

# Bristol DAQ ML Studies

---

Bristol DUNE Meeting 16.2.20

Joel Greer + Raul Stein

# Contents

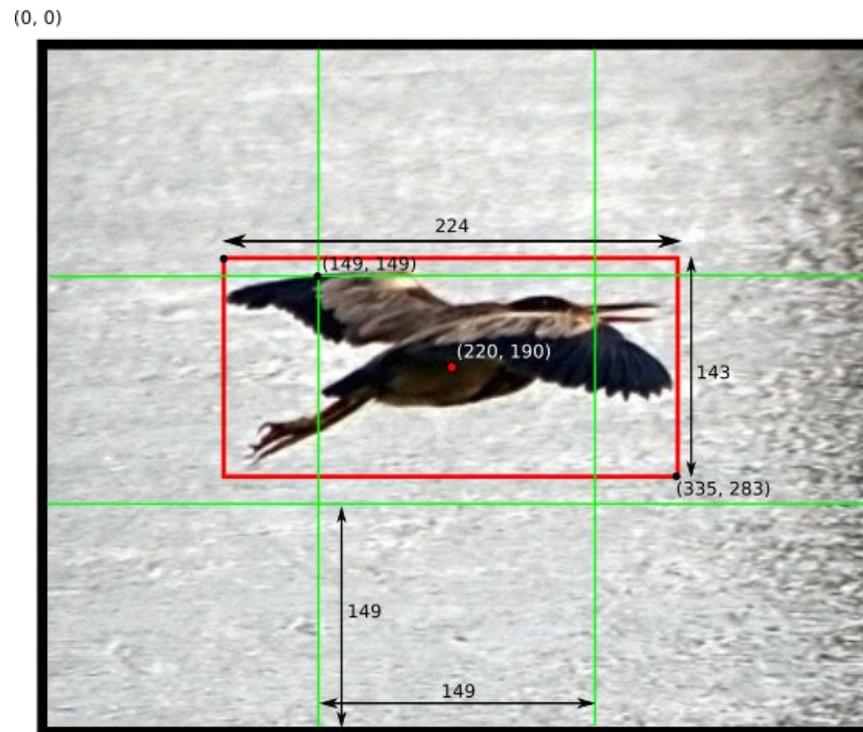
---

1. YOLO network
  - a. Overview
  - b. Thresholds
  - c. Loss
  - d. Truth
  - e. Metrics
2. Investigation of what the network is learning as it trains
  - a. Raccoons
  - b. SNNu

# YOLOv3 Network - overview

You Only Look Once (YOLO) network predicts both the class and localization of features it identifies in an image.

- YOLOv3 looks at the image at 3 different scales
- For each scale the network divides an image into  $S \times S$  grid cells (15x15, 30x30, 60x60 in our cases)
- Each grid cell makes 3 prediction boxes, each guided by a reference box for that scale
- Predictions have  $(x, y, w, h, \text{obj\_score}, \text{class\_scores})$
- Predictions are thresholded based on the class score (we have only 1 class, so we use the object score instead)
- Then the remaining predictions from all 3 scales are recombined to get rid of overlapping predictions
- The remaining predictions are the network outputs



$$x = (220 - 149) / 149 = 0.48$$

$$y = (190 - 149) / 149 = 0.28$$

$$w = 224 / 448 = 0.50$$

$$h = 143 / 448 = 0.32$$

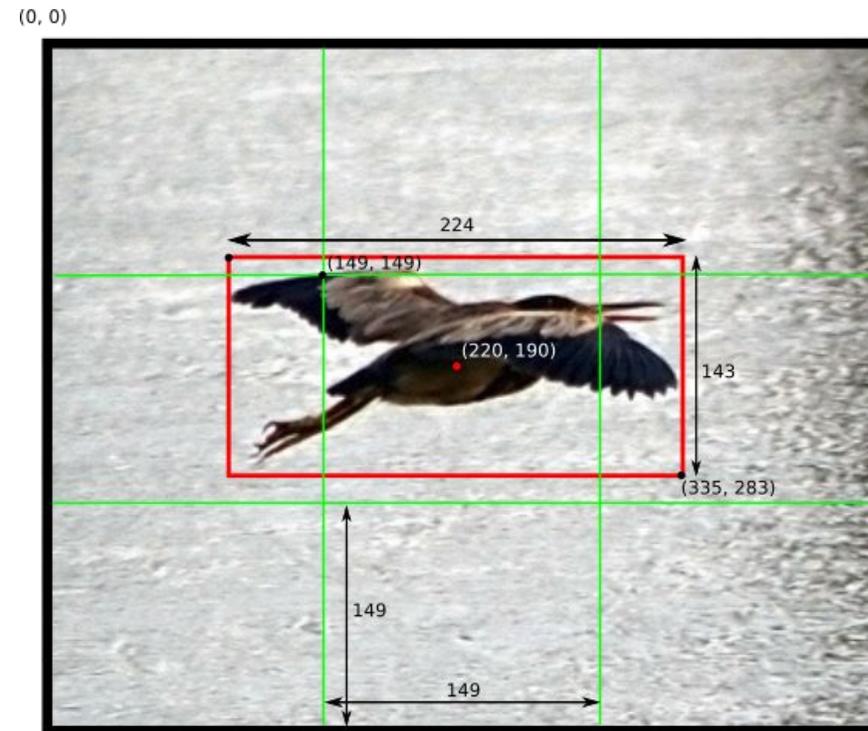
# YOLOv3 Network - thresholds

Given a trained model we therefore have 2 discriminators which we can toggle to affect the performance of our network:

1. The object score threshold
2. The algorithm used to discard remaining predictions which overlap

The algorithm for discarding remaining predictions:

- Get prediction with highest object score
- Compute its intersection over union (IOU) with the other predicted boxes
- Discard boxes which have an IOU larger than a chosen threshold
- Repeat for the remaining box with the next largest object score



(447, 447)

$$x = (220-149) / 149 = 0.48$$
$$y = (190-149) / 149 = 0.28$$
$$w = 224 / 448 = 0.50$$
$$h = 143 / 448 = 0.32$$

# YOLOv3 Network - loss

During training we minimise the loss, which has contributions from 4 different aspects the network is trying to optimise

1. Bbox x,y positional term, only uses the bbox prediction with maximal IOU, and is only non-zero if an object is in the cell
2. BBox w,h term, stronger contribution from larger bboxes for equivalent w,h difference to truth
3. Class prediction term, the square of Confidence - IOU
4. Classification loss term, only contributes when an object is present in the cell

S: number of cells along one axis

B: number of bboxes per cell

$$\lambda_{coord} \sum_{i=0}^{S^2} \sum_{j=0}^B \mathbb{1}_{ij}^{obj} (x_i - \hat{x}_i)^2 + (y_i - \hat{y}_i)^2$$

$$\lambda_{coord} \sum_{i=0}^{S^2} \sum_{j=0}^B \mathbb{1}_{ij}^{obj} (\sqrt{w_i} - \sqrt{\hat{w}_i})^2 + (\sqrt{h_i} - \sqrt{\hat{h}_i})^2$$

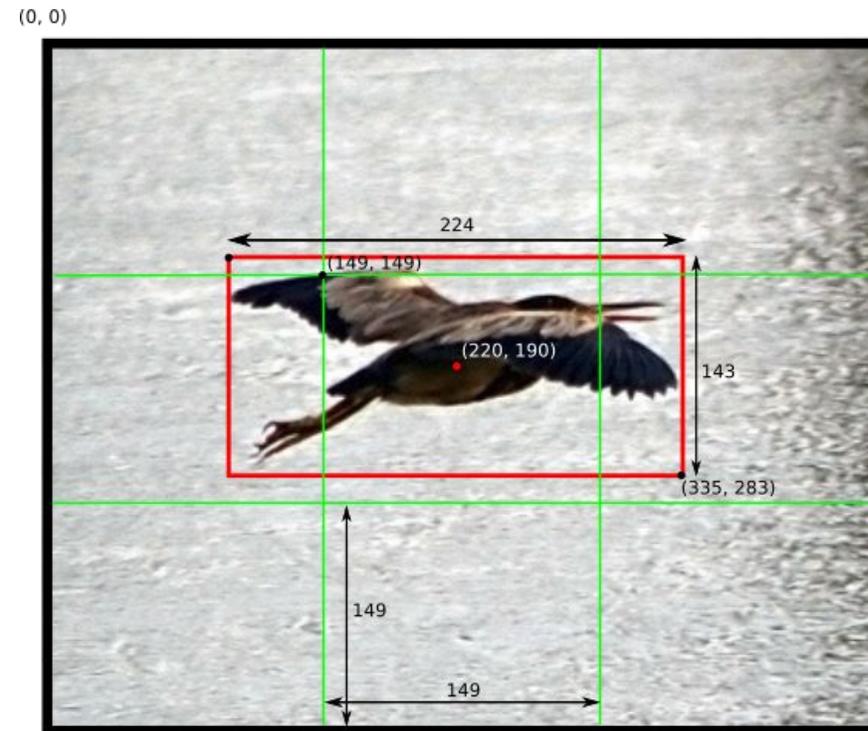
$$\sum_{i=0}^{S^2} \sum_{j=0}^B \mathbb{1}_{ij}^{obj} (C_i - \hat{C}_i)^2 + \lambda_{noobj} \sum_{i=0}^{S^2} \sum_{j=0}^B \mathbb{1}_{ij}^{noobj} (C_i - \hat{C}_i)^2$$

$$\sum_{i=0}^{S^2} \mathbb{1}_i^{obj} \sum_{c \in classes} (p_i(c) - \hat{p}_i(c))^2$$

# YOLOv3 Network - truth

A truth box is a box containing information about the location of an object in an image which the network should learn to identify.

- A truth box containing a truth object is assigned to 1 scale based on which one of the reference boxes (anchors) is closest in size to it.
- It is also assigned to the grid cell at that scale which contains the centre of the truth box
- The object score and IOU with the truth box for the prediction box produced by this grid cell anchor is what the loss is designed to improve during training.
- Prediction boxes made by grid cell anchors with no associated truth box should contribute in a positive way to the loss function if they have low objectness scores and a low IOU with truth boxes



(447, 447)

$$x = (220 - 149) / 149 = 0.48$$

$$y = (190 - 149) / 149 = 0.28$$

$$w = 224 / 448 = 0.50$$

$$h = 143 / 448 = 0.32$$

# YOLOv3 Network - metrics

---

We want to understand:

1. Whether the performance of the network is improving during training and, if not, why
2. At which epoch was the network performing best
3. What thresholds on object score and IOU of predicted boxes produce the best performance

To understand whether performance of the network is improving during training, we can look at the average object scores for:

- predictions which are made using an anchor in a grid cell which has an associated truth box
- predictions made using an anchor in a grid cell which has no associated truth box

We can also look at the distribution of object scores for these two cases

# YOLOv3 Network - metrics

---

We can also look at the IOU with truth boxes of these 2 cases. Note that this is a different concept than the IOU of predicted boxes with one another!

We can use the IOU of predicted boxes with truth boxes to evaluate whether the network learns to predict boxes with the right size and locality. We can look at the distribution and might want to set a minimum threshold on this when evaluating false positive rate, recall and precision.

In future we could potentially determine the % of the SADC of the signal from a SNNu within the prediction box.

In order to determine good object score and IOU score thresholds, we need to look at the recall, precision and false positive rate of the network

- Recall: proportion of truth boxes which are identified
- Precision: proportion of positive predictions which are true positives

# Network Training Investigation

---

We are currently trying to understand what is happening during training. We are still looking only at metrics computed from a single batch (8 images) in this presentation.

We expect the object scores from predictions with associated truth boxes to be higher than for predictions without associated truth boxes, and would naively expect the separation between the two distributions to increase as the network trains and its loss is minimized.

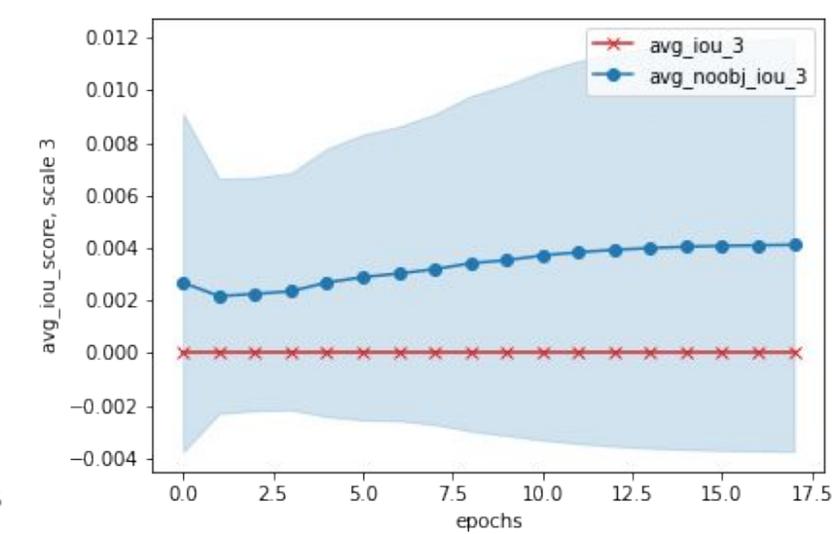
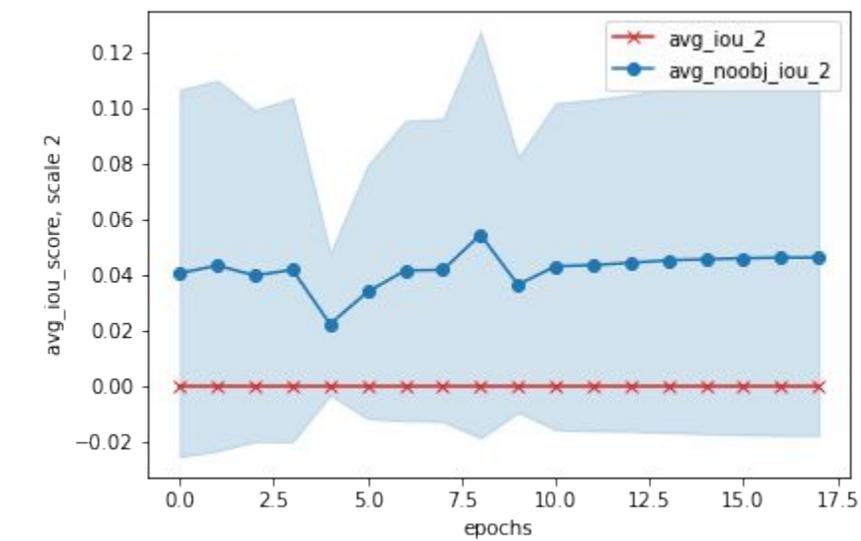
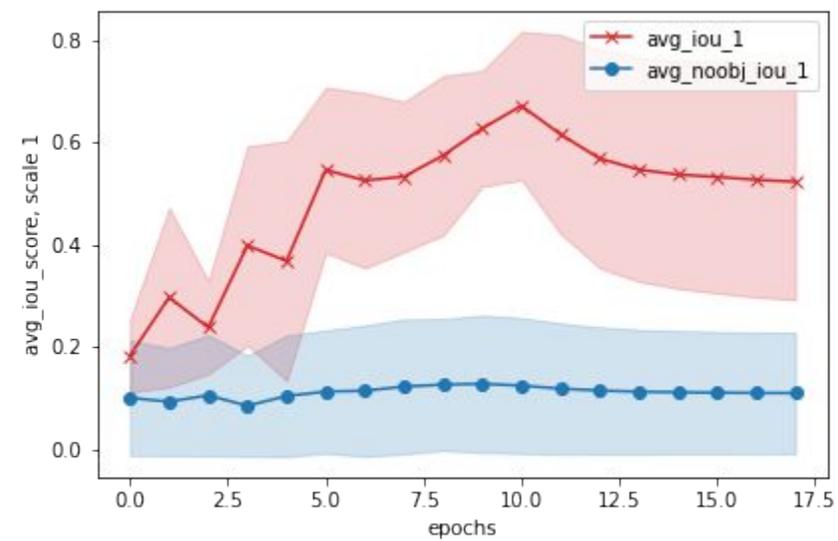
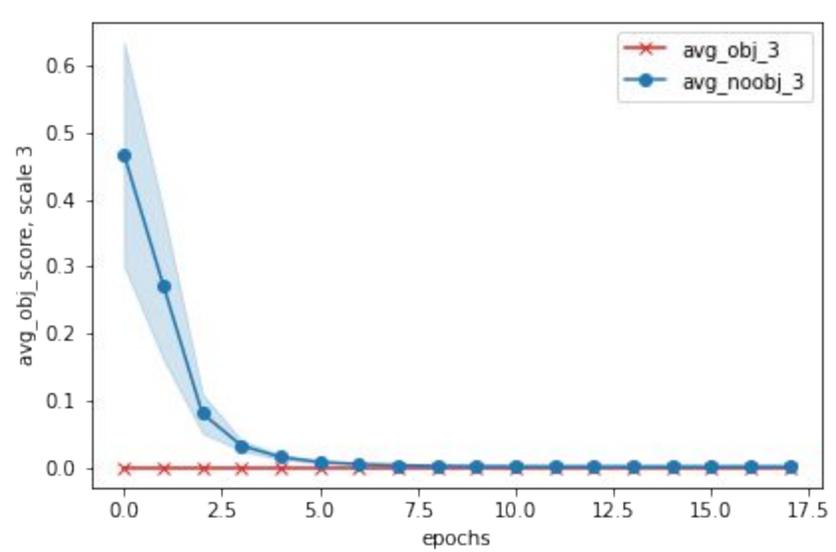
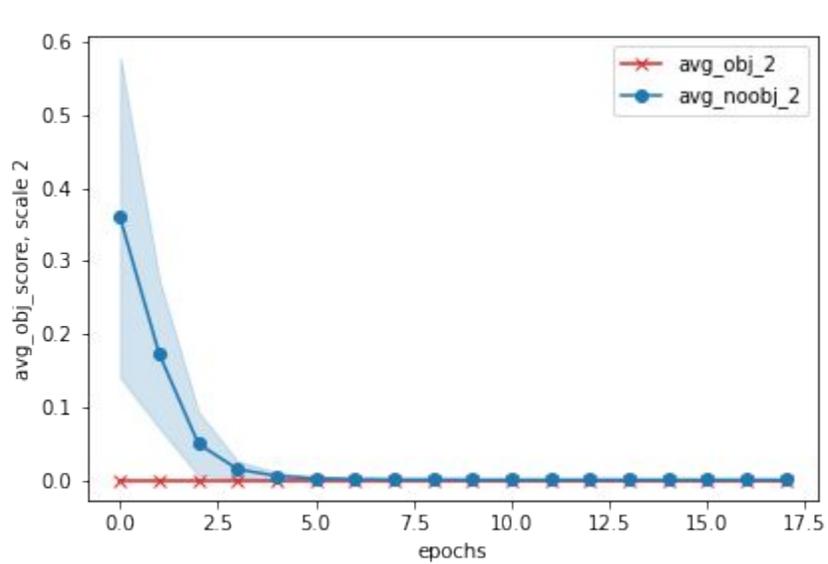
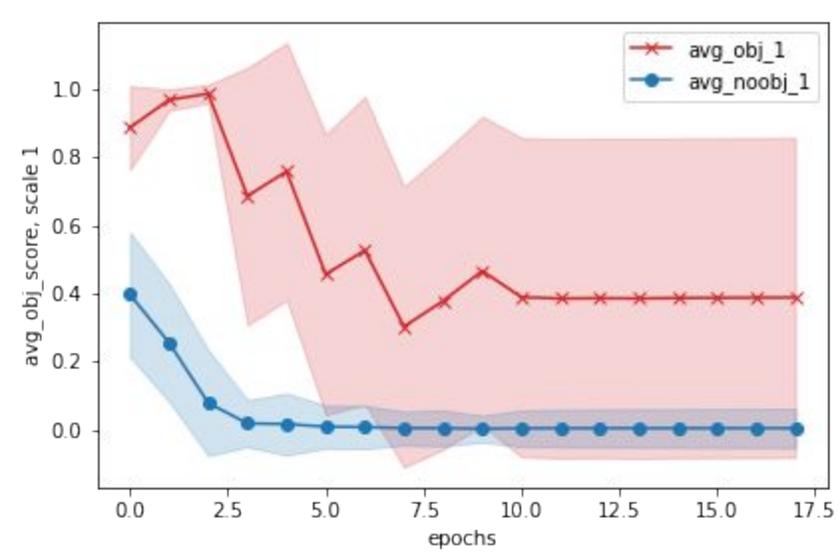
The same is expected for the IOU of predicted boxes with truth boxes.

# Network Training On Raccoons

---

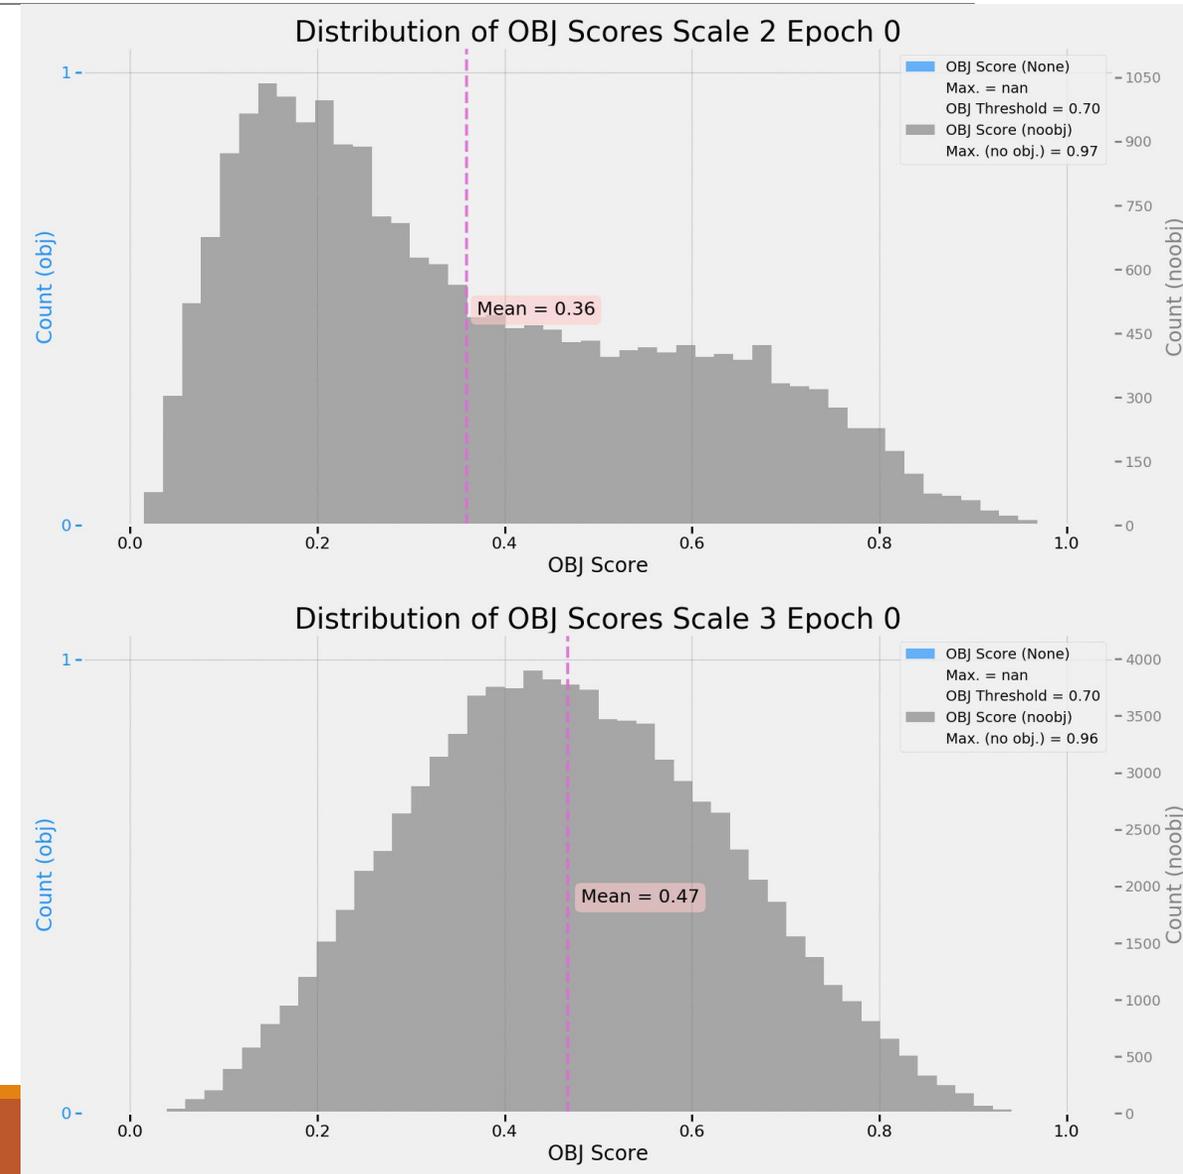
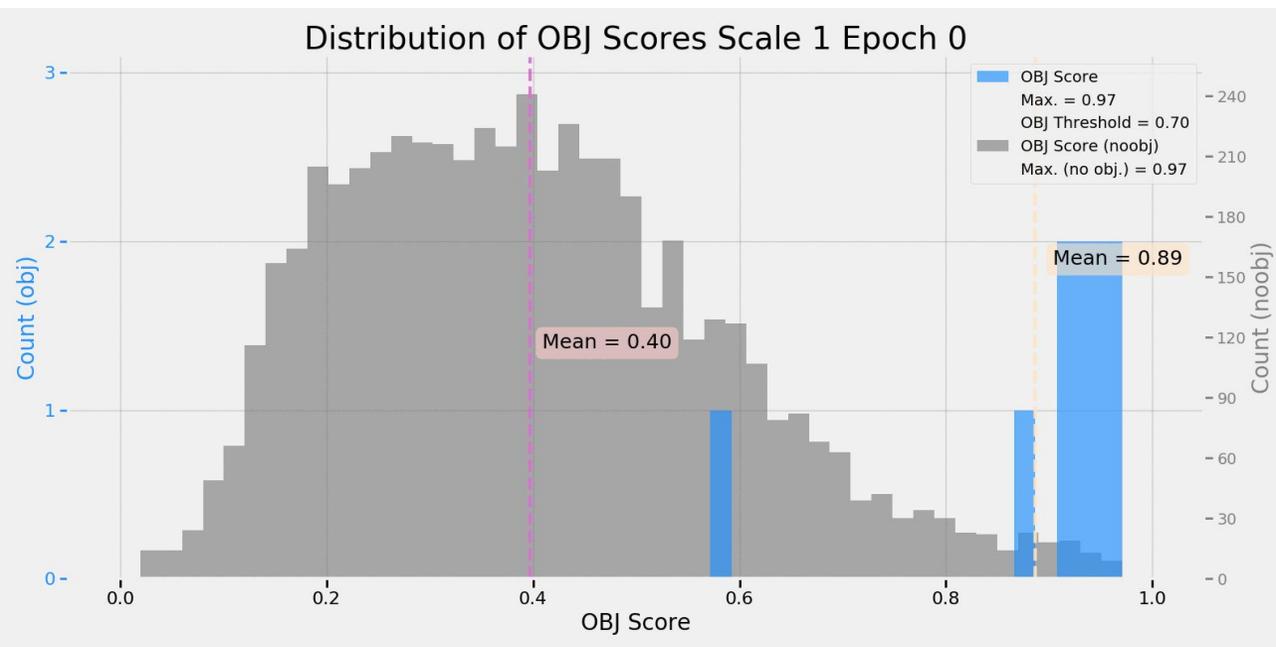
- We have a total of 8 truth boxes in this dataset (1 per image in the batch).
- All the truth boxes are closest in size to a large anchor box and therefore are associated with scale 1 (the coarsest scale - with the fewest grid cells).

On the following slide we plot the average object score and iou score for predictions from all 3 scales, along with error bars of 1 standard deviation.



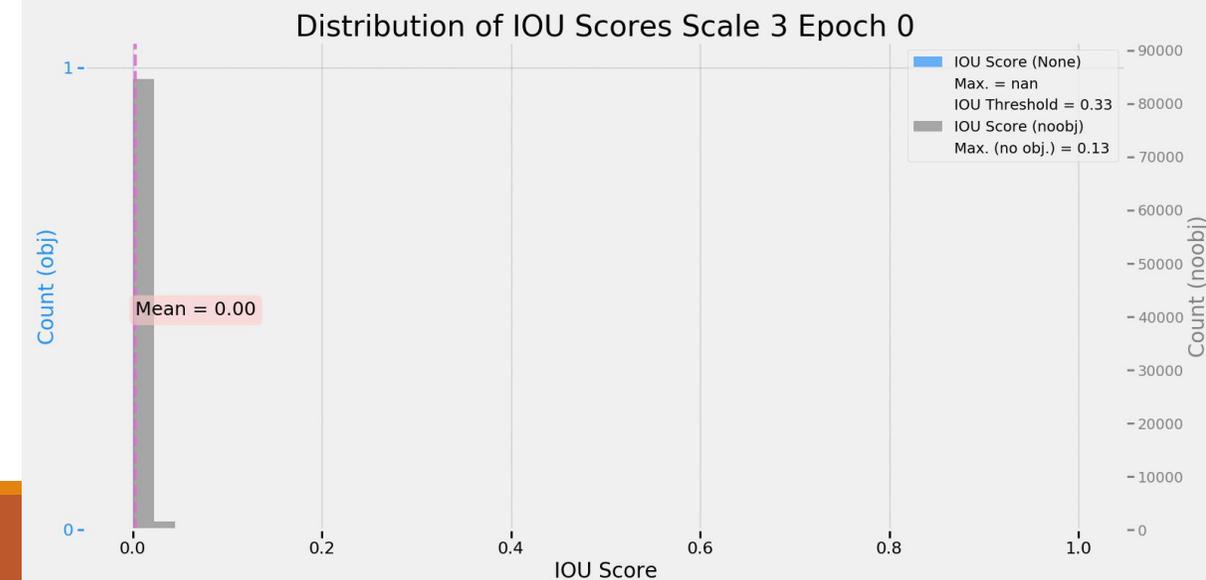
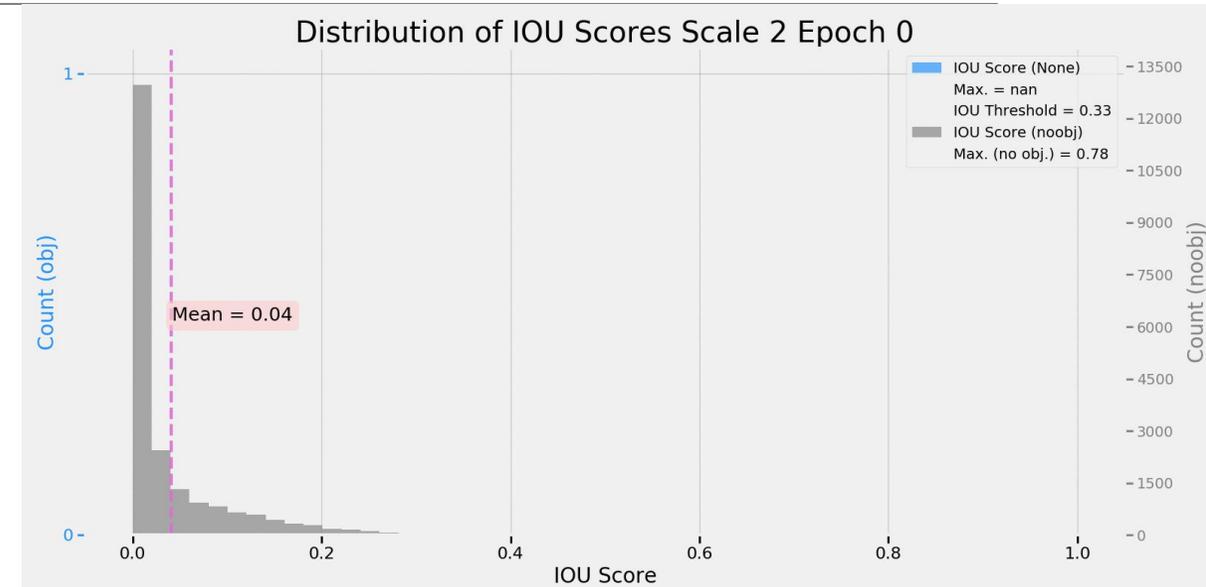
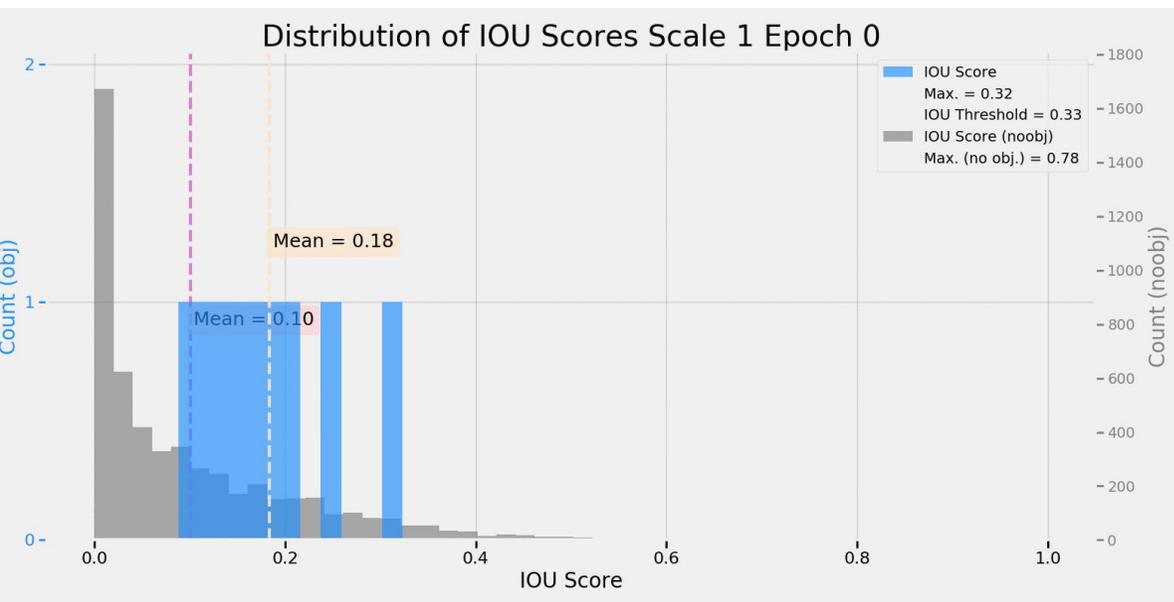
# Network Training On Raccoons

What do the distributions of object scores and iou scores look like as we train for predictions from each scale?



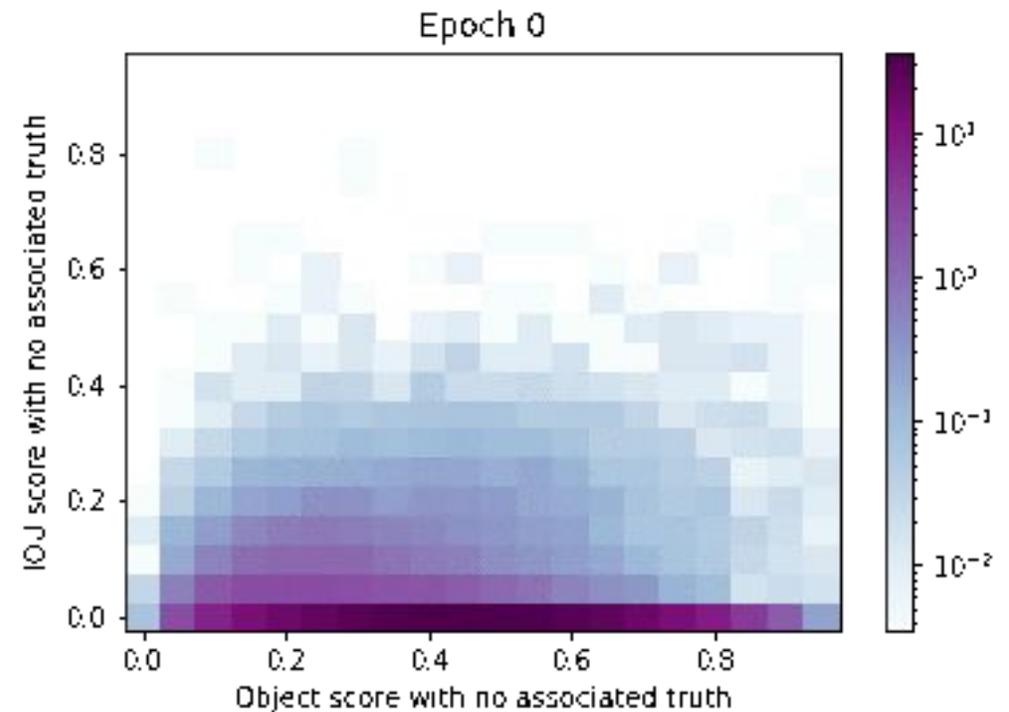
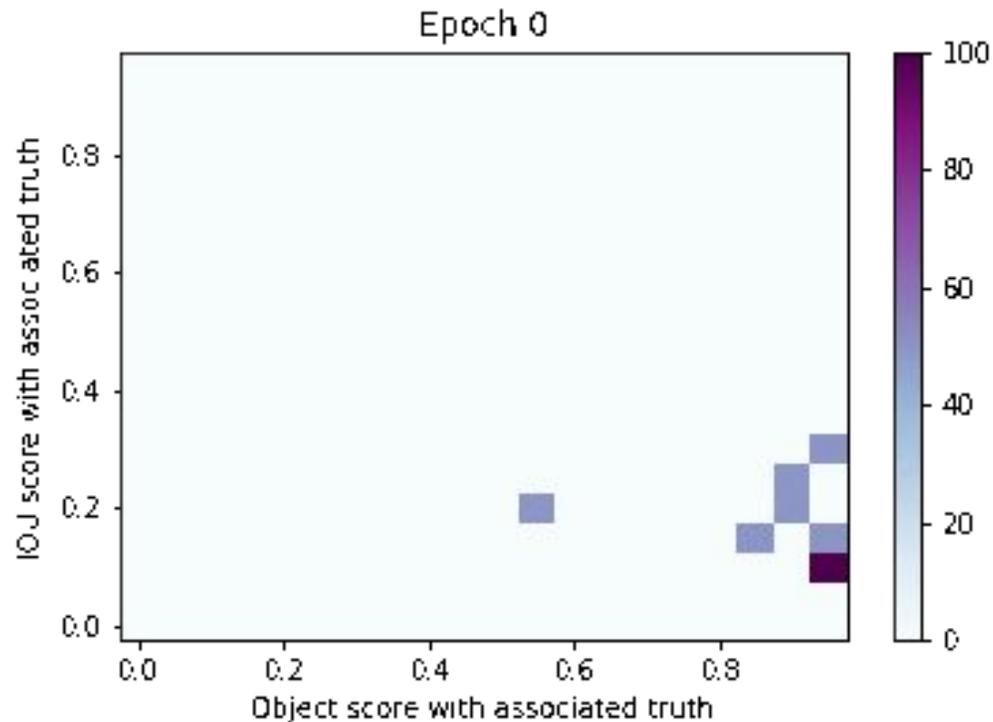
# Network Training On Raccoons

What do the distributions of object scores and iou scores look like as we train for predictions from each scale?



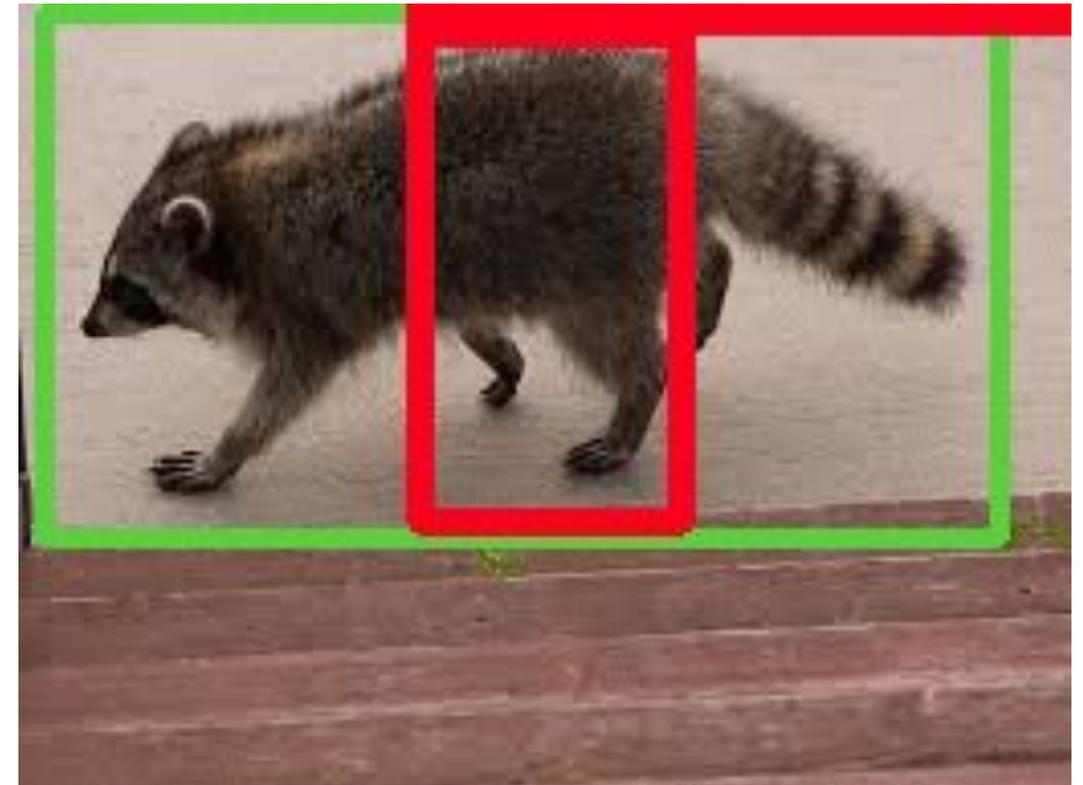
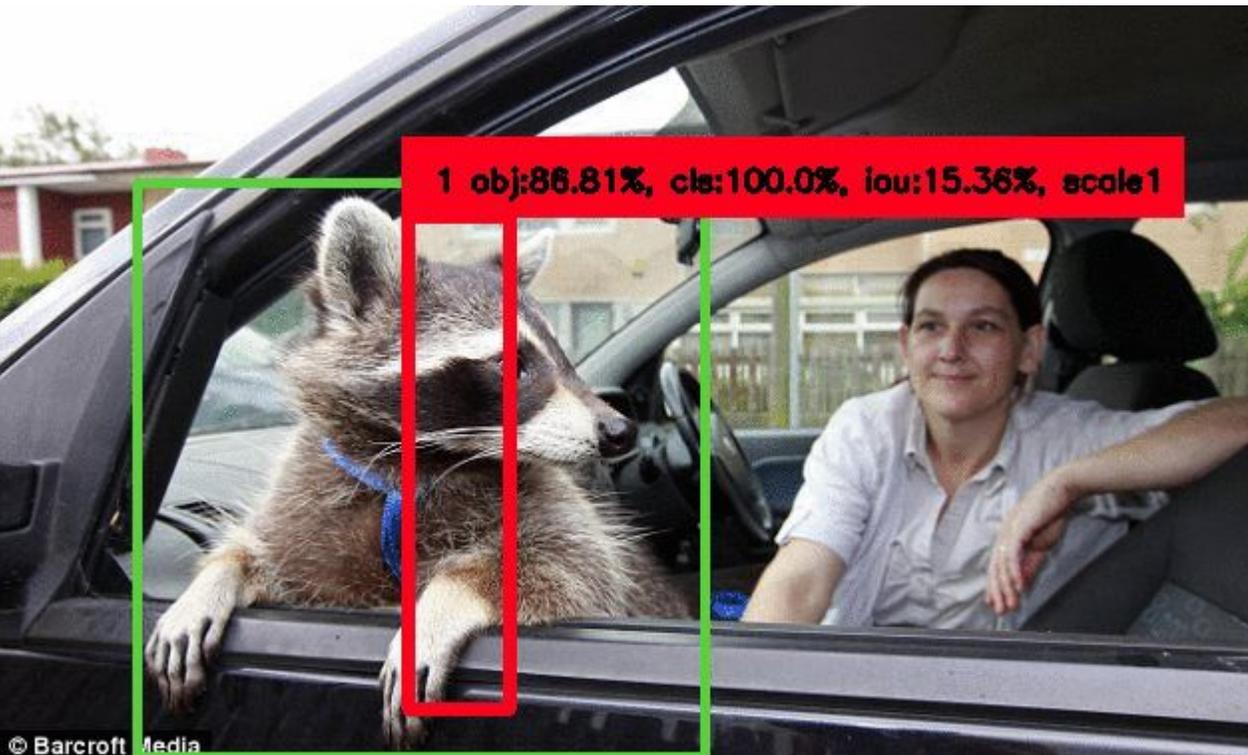
# Network Training On Raccoons

What is the spread of object scores and iou scores for all predictions from all scales?



# Network Training On Raccoons

Let's look at a couple of examples for the prediction boxes with associated truth boxes throughout training



# Network Training On Raccoons

- The average OBJ scores for boxes with and without associated truth boxes are within 1 std deviation in later epochs because the average OBJ score for predictions with associated truth decreases after the 3rd epoch. Std deviation might decrease with the use of more batches to compute these metrics
- The OBJ scores for predicted boxes with associated truth boxes appear to separate well from those without associated truth boxes when looking at the histogram for scale 1. For scales 2 and 3 the object scores for predictions quickly reduce to near 0.
- However, when we plot the OBJ score vs the IOU score, we see a number of examples where the OBJ score is large when we use a log scale. These examples tend to have IOU scores which are similar to those of the predictions which do have an associated truth box. Perhaps they are made by other anchor boxes from the same scale in the same/adjacent grid box. Some of these predictions might be removed when applying an IOU threshold for the overlap of predicted boxes
- The IOU score distribution for predictions with no associated truth box appears to have a longer tail for predictions made by scale 1 of the network. The tail is also longer for scale 2 than scale 3. This makes sense as scale 2 and 3 make predictions whose size guided by smaller bounding boxes.

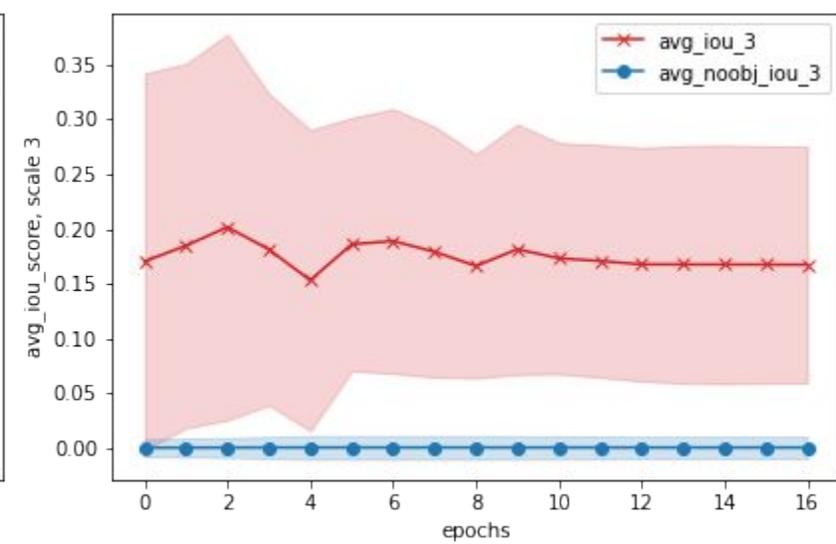
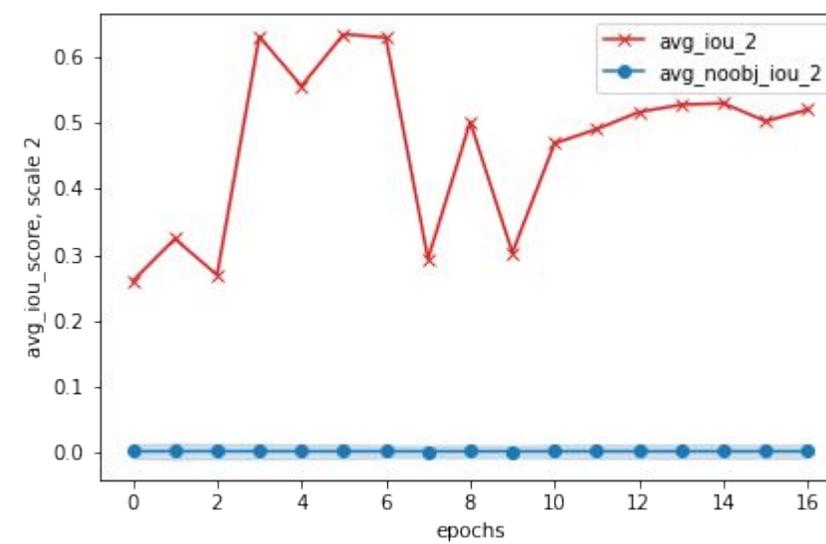
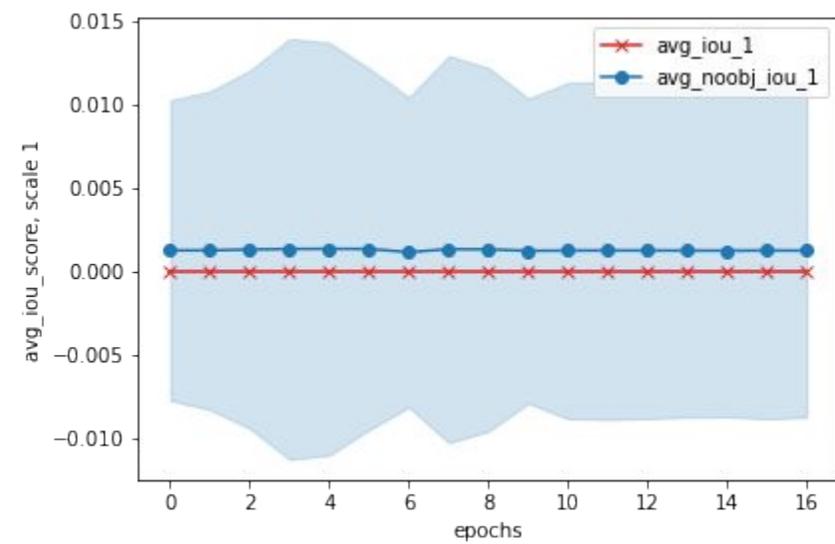
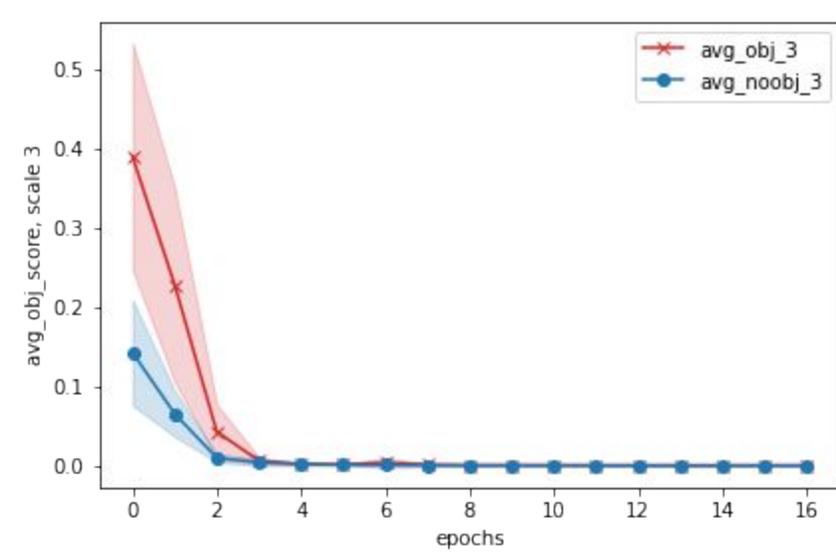
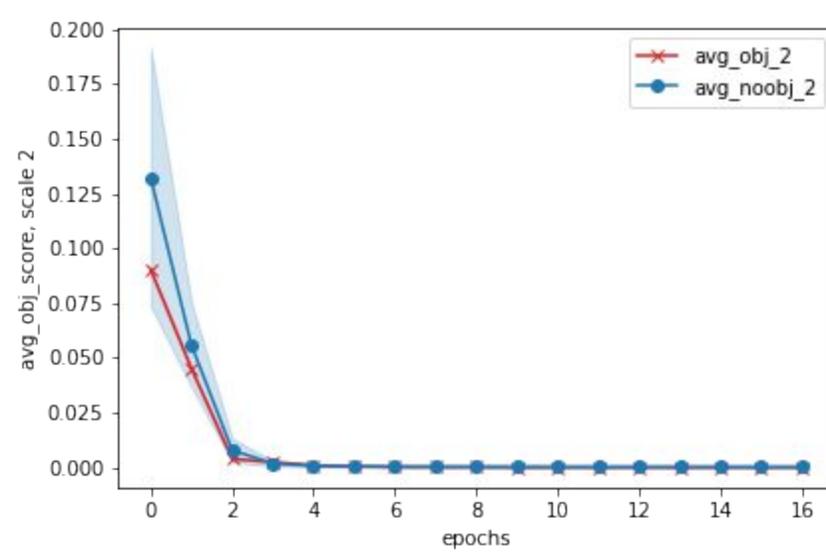
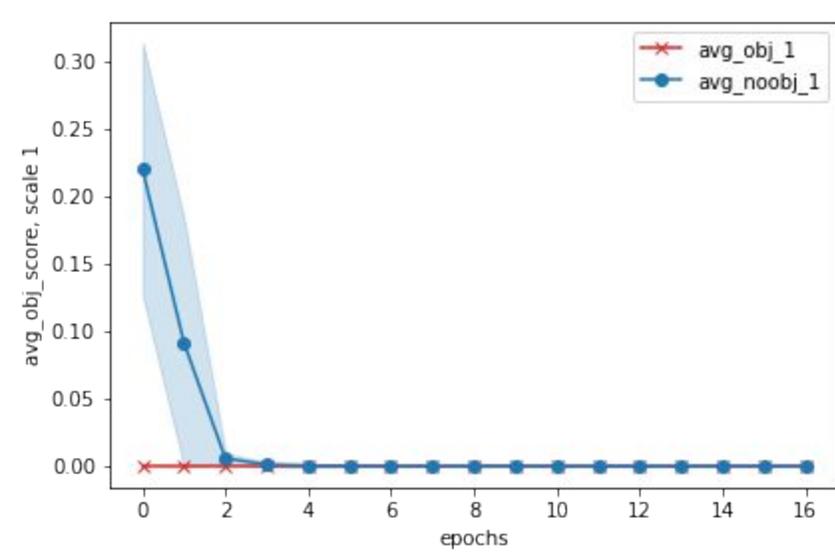
# Network Training On SNNu

---

Do we see something similar when we train on the SNNu dataset? The main differences between this and the raccoons dataset are the increased sparsity and the different downsampling method.

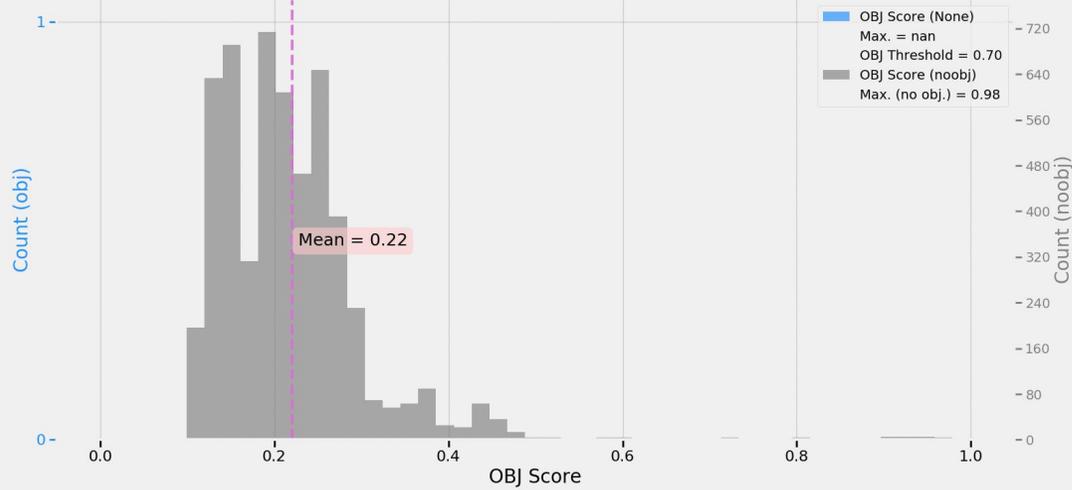
- We have a total of 12 truth boxes in this dataset
  - 1 for scale 2, the medium granularity
  - 11 for scale 3, the finest scale

Once again, on the following slide we plot the average object score and iou score for predictions from all 3 scales, along with error bars of 1 standard deviation.

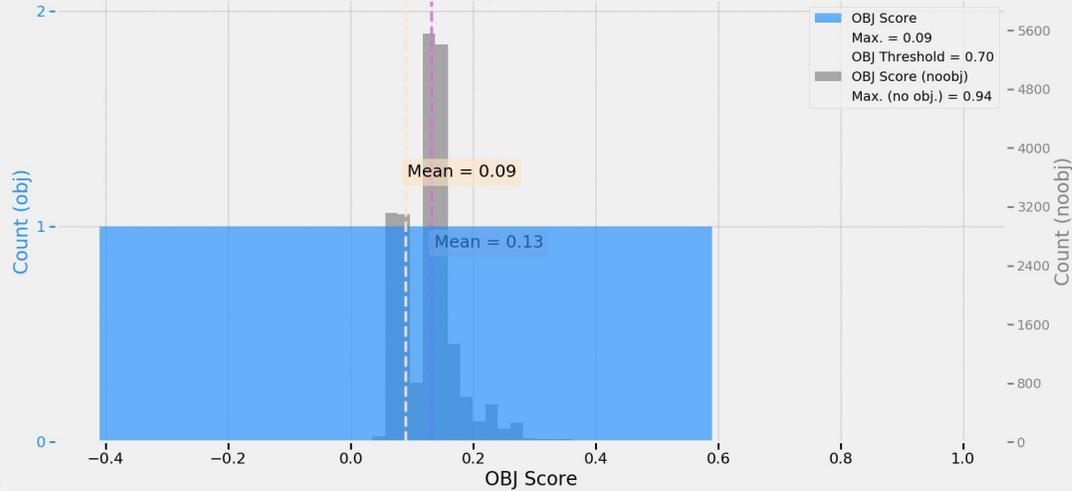


# Network Training On SNNu

Distribution of OBJ Scores Scale 1 Epoch 0

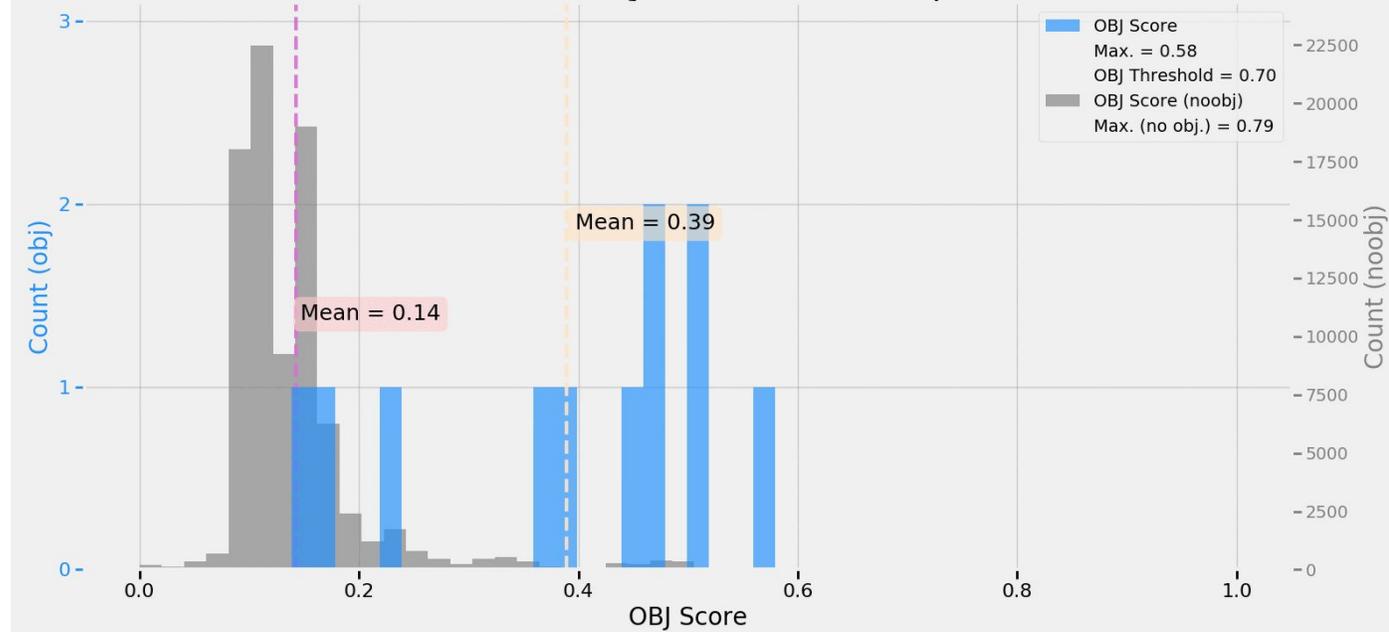


Distribution of OBJ Scores Scale 2 Epoch 0

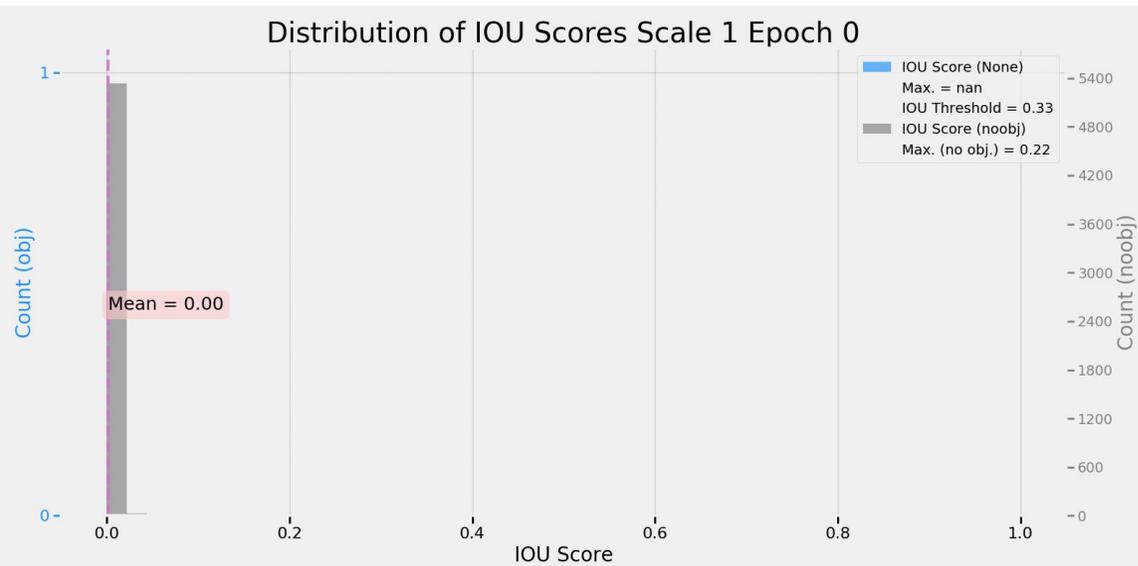


What do the distributions of object scores and iou scores look like as we train for predictions from each scale?

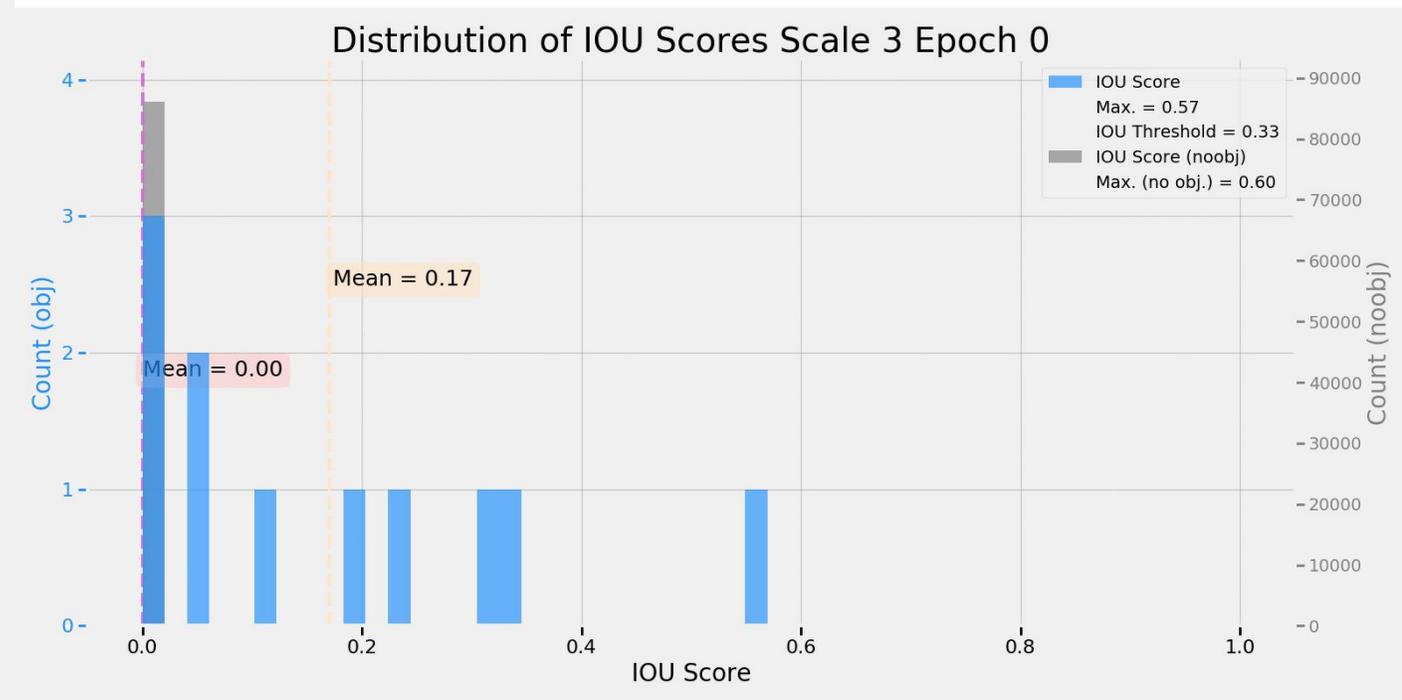
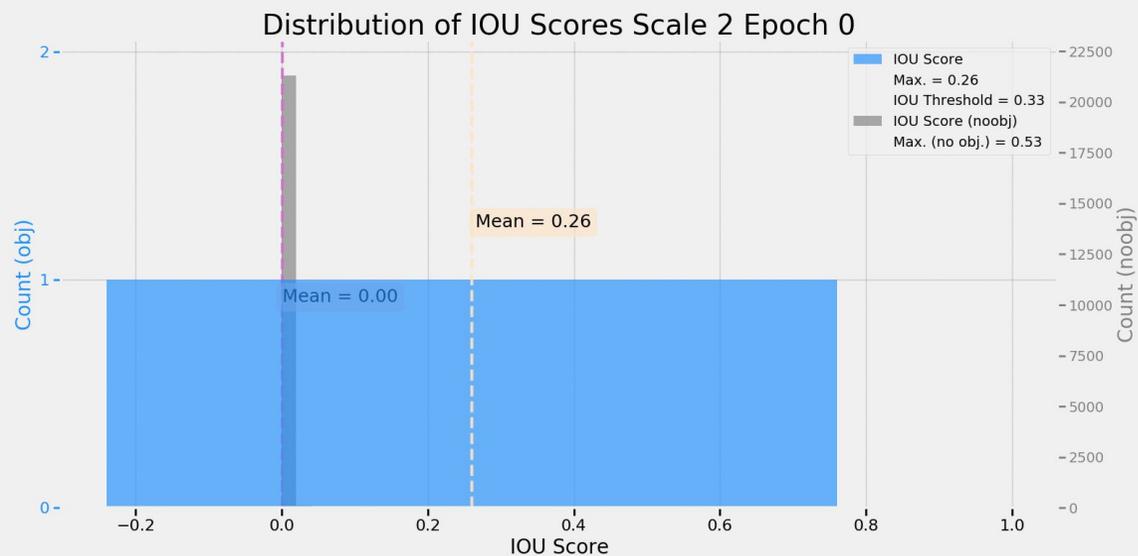
Distribution of OBJ Scores Scale 3 Epoch 0



# Network Training On SNNu

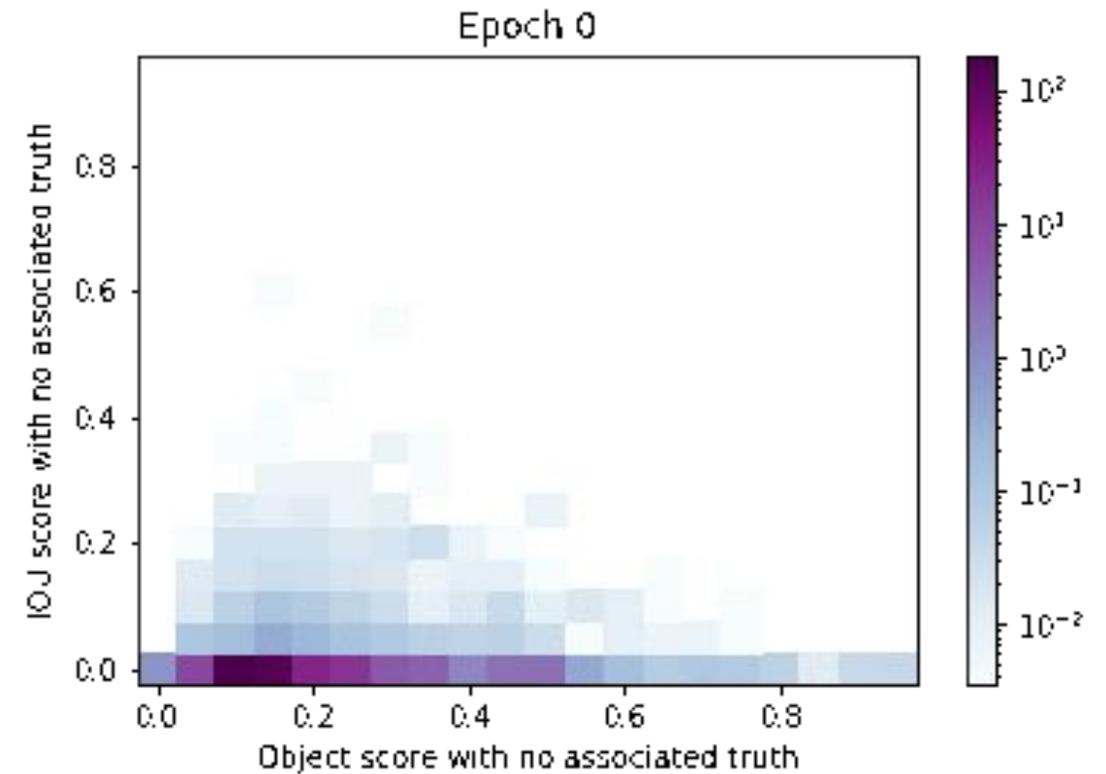
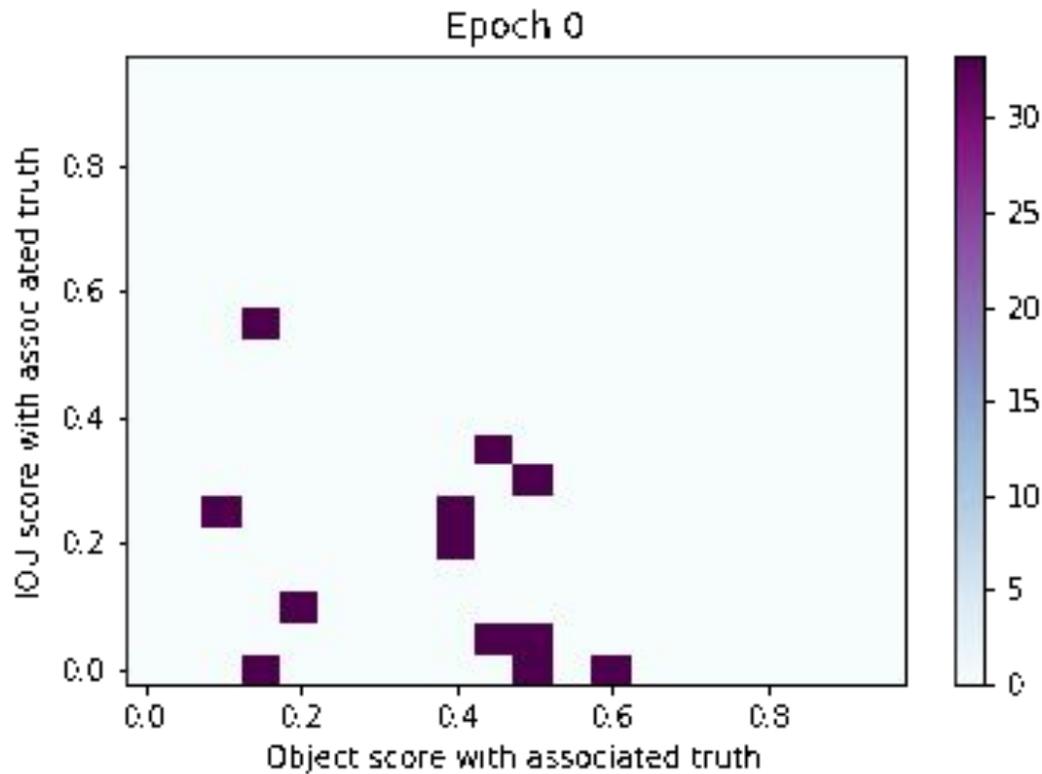


What do the distributions of object scores and iou scores look like as we train for predictions from each scale?



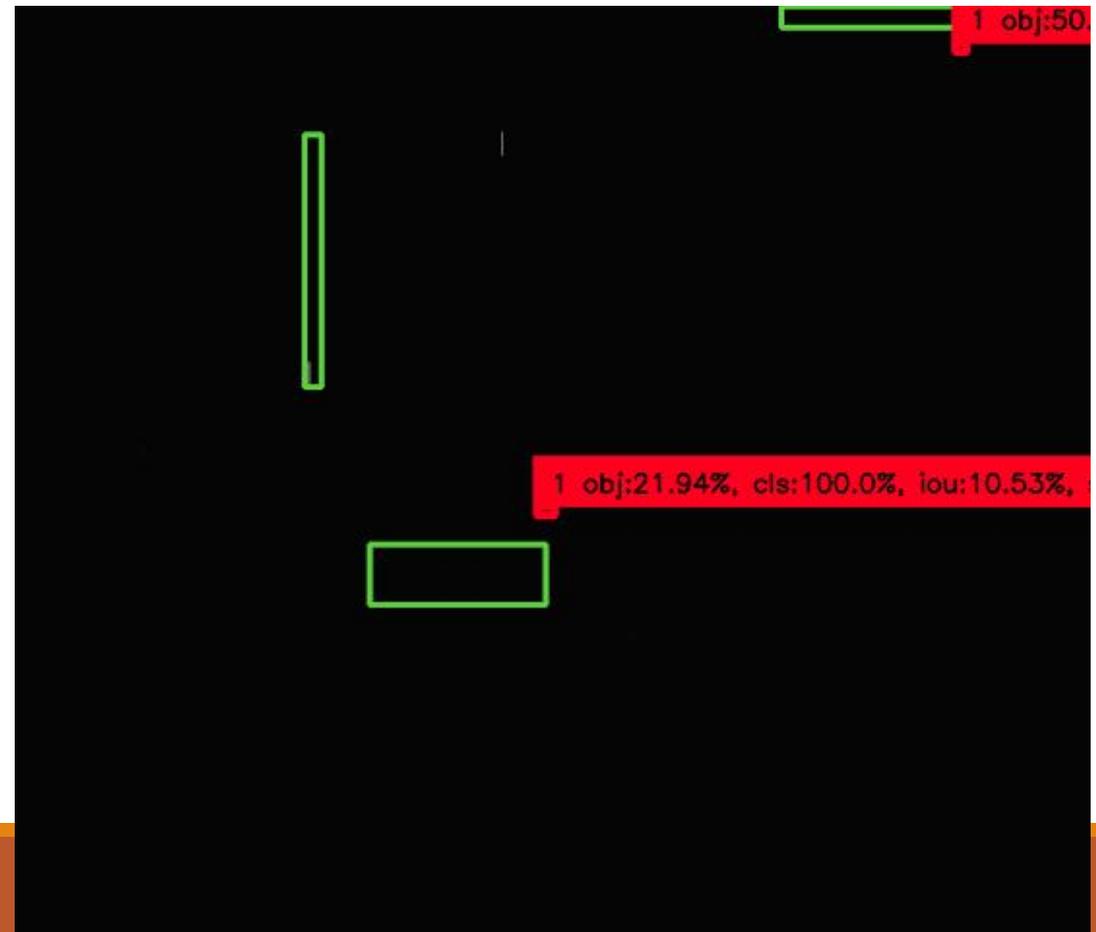
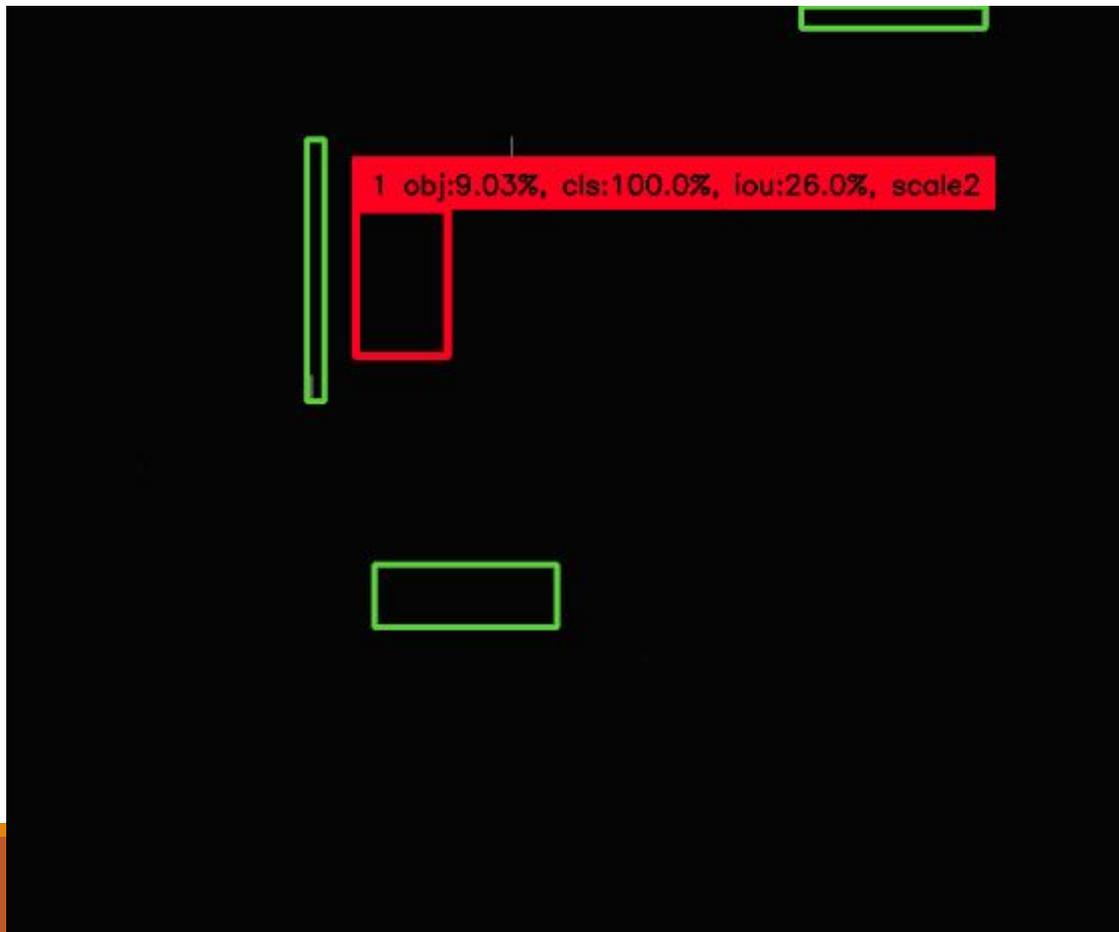
# Network Training On SNNu

What is the spread of object scores and iou scores for all predictions from all scales?



# Network Training On SNNu

Let's look at prediction boxes which have an associated truth box for an image which has 1 truth box associated with scale 2 and 2 truth boxes associated with scale 3



# Network Training On SNNu

- The average OBJ scores for boxes with and without associated truth boxes rapidly decrease to be comparable to predictions with no associated truth box
- The OBJ scores for predicted boxes with associated truth boxes rapidly become similar to those from predictions without associated truth boxes for both scales 2 and 3
- The IOU scores for predictions with associated truth boxes have a higher average than those without for scales 2 and 3
- Plotting the OBJ score vs the IOU score, doesn't appear to give use anything else (at least with bin widths of 0.05).

# Summary

---

- We've looked at how the object score and IOU scores change for predictions which do have an associated truth box and those which do not for raccoons and SNNu datasets.
- For both datasets our main problem is that the object scores for predictions which we know should produce high objectness scores (when the network is functioning/learning correctly) are very low.
- Even more confusingly, they start high. Perhaps we should try upweighting the contribution of the object score to the loss function.
- Our network's performance depends heavily on the object score. It doesn't bode well for performance if this isn't optimised during training. We could look at the FPR, precision and recall if we can identify a 'best epoch' which it would be worth doing this for.
- Will likely get better insight when we compute these metrics over many batches instead of just 1
- It is possible that IOU scores appear more reasonable than objectness scores for predictions which have an associated truth box just because they are guided by a grid cell location and a guide size (from the anchor box)

# YOLOv3 Network - progress + to do

- Added calculation of IOU score for predicted boxes with no associated truth box
- Added distributions of prediction obj and iou with truth scores split by scale and whether a truth box is associated with the prediction
- Added output of the truth boxes and predicted boxes during each epoch to allow us to determine performance metrics external to our training code

## To do:

- Look at these metrics for the output of training on bottles dataset
- Use output truth and predicted boxes to determine FPR, precision and IOU for different object score thresholds and no other cuts
- Do the same but in addition for a range of thresholds on the IOU of predicted boxes when we discard overlapping predictions
- Add log scale (and single y axis with normalization) to obj and IOU distribution histograms
- Bug fix to allow saving of images of predictions made which have no associated truth box during each epoch of training
- Extend metrics callback over multiple batches
- Reproduce the same plots and distributions with higher statistics from multiple batches
- Investigate the magnitude of different contributions to loss during training

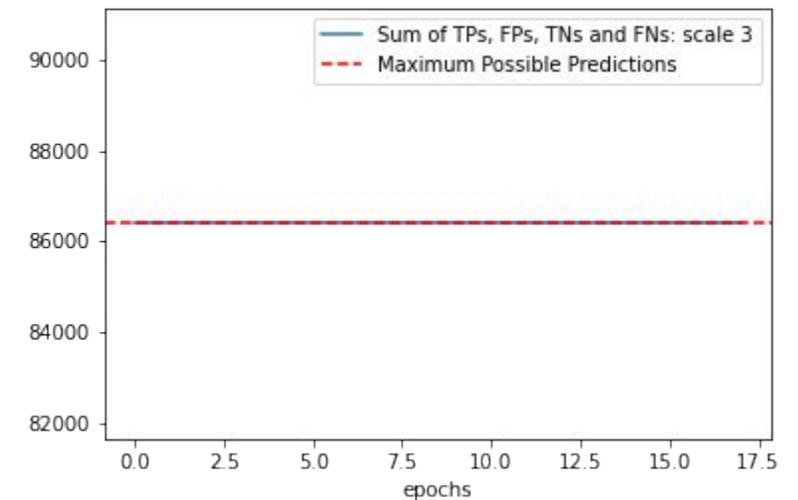
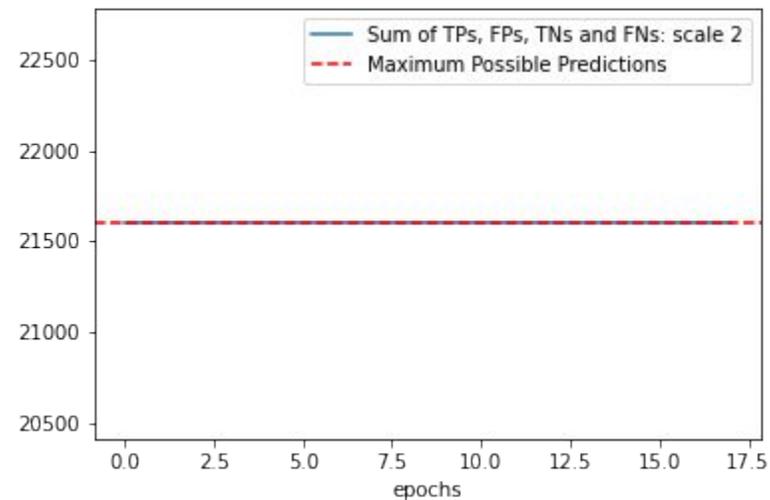
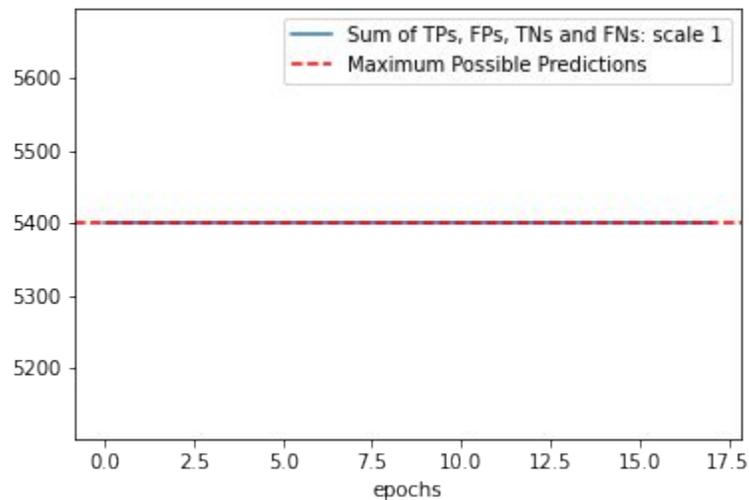
# Backup Slides

---

# Network Training On Raccoons

Counting the total number of predictions made for each scale. We should have:

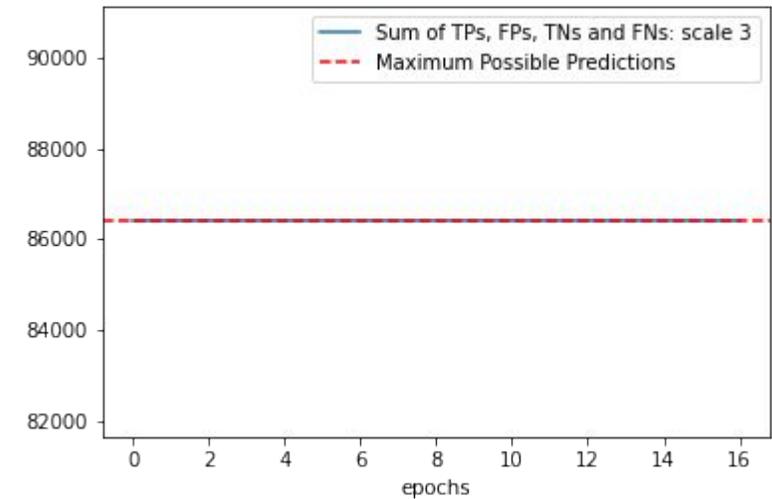
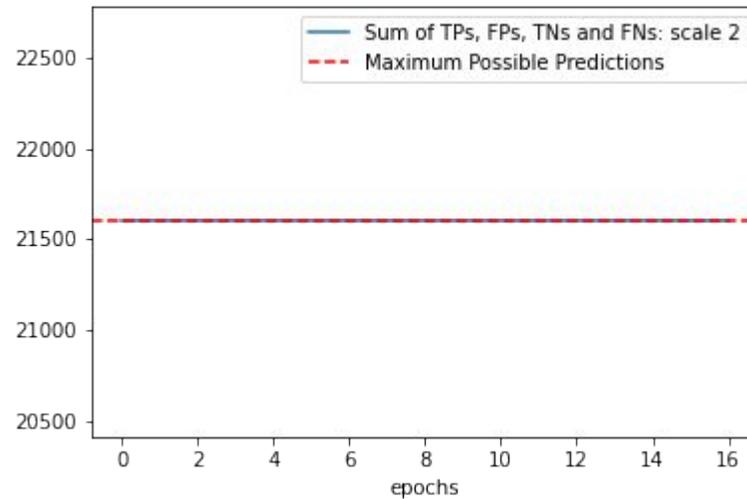
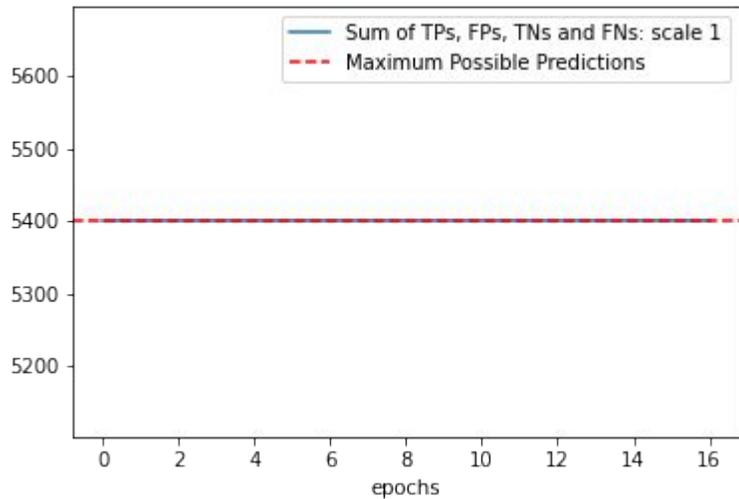
- 5400 for scale 1
- 21600 for scale 2
- 86400 for scale 3



# Network Training On SNNu

Counting the total number of predictions made for each scale. We should have:

- 5400 for scale 1
- 21600 for scale 2
- 86400 for scale 3



# Network Training On Bottles

---

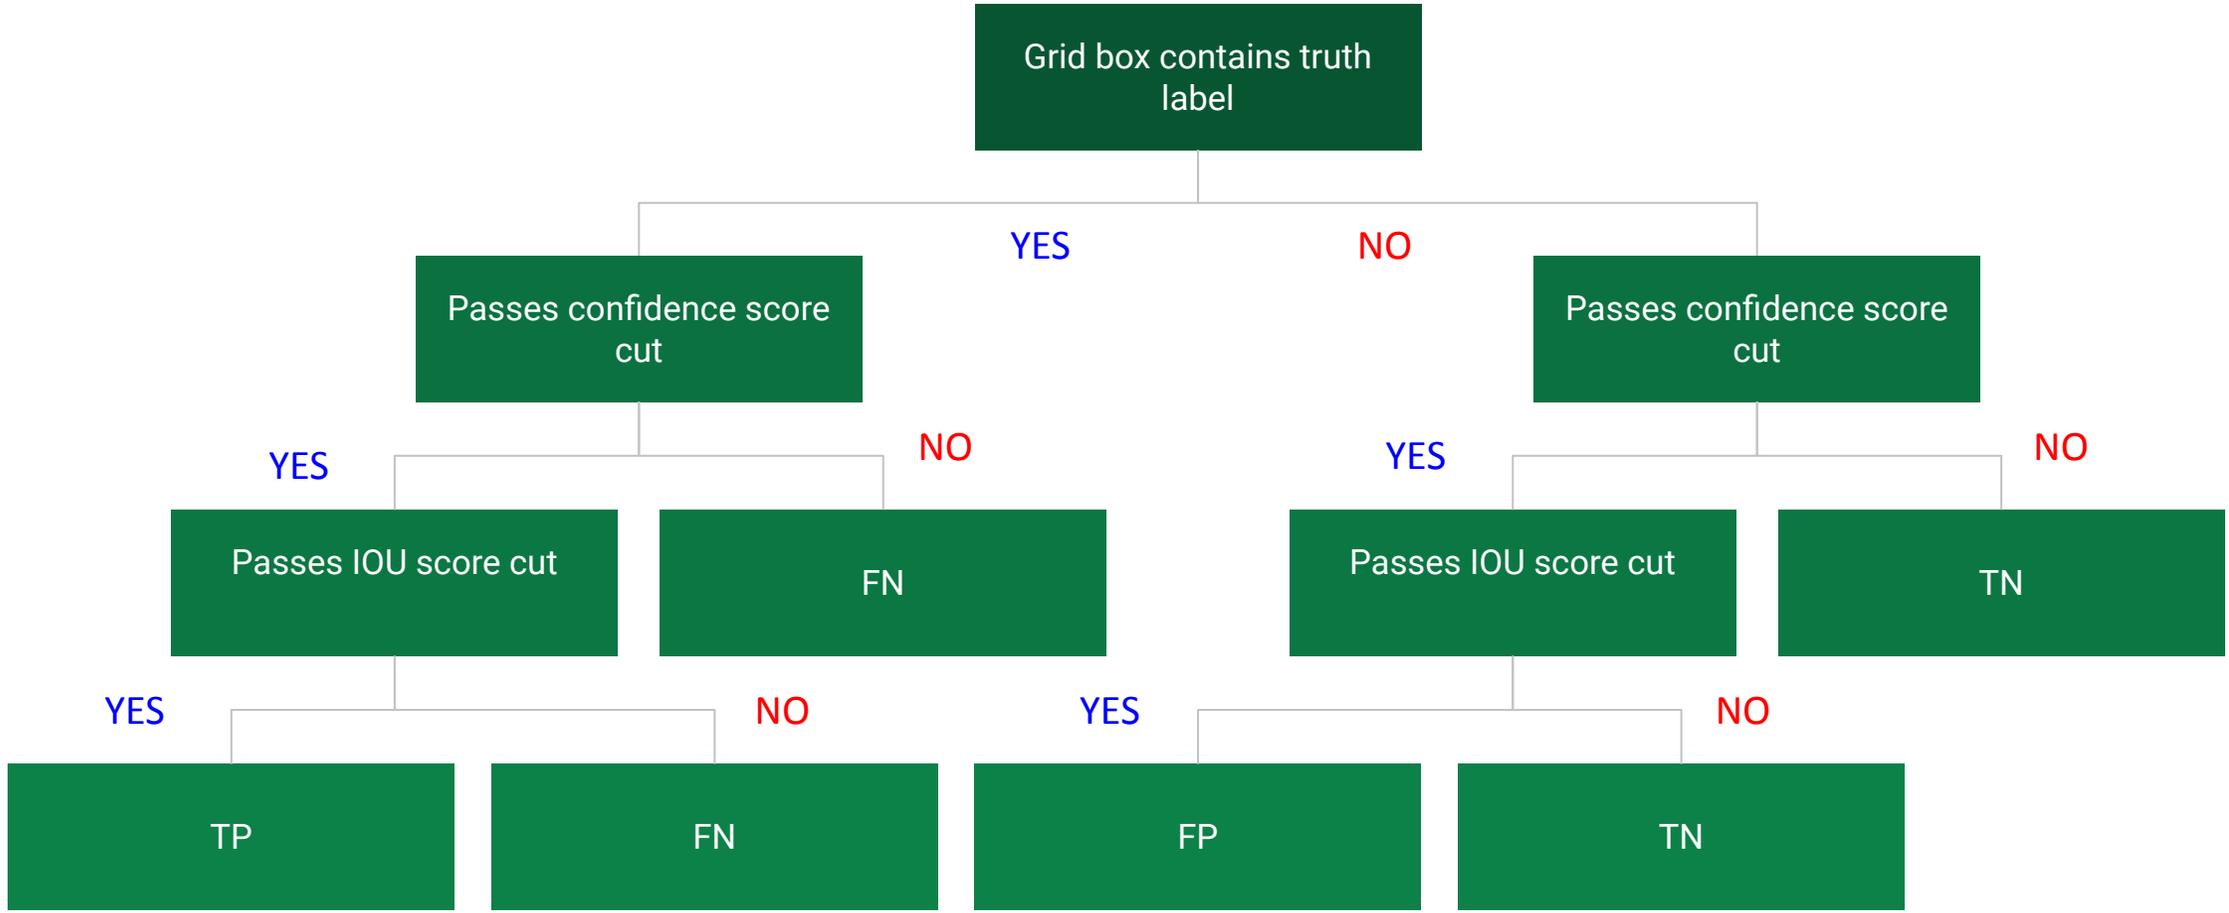
# Network Training On Bottles

---

Counting the total number of predictions made for each scale. We should have:

- 5400 for scale 1
- 21600 for scale 2
- 86400 for scale 3

# Metrics Flowchart



# Predictions out of training

All boxes that passed the obj. conf. cut need to be ordered in descending order of obj. score

For each predicted box:

Passes object confidence score cut?

YES

NO

Is the box with highest obj. score of all unassigned boxes?

YES

NO

Overlap with unassigned truth box > 0?

YES

NO

Assign truth box to this predicted box

Passes IOU (with truth box) threshold?

YES

TP

NO

FP & go to start

IOU with better boxes less than NMS threshold?

YES

NO

Overlap with unassigned truth box > 0?

YES

NO

More boxes exist to assign to truth?

YES

NO

FP & FN for unassigned truth box

Discard box and look at next box

Box(es) exist

Go to start

NO boxes exist

Truth box(es) exist?

YES

NO

FN for each truth box

TN

# Summary + Next Steps

---

We think we understand how the network is reading the truth information on a per-batch level. We have run our new callback with a small dataset of 8 images only to test it and have looked at the output metrics.

We are going to ensure we understand what is going on for individual images by plotting the predicted box and the truth box associated with it on the image for each epoch. We can then see how the network learns by visualising it

The weights learned by the network appear to function poorly for predicting the objectness score. Due to the callback using only 8 images, it's not really possible to tell if this is due to the one factor or another yet, but it doesn't correlate with the reported metrics from the training and validation sets which we can still access from the logging layer method.

We are adding a callback which evaluates the same metrics but for the entire training/validation dataset. We are also running the metrics training on our bottles and SNNu dataset

In future, we will also be able to calculate more complex metrics combining the outputs from all 3 scales from within a callback in a similar manner.