Runway Recognition

Michael Hardt
mwhardt@stanford.edu

Carlos Querejeta
carlosq@stanford.edu

Grzegorz Kawiecki
kawiecki@stanford.edu

https://youtu.be/bKq6iCp_MAw

1. Introduction

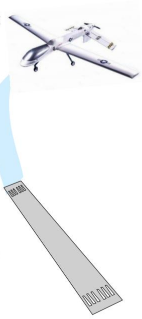

- Autonomous aircraft will require redundant and robust solutions for navigation.
- Without GPS or other aids, need new tools to rely upon its camera for landing.
- Assume the runway's position is known, but the aircraft doesn't know its position well.
- **Question: Is the runway in view in the forward-looking camera?**
- **If so, how large is the runway, what is its relative angle within the image?**
- *Use Deep Learning to answer.* From camera calibration and barometric altitude, one can determine aircraft position relative to the runway and then plan for landing.

2. Data

- Two data sources: Google Earth and osgEarth.
- Google Earth images are high quality, but need manual preparation.
- osgEarth transforms satellite orthoimagery from free, online databases to any camera location and perspective which we automated.
- 16 single runway airports selected, 1186 random samples of camera positions: 83% viewing runway, 17% negative samples without runway.
- Pixel masks calculated.
- 1024x768 resolution scaled and padded to 416x416.
- Dataset split: 80%-10%-10% for training / dev / test sets.

3. Methods

- Yolo (v3) model
- 106-layer fully convolutional architecture
- Multilabel classification
- Residual blocks & skip connections for bigger network
- Parallel object detections on 3 scales
- 416 x 416 native resolution
- Tiny-Yolo also tested
- Enhanced detection of small objects

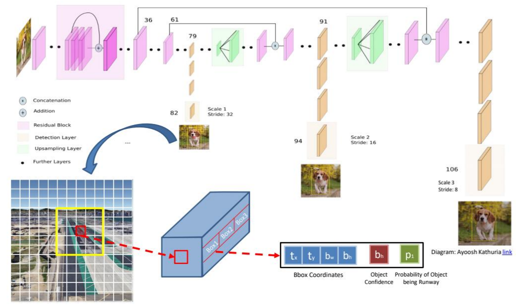

4. Results

- Successful application of Transfer Learning of Yolo and Tiny-Yolo pretrained on COCO dataset to runway recognition task.
- Hyperparameter search conducted with following results:
  - number of trained layers: **all**
  - learning rate: **10^{-4}**
  - batch size: **8**
- Excellent runway detection capability: **mAP = 0.94** on test set. Even detected unexpectedly multiple runways which were not initially labeled, and some runways at a large distance that even a human can't easily recognize.

- Tiny-Yolo also did well with a **mAP = 0.72** with only about a third of the number of layers of the full Yolo.
- Semantic segmentation attempted which classifies each pixel in image and indicates shape, but no conclusive results obtained.

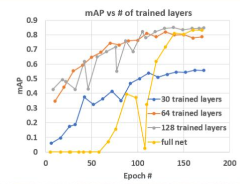

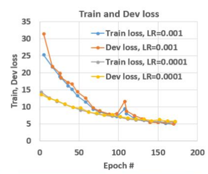

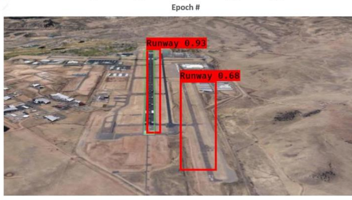

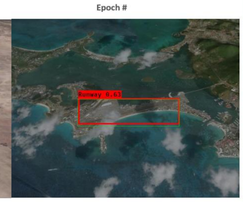

5. Conclusions

- Yolo achieved excellent, promising results. Its detection capability in some cases beat the human eye!
- Now need to validate model on real flight imagery and check its sensitivity to varying conditions, i.e. lighting, terrain, distance.
- **How does Yolo answer the initial question? Yolo tells us first if the runway is in the image.**
- Then, in all but very few cases when the aircraft is already perfectly aligned with the runway, the runway appears as a diagonal of the bounding box. **So the diagonal length gives the runway length in pixels!**

- Still need the runway orientation. Our current plan is to extend Yolo with a binary variable or probability which indicates whether runway appears as left or right diagonal.