Find the Polar Bear

Objective

Out of a set of pictures from Svalbard (Spitsbergen) provided by Dominique Ueltschi, the objective was, to find those pictures including a polar bear. Furthermore, to label the bear and give his location within the picture.

System choice

There are various object detection systems present at the moment, following different approaches. Given the course is deep learning, we focused on those techniques. Further, we decided not to follow any guides to build our own detection system, but to train one of the currently leading systems to our purpose.

Decision: YOLOv3

Reasons:
- good performance on different platforms
- simple to retrain to our needs

A second approach on detectron[1] was abandoned due to the need of GPU support.

Leading object detection systems

- Region Proposals (R-CNN, Fast R-CNN,Faster R-CNN)
- Single Shot MultiBox Detector (SSD)
- You Only Look Once (YOLO)

How does YOLOv3 work?

YOLOv3 is a real-time object detection system that applies a *single neural network* to an image. The network divides the image into regions and *predicts bounding boxes and probabilities* for each region. These bounding boxes are weighted by the predicted probabilities. Each box predicts the classes the box may contain using multilabel classification.

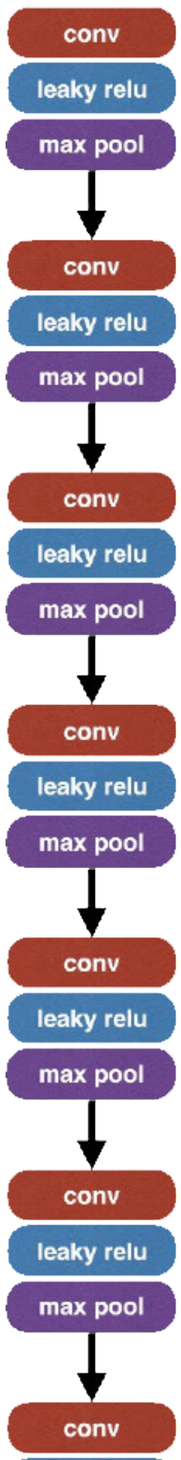

YOLO Architecture

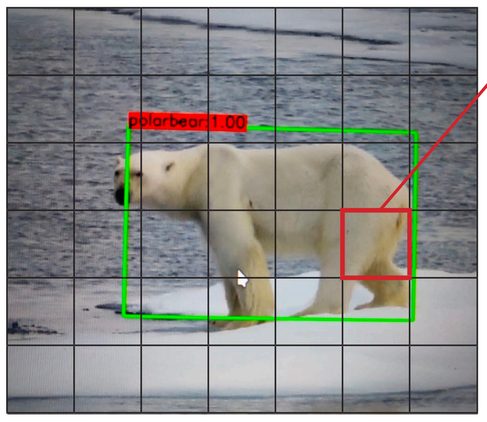

bounding boxes | class prediction

1 - x-coordinates of the bounding box
2 - y-coordinates of the bounding box
3 - width of the bounding box
4 - height of the bounding box
5 - bounding box confidence
6-85 - class prediction

To solve the problem of multiple objects within one grid, YOLOv3 uses the concept of anchor boxes. This adds one more "dimension" to the output labels by pre-defining a number of anchor boxes. Therefore to be able to assign one object to each anchor box.

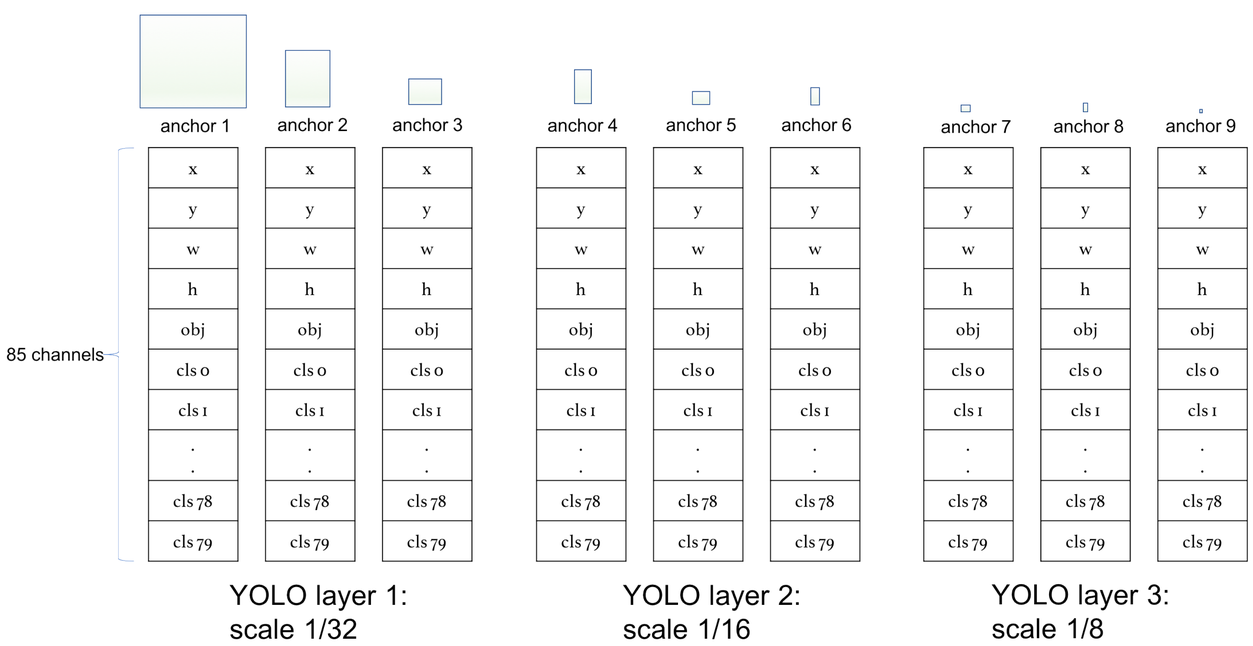

Training & Testing Data[2]

Training Pictures: 441
Testing Pictures: 49

Training iterations: 5200

Data augmentation:
```
angle=0
saturation=1.5
exposure=1.5 hue=.1
```

Other parameters:
```
momentum=0.9
decay=0.0005
learning_rate=0.001
policy=steps
steps=3800
scales=.1
burn_in=400
```

Training loss:

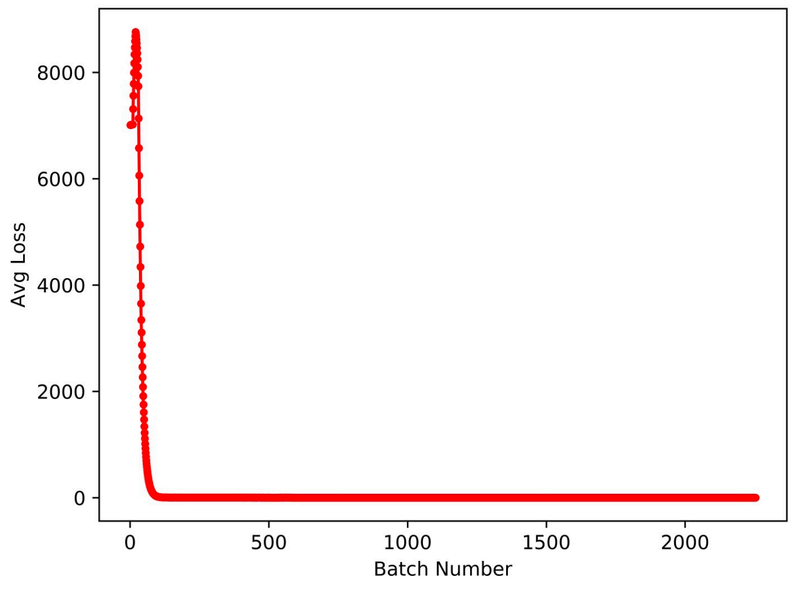

Training YOLOv3

YOLOv3 comes with a set of pretrained weights and configuration files, for different purposes. None of them was able to detect the polar bears on our photos as such. Out-of-the-box YOLOv3 doesn't give any option to improve those weights to add new classes. Therefore, we prepared a set of pictures to train the network on polar bears.

Results and Challenges

Testing on our final trained weights YOLO is able to correctly label a polar bear. Although it still struggles on pictures from far away or from behind.

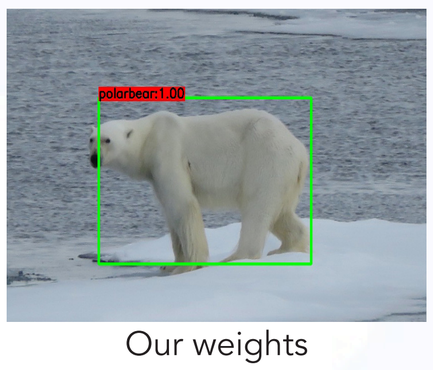

Our weights

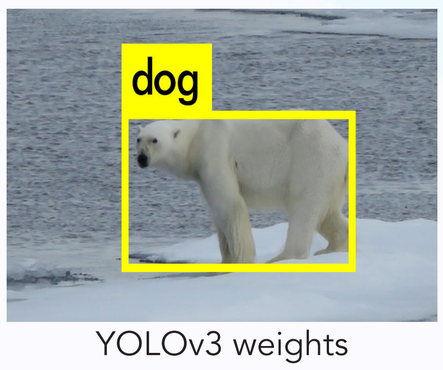

YOLOv3 weights

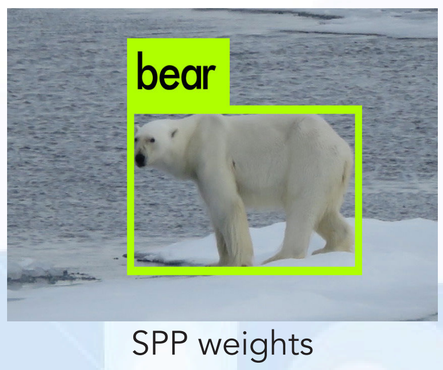

SPP weights

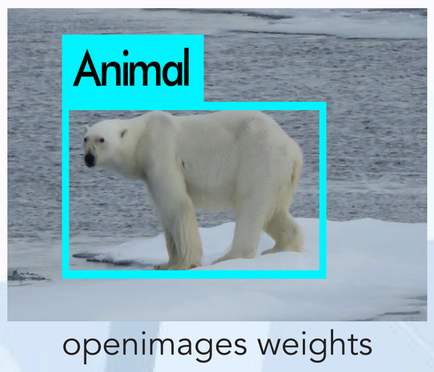

openimages weights

Lime predictor Results

Top 1 Featrue | Top 10 Featrues | Top 5 Classes

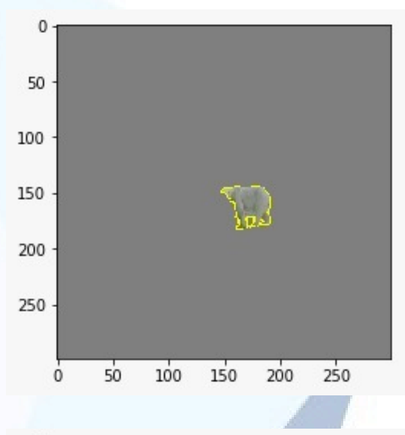

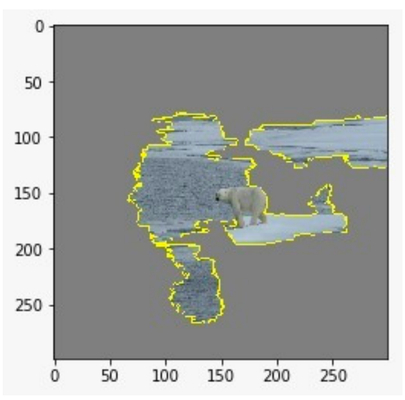

```
295 brown bear 6.408
208 golden retriever
    6.453
843 swimming trunks
    0.001
209 Labrador retriever
    0.001
297 polar bear 0.998
```

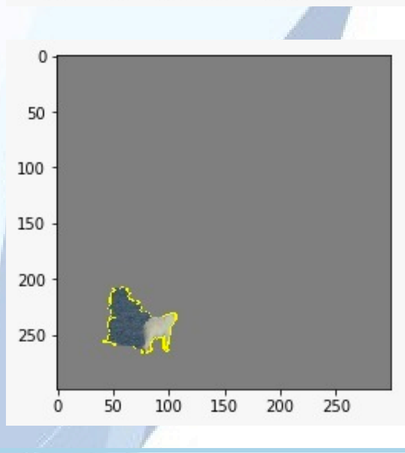

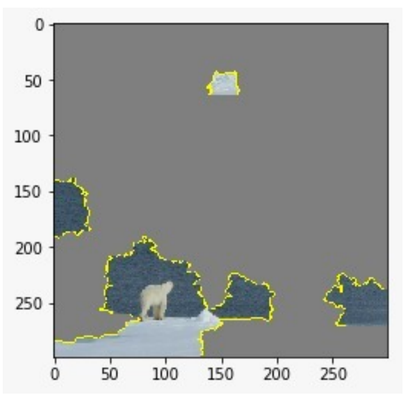

```
280 arctic fox, 0.012
259 samoyed 0.015
273 coyote 0.018
271 white wolf 0.149
297 polar bear 0.723
```

Project team: Dominique Ueltschi & Isabelle Kluser

1 https://github.com/facebookresearch/Detectron/tree/master/detectron
2 https://www.learnopencv.com/training-yolov3-deep-learning-based-custom-object-detector/