPU inference only (excluding pre- and postprocessing). Label *leakyA* represent optimized models using leaky ReLU plugin layers, while *leakyB* represents optimized models using our alternative leaky ReLU depicted in figure 3. Benchmarks labeled with *relu* represent optimizations with standard ReLU. The blue dots represent the baseline (unoptimized) models, ran on the same hardware using our training framework with all GPU optimizations enabled.

On the TX2, the *leakyB* optimization is 15% faster than the *leakyA* optimization in half float mode (FP16), but still lacks behind our fastest TX2 model (*relu* half float) which finishes in 79ms, outperforming its baseline by 4.2 $\times$.

On the Xavier platform, *leakyB* is only 6% faster in 8-bit integer mode compared to the plugin version. The fastest model finishes inference in about 12ms with a 608 input width and in 26ms with a 960 input width, boosting inference speed 11 $\times$ compared to the baseline model.

Note that the accuracy drop of the optimized models compared to their baseline is negligible, even for the 8-bit quantized models.

5 Conclusion

In this paper, we proposed a methodology to make object detection in UAV imagery both fast and super accurate. We achieved this by introducing a multi-dataset learning strategy and demonstrated that this methodology significantly outperforms transfer-learning by 3.5% mAP on Visdrone2018 and by 50% mAP on MS COCO, resulting in a generic model usable for all UAV operations. Furthermore we applied extensive optimization steps to achieve minimal latency on our embedded target platforms, the NVIDIA Jetson TX2 and Xavier. By fusing layers and quantizing calculations, we achieve a speed-wise performance boost of 4.2 $\times$ on the TX2 and 11 $\times$ on the Xavier hardware without giving in on accuracy.

In future work, we plan to integrate our model with on-board trackers.

6 Acknowledgement

This work is partially supported by the 3DSafeGuard and Mirador project.

References

- [1] Tijtgat et al. “Embedded real-time object detection for a UAV warning system.” *ICCV*, 2017. p. 2110-2118.
- [2] Han et al. “Deep Drone: object detection and tracking for smart drones on embedded system.” *Stanford University* 2015
- [3] Blanco-Filgueira et al. “Deep Learning-Based Multiple Object Visual Tracking on Embedded System for

Figure 4: Detection results of the *coco_608*, *visdrone_coco_init_608* and *joint_608* models on validation images of MS COCO (top left), Visdrone2018 (top right), UAVDT (bottom left) [24] and our own (bottom right).

Figure 5: Accuracy vs. latency results on the TX2 and Xavier platforms

- IoT and Mobile Edge Computing Applications.” *IEEE IOTJ*, 2019.
- [4] Gao et al. “IFQ-Net: Integrated Fixed-point Quantization Networks for Embedded Vision.” *CVPR*, 2018. p. 607-615.
 - [5] Gong et al. “Highly Efficient 8-bit Low Precision Inference of Convolutional Neural Networks with Intel-Caffe.” *ACM ReQuEST*, 2018. p. 2.
 - [6] He et al. “Progressive Deep Neural Networks Acceleration via Soft Filter Pruning.” *IJCAI*, 2018. p. 2234-2240.
 - [7] Molchanov et al. “Pruning Convolutional Neural Networks for Resource Efficient Inference” *ICLR*, 2017
 - [8] Li et al. “Pruning filters for efficient convnets” *ICLR*, 2017
 - [9] Howard et al. “MobileNets: Efficient convolutional neural networks for mobile vision applications.” *arXiv:1704.04861*, 2017
 - [10] Iandola et al. “SqueezeNet: Alexnet-level accuracy with 50x fewer parameters and < 0.5 mb model size.” *arXiv:1602.07360*, 2016
 - [11] Mehta et al. “Object detection at 200 Frames Per Second.” *ECCV*, Springer, 2018. p. 659-675
 - [12] Ren et al. “Faster R-CNN: Towards real-time object detection with region proposal networks.” *NIPS*, 2015. p. 91-99.
 - [13] Liu et al. “SSD: Single shot multibox detector.” *ECCV*, Springer, 2016. p. 21-37.
 - [14] Redmon et al. “You Only Look Once: Unified, real-time object detection.” *CVPR*, 2016. p. 779-788.
 - [15] Lin et al. “Feature Pyramid Networks for Object Detection.” *CVPR*, 2017. p. 2117-2125.
 - [16] He et al. “Deep residual learning for image recognition.” *CVPR*, 2016. p. 770-778.
 - [17] Xie et al. “Aggregated residual transformations for deep neural networks.” *CVPR*, 2017. p. 1492-1500.
 - [18] Lin et al. “Focal loss for dense object detection.” *ICCV*, 2017. p. 2980-2988.
 - [19] He et al. “Mask R-CNN.” *ICCV*, 2017. p. 2961-2969
 - [20] Redmon et al. “YOLOV3: An incremental improvement.” *arXiv:1804.02767*, 2018
 - [21] Pengfei et al. “Vision Meets Drones: A Challenge.” *arXiv:1804.07437*, 2018
 - [22] Robicquet et al. “Learning Social Etiquette: Human Trajectory Prediction In Crowded Scenes” *ECCV*, Springer 2016. p. 549-565.
 - [23] UAV123 Dataset <https://ivul.kaust.edu.sa/Pages/Dataset-UAV123.aspx>
 - [24] Du et al. “The Unmanned Aerial Vehicle Benchmark: Object Detection and Tracking.” *ECCV*, 2018. p. 370-386.
 - [25] Lin et al. “Microsoft coco: Common objects in context.” *ECCV*. Springer 2014. p. 740-755.
 - [26] 8-bit inference with TensorRT <http://on-demand.gputechconf.com/gtc/2017/presentation/s7310-8-bit-inference-with-tensorrt.pdf>
 - [27] Torralba et al. “Unbiased look at dataset bias.” *CVPR*, 2011, Vol 1, p. 1521-1528
 - [28] Xu et al. “Empirical Evaluation of Rectified Activations in Convolutional Network” *arXiv:1505.00853*, 2015