

Computer vision and machine learning for the material scientist

Lecture 9. *Object Detection*

Romain Vo



*slides adapted from [CS231n](#)

Computer Vision Tasks

Classification



DOG



Classify the image

Computer Vision Tasks

Classification




DOG



Classify the image

Semantic segmentation



DOG, CAT, BG



Classify each pixel

Computer Vision Tasks

Classification




DOG



Classify the image

Semantic segmentation



DOG, CAT, BG



Classify each pixel

Instance Segmentation



SMTH, SMTH, SMTH



Segment independent instances

Panoptic segmentation



DOG, DOG, CAT



Segment & Classify independent instances

Computer Vision Tasks

Classification




DOG



Classify the image

Semantic segmentation




DOG, CAT, BG



Classify each pixel

Instance Detection



SMTH, SMTH, SMTH



Detect independent instances

Panoptic segmentation



DOG, DOG, CAT



Segment & Classify independent instances

Computer Vision Tasks

Classification




DOG



Classify the image

Semantic segmentation




DOG, CAT, BG



Classify each pixel

Instance Detection

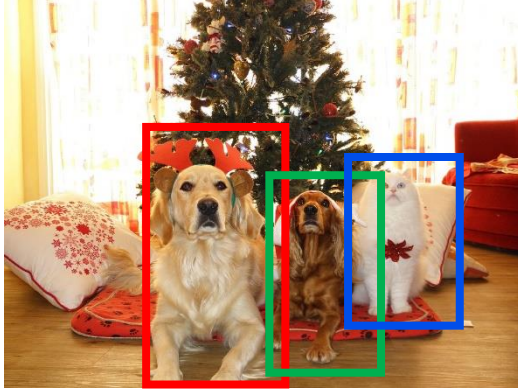


SMTH, SMTH, SMTH



Detect independent instances

Panoptic Detection



DOG, DOG, CAT



Detect & Classify independent instances

Computer Vision Tasks

Classification




DOG



Classify the image

Semantic segmentation




DOG, CAT, BG



Classify each pixel

Instance Detection




SMTH, SMTH, SMTH



Detect independent instances

Object Detection



DOG, DOG, CAT



Detect & Classify independent instances

Object Detection


Classification



DOG

Classify the image


Semantic segmentation



DOG, CAT, BG

Classify each pixel


Instance Detection



SMTH, SMTH, SMTH

Detect independent instances


Object Detection



DOG, DOG, CAT

Detect & Classify independent instances



Object Detection : metrics




Ground-truth bounding box

Predicted box

$$IoU = \frac{Intersection}{Union}$$



Object Detection : metrics



Ground-truth
bounding box



Predicted
box

$$IoU = \frac{\text{Intersection}}{\text{Union}}$$

A box is correctly detected when $IoU > t$

$t = \text{user-defined threshold}$

(and when the label associated is correct)




Object Detection : Precision / Recall



Ground-truth

Prediction

True positive



Ground-truth

Prediction

False positive



Ground-truth

Prediction

False negative

Precision: $P = \frac{TP}{TP+FP}$


Object Detection : Precision / Recall



Ground-truth

Prediction

True positive



Ground-truth

Prediction

False positive



Ground-truth

Prediction

False negative

Precision: $P = \frac{TP}{TP+FP}$

i.e rate of well-detected objects among **detected objects**


Object Detection : Precision / Recall



Ground-truth

Prediction

True positive



Ground-truth

Prediction

False positive



Ground-truth

Prediction

False negative

$$\text{Precision: } P = \frac{TP}{TP+FP}$$

i.e rate of well-detected objects among **detected objects**

$$\text{Recall: } R = \frac{TP}{TP+FN}$$


Object Detection : Precision / Recall



Ground-truth

Prediction

True positive



Ground-truth

Prediction

False positive



Ground-truth

Prediction

False negative

Precision: $P = \frac{TP}{TP+FP}$

i.e rate of well-detected objects among **detected objects**

Recall: $R = \frac{TP}{TP+FN}$

i.e rate of well-detected objects among **ground-truth objects**


Object Detection : Precision / Recall



Ground-truth

Prediction

True positive



Ground-truth

Prediction

False positive



Ground-truth

Prediction

False negative

Precision: $P = \frac{TP}{TP+FP}$

i.e rate of well-detected objects among **detected objects**

Recall: $R = \frac{TP}{TP+FN}$

i.e rate of well-detected objects among **ground-truth objects**

Intuition : if t is very low or close to 0, all predictions are considered True Positive, but we increase the number of False Positive (and inversely).

Precision/Recall trade-off


Object Detection : Precision / Recall curve and mAP

Precision: $P = \frac{TP}{TP+FP}$

Recall: $R = \frac{TP}{TP+FN}$

True Positive if $IoU > t$

- By varying the confidence threshold t for the detection, one can obtain so-called precision-recall curve




Object Detection : Precision / Recall curve and mAP

Precision: $P = \frac{TP}{TP+FP}$

Recall: $R = \frac{TP}{TP+FN}$

True Positive if $IoU > t$

- By varying the confidence threshold t for the detection, one can obtain so-called precision-recall curve
- The **average precision (AP)** is defined as the area under the curve




Object Detection : Precision / Recall curve and mAP

Precision: $P = \frac{TP}{TP+FP}$

Recall: $R = \frac{TP}{TP+FN}$

True Positive if $IoU > t$

- By varying the confidence threshold t for the detection, one can obtain so-called precision-recall curve
- The **average precision (AP)** is defined as the area under the curve



Object Detection : Precision / Recall curve and mAP


Precision: $P = \frac{TP}{TP+FP}$

Recall: $R = \frac{TP}{TP+FN}$

True Positive if $IoU > t$

- By varying the confidence threshold t for the detection, one can obtain so-called precision-recall curve
- The **average precision (AP)** is defined as the area under the curve
- We say **mean average precision (mAP)** because the **AP** is computed for all object classes

mAP = $mean(AP_{dog}, AP_{cat}, AP_{human})$



Object Detection : Box Regression

a box is defined by 4 parameters:


- center (x, y)
- height h
- width w



w

h

Object Detection : Single object




a box is defined by 4 parameters:

- center (x, y)
- height h
- width w

and, we assign a class label \underline{c} to this box


Object Detection : Single object



a box is defined by 4 parameters:

- center (x, y)
- height h
- width w

and, we assign a class label \underline{c} to this box




class vector $z \in \mathbb{R}^C$ \longrightarrow cross-entropy loss

box prediction (x', y', h', w') \longrightarrow ℓ_2 regression loss

+

Total loss


Object Detection : Single object



a box is defined by 4 parameters:

- center (x, y)
- height h
- width w

and, we assign a class label \underline{c} to this box




class vector $z \in \mathbb{R}^C$ \rightarrow cross-entropy loss

box prediction (x', y', h', w') \rightarrow ℓ_2 regression loss

$+$ \rightarrow Total loss

Object Detection : Multi objects

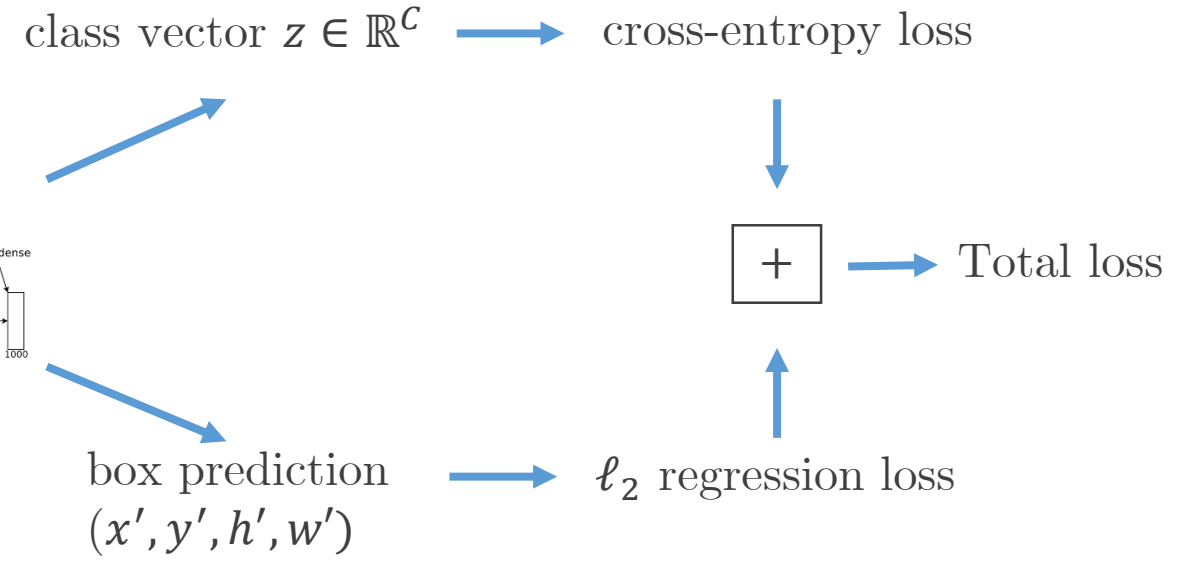



a box is defined by 4 parameters:

- center (x, y)
- height h
- width w

and, we assign a class label \underline{c} to this box

Problem: how to treat a varying number of Objects ??



Object Detection : Multi objects



a box is defined by 4 parameters:

- center (x, y)
- height h
- width w

and, we assign a class label \underline{c} to this box

Naive approach : sample all possible crops on the image and evaluate the classification/regression network

Object Detection : Multi objects




a box is defined by 4 parameters:

- center (x, y)
- height h
- width w

and, we assign a class label \underline{c} to this box

Naive approach : sample all possible crops on the image and evaluate the classification/regression network

Object Detection : Multi objects



a box is defined by 4 parameters:

- center (x, y)
- height h
- width w

and, we assign a class label \underline{c} to this box

Naive approach : sample all possible crops on the image and evaluate the classification/regression network

Object Detection : Multi objects



a box is defined by 4 parameters:

- center (x, y)
- height h
- width w

and, we assign a class label \underline{c} to this box

Naive approach : sample all possible crops on the image and evaluate the classification/regression network

Object Detection : Multi objects



a box is defined by 4 parameters:

- center (x, y)
- height h
- width w

and, we assign a class label \underline{c} to this box

Naive approach : sample all possible crops on the image and evaluate the classification/regression network

Object Detection : Multi objects



a box is defined by 4 parameters:

- center (x, y)
- height h
- width w

and, we assign a class label \underline{c} to this box

Naive approach : sample all possible crops on the image and evaluate the classification/regression network

Object Detection : Multi objects



a box is defined by 4 parameters:

- center (x, y)
- height h
- width w

and, we assign a class label \underline{c} to this box

Naive approach : sample all possible crops on the image and evaluate the classification/regression network

Object Detection : Multi objects



a box is defined by 4 parameters:

- center (x, y)
- height h
- width w

and, we assign a class label \underline{c} to this box

Naive approach : sample all possible crops on the image and evaluate the classification/regression network

Object Detection : Multi objects



a box is defined by 4 parameters:

- center (x, y)
- height h
- width w

and, we assign a class label \underline{c} to this box

Naive approach : sample all possible crops on the image and evaluate the classification/regression network

Object Detection : Multi objects



a box is defined by 4 parameters:

- center (x, y)
- height h
- width w

and, we assign a class label \underline{c} to this box

Naive approach : sample all possible crops on the image and evaluate the classification/regression network

Object Detection : Region proposals



a box is defined by 4 parameters:

- center (x, y)
- height h
- width w

and, we assign a class label c to this box


Region of Interest (RoI) proposal approach : use SOTA algorithms (in 2015) to evaluate objectness of the image and propose a reduced set of patches

2012, Alexe et al, "Measuring the objectness of image windows"

2013, Uijlings et al, "Selective Search for Object Recognition"

2014, Zitnick and Dollar, "Edge boxes: Locating object proposals from edges"

Object Detection : Region of Interest proposals



a box is defined by 4 parameters:

- center (x, y)
- height h
- width w

and, we assign a class label \underline{c} to this box


Region of Interest (RoI) proposal approach : use SOTA algorithms (in 2015) to evaluate objectness of the image and propose a reduced set of patches

2012, Alexe et al, "Measuring the objectness of image windows"

2013, Uijlings et al, "Selective Search for Object Recognition"

2014, Zitnick and Dollar, "Edge boxes: Locating object proposals from edges"

Object Detection : Region of Interest proposals



a box is defined by 4 parameters:

- center (x, y)
- height h
- width w

and, we assign a class label \underline{c} to this box

Region of Interest (RoI) proposal approach : use SOTA algorithms (in 2015) to evaluate objectness of the image and propose a reduced set of patches



CNN

class vector $z \in \mathbb{R}^C$


box correction (x', y', h', w')

2012, Alexe et al, "Measuring the objectness of image windows"

2013, Uijlings et al, "Selective Search for Object Recognition"

2014, Zitnick and Dollar, "Edge boxes: Locating object proposals from edges"

Object Detection : Region of Interest proposals




a box is defined by 4 parameters:

- center (x, y)
- height h
- width w

and, we assign a class label \underline{c} to this box

Region of Interest (RoI) proposal approach : use SOTA algorithms (in 2015) to evaluate objectness of the image and propose a reduced set of patches



CNN

class vector $z \in \mathbb{R}^C$

box correction (x', y', h', w')

Problem: still super slow !

Need to run a CNN a few 1000 times per image

2012, Alexe et al, "Measuring the objectness of image windows"

2013, Uijlings et al, "Selective Search for Object Recognition"

2014, Zitnick and Dollar, "Edge boxes: Locating object proposals from edges"

Object Detection : From R-CNN to RetinaNet


1. R-CNN to Fast R-CNN
2. Fast R-CNN to Faster R-CNN
3. Faster R-CNN to RetinaNet

Object Detection : RoI pooling

1. R-CNN to Fast R-CNN : run a CNN one time on the image, and extract the patches from the feature map not the image. Then small CNN for classification and box regression

Object Detection : From R-CNN to RetinaNet

1. R-CNN to Fast R-CNN : run a CNN one time on the image, and extract the patches from the feature map not the image. Then small CNN for classification and box regression



Project RoI proposal onto feature map : RoI pooling

Object Detection : Region Proposal Network



1. R-CNN to Fast R-CNN : run a CNN one time on the image, and extract the patches from the feature map not the image. Then small CNN for classification and box regression
2. From Fast R-CNN to Faster R-CNN : Use a Region Proposal Network (instead of handcrafted algorithm) to propose RoI

Object Detection : Anchor-based

1. R-CNN to Fast R-CNN : run a CNN one time on the image, and extract the patches from the feature map not the image. Then small CNN for classification and box regression
2. From Fast R-CNN to Faster R-CNN : Use a Region Proposal Network (instead of handcrafted algorithm) to propose RoI



$H \times W$



$$D \times \frac{H}{16} \times \frac{W}{16}$$

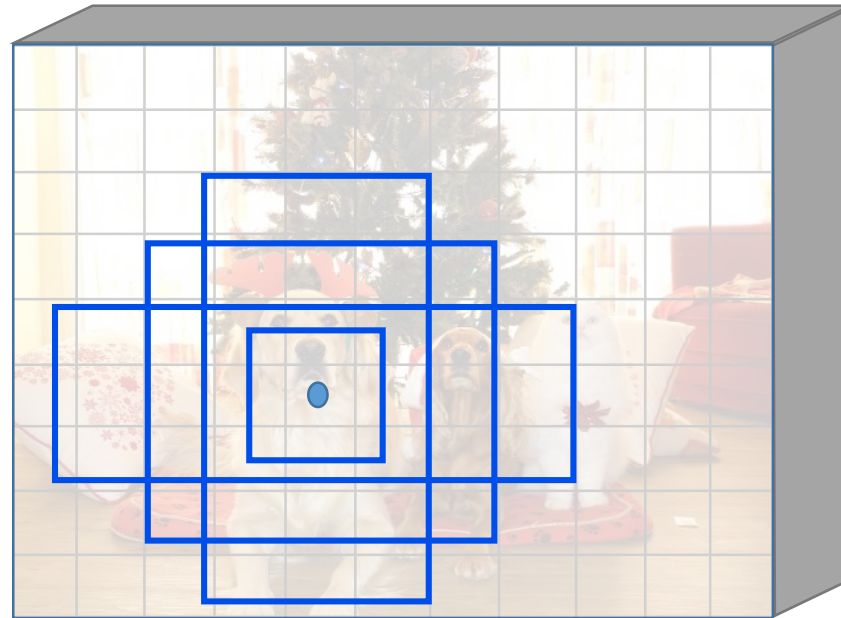

Associate \underline{K} anchor boxes of different size/scale at each point

Object Detection : Anchor-based

1. R-CNN to Fast R-CNN : run a CNN one time on the image, and extract the patches from the feature map not the image. Then small CNN for classification and box regression
2. From Fast R-CNN to Faster R-CNN : Use a Region Proposal Network (instead of handcrafted algorithm) to propose RoI



$H \times W$



$D \times \frac{H}{16} \times \frac{W}{16}$

Associate K anchor boxes of different size/scale at each point

Objectness ?


$$K \times \frac{H}{16} \times \frac{W}{16}$$

Box correction

$$4K \times \frac{H}{16} \times \frac{W}{16}$$

Object Detection : Anchor-based RPN

1. R-CNN to Fast R-CNN : run a CNN one time on the image, and extract the patches from the feature map not the image. Then small CNN for classification and box regression
2. From Fast R-CNN to Faster R-CNN : Use a Region Proposal Network (instead of handcrafted algorithm) to propose RoI





Object Detection : Anchor-based RPN

1. R-CNN to Fast R-CNN : run a CNN one time on the image, and extract the patches from the feature map not the image. Then small CNN for classification and box regression
2. From Fast R-CNN to Faster R-CNN : Use a Region Proposal Network (instead of handcrafted algorithm) to propose RoI



$H \times W$



$D \times \frac{H}{16} \times \frac{W}{16}$



Objectness ?

$$K \times \frac{H}{16} \times \frac{W}{16}$$

Box correction

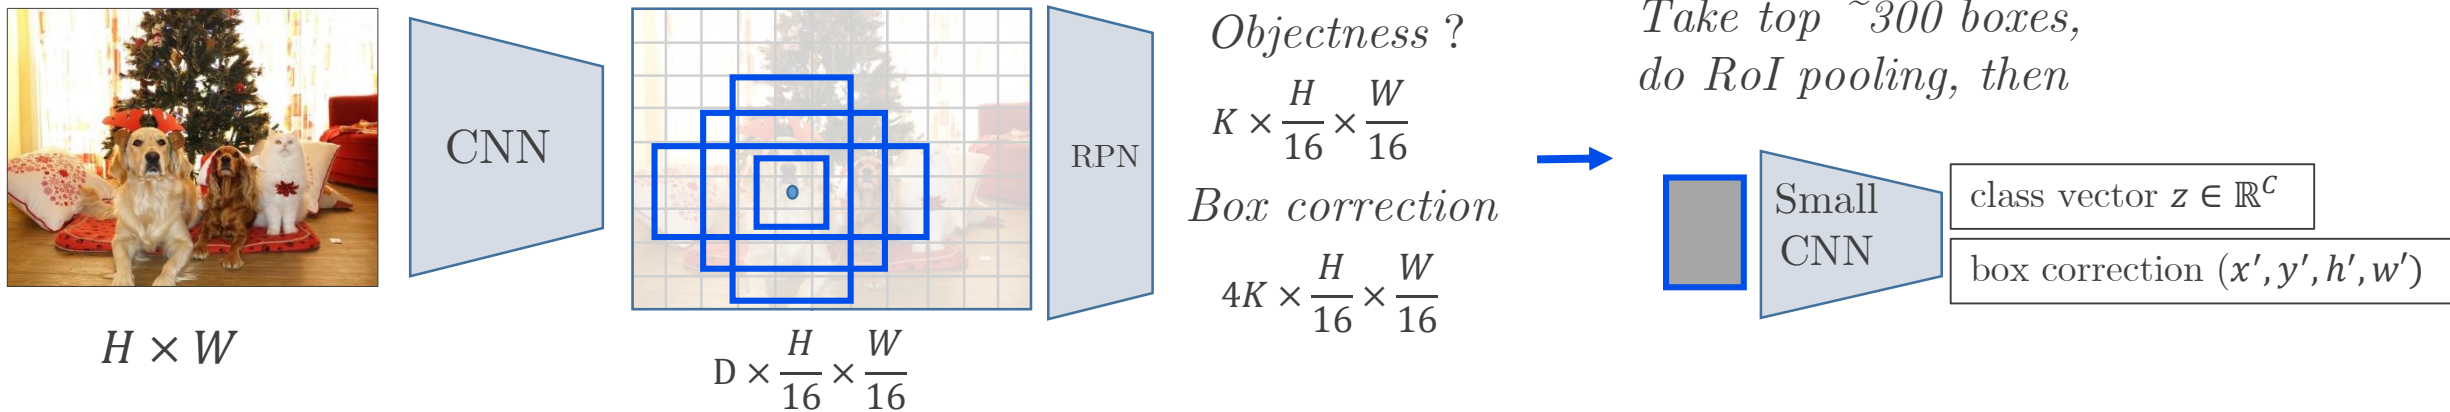
$$4K \times \frac{H}{16} \times \frac{W}{16}$$



*Take top ~300 boxes,
do RoI pooling, then*

Object Detection : Anchor-based RPN

1. R-CNN to Fast R-CNN : run a CNN one time on the image, and extract the patches from the feature map not the image. Then small CNN for classification and box regression
2. From Fast R-CNN to Faster R-CNN : Use a Region Proposal Network (instead of handcrafted algorithm) to propose RoI

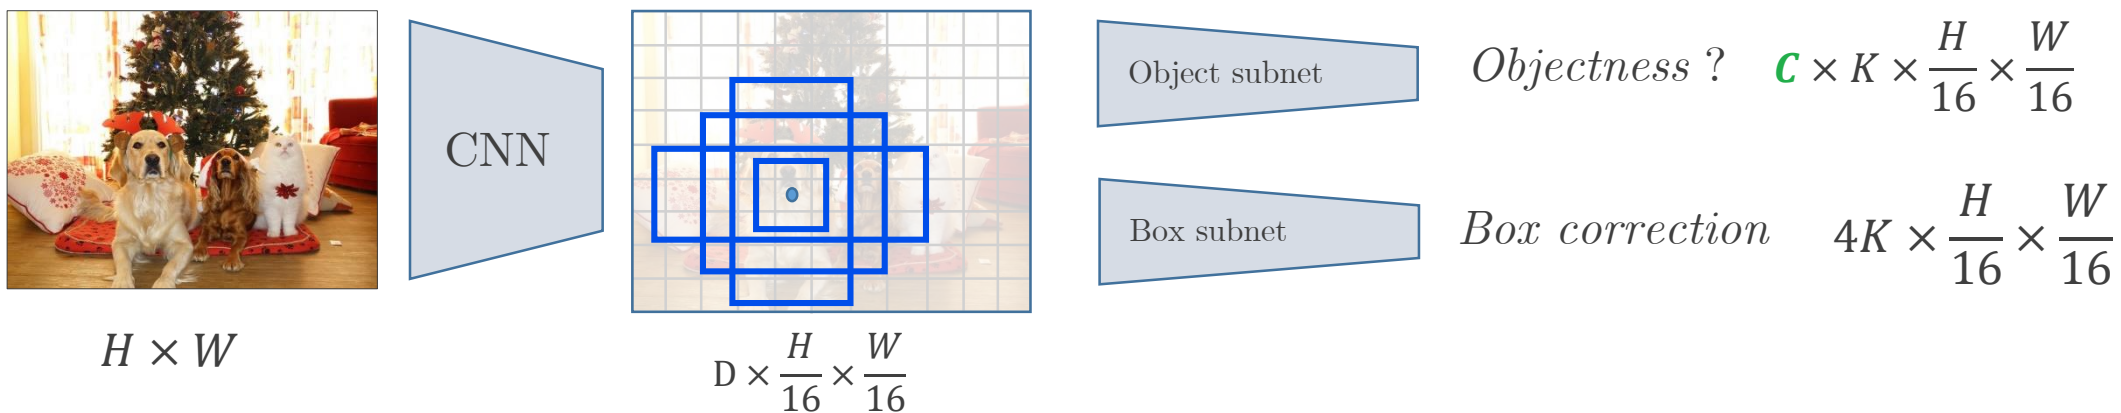


Object Detection : RetinaNet – One stage detector


1. R-CNN to Fast R-CNN : run a CNN one time on the image, and extract the patches from the feature map not the image. Then small CNN for classification and box regression
2. From Fast R-CNN to Faster R-CNN : Use a Region Proposal Network (instead of handcrafted algorithm) to propose RoI
3. From Faster R-CNN to RetinaNet : Remove the RoI pooling and do a 1-stage approach

Object Detection : RetinaNet – One stage detector

1. R-CNN to Fast R-CNN : run a CNN one time on the image, and extract the patches from the feature map not the image. Then small CNN for classification and box regression
2. From Fast R-CNN to Faster R-CNN : Use a Region Proposal Network (instead of handcrafted algorithm) to propose RoI
3. From Faster R-CNN to RetinaNet : Remove the RoI pooling and do a 1-stage approach




Object Detection : Non-Maximum Suppression



Skipped a lot of modelisation details :

- NMS or *non-maximum suppression*

Object Detection : Non-Maximum Suppression



Skipped a lot of modelisation details :


- NMS or *non-maximum suppression*

Algorithm 1 Non-Max Suppression


```

1: procedure NMS( $B, c$ )
2:    $B_{nms} \leftarrow \emptyset$ 
3:   for  $b_i \in B$  do
4:      $discard \leftarrow \text{False}$ 
5:     for  $b_j \in B$  do
6:       if  $\text{IoU}(b_i, b_j) > \lambda_{nms}$  then
7:         if  $\text{score}(c, b_j) > \text{score}(c, b_i)$  then
8:            $discard \leftarrow \text{True}$ 
9:       if not  $discard$  then
10:         $B_{nms} \leftarrow B_{nms} \cup b_i$ 
11:  return  $B_{nms}$ 

```




Object Detection : Non-Maximum Suppression




Skipped a lot of modelisation details :

- NMS or *non-maximum suppression*



Object Detection : Non-Maximum Suppression



Skipped a lot of modelisation details :


- NMS or *non-maximum suppression*

Algorithm 1 Non-Max Suppression


```

1: procedure NMS( $B, c$ )
2:    $B_{nms} \leftarrow \emptyset$ 
3:   for  $b_i \in B$  do
4:      $discard \leftarrow \text{False}$ 
5:     for  $b_j \in B$  do
6:       if  $\text{IoU}(b_i, b_j) > \lambda_{nms}$  then
7:         if  $\text{score}(c, b_j) > \text{score}(c, b_i)$  then
8:            $discard \leftarrow \text{True}$ 
9:       if not  $discard$  then
10:         $B_{nms} \leftarrow B_{nms} \cup b_i$ 
11:  return  $B_{nms}$ 

```



Object Detection : Non-Maximum Suppression




Skipped a lot of modelisation details :

- NMS or *non-maximum suppression*


Algorithm 1 Non-Max Suppression

```

1: procedure NMS( $B, c$ )
2:    $B_{nms} \leftarrow \emptyset$ 
3:   for  $b_i \in B$  do
4:      $discard \leftarrow \text{False}$ 
5:     for  $b_j \in B$  do
6:       if  $\text{IoU}(b_i, b_j) > \lambda_{nms}$  then
7:         if  $\text{score}(c, b_j) > \text{score}(c, b_i)$  then
8:            $discard \leftarrow \text{True}$ 
9:       if not  $discard$  then
10:         $B_{nms} \leftarrow B_{nms} \cup b_i$ 
11:   return  $B_{nms}$ 
  
```




Object Detection : Non-Maximum Suppression



Skipped a lot of modelisation details :

- NMS or *non-maximum suppression*


Object Detection : Non-Maximum Suppression



Skipped a lot of modelisation details :

- NMS or *non-maximum suppression*
- Solve *imbalance* between negative and positive anchors, *i.e* among $K \sim 50000$ anchors most of them do not bind an object


Object Detection : Foreground-Background imbalance



Skipped a lot of modelisation details :

- NMS or *non-maximum suppression*
- Solve *imbalance* between negative and positive anchors, *i.e* among $K \sim 50000$ anchors most of them do not bind an object
 - Sampling heuristics: fixed foreground-background ratio of anchors, online hard example mining


Object Detection : Focal loss



Skipped a lot of modelisation details :

- NMS or *non-maximum suppression*
- Solve *imbalance* between negative and positive anchors, *i.e* among $K \sim 50000$ anchors most of them do not bind an object
 - Sampling heuristics: fixed foreground-background ratio of anchors, online hard example mining
 - Focal loss for 1-stage detectors is the turning point

Object Detection : Focal loss



Skipped a lot of modelisation details :

- NMS or *non-maximum suppression*
- Solve *imbalance* between negative and positive anchors, *i.e* among $K \sim 50000$ anchors most of them do not bind an object
 - Sampling heuristics: fixed foreground-background ratio of anchors, online hard example mining
 - Focal loss for 1-stage detectors is the turning point




Figure 1. We propose a novel loss we term the *Focal Loss* that adds a factor $(1 - p_t)^\gamma$ to the standard cross entropy criterion. Setting $\gamma > 0$ reduces the relative loss for well-classified examples ($p_t > .5$), putting more focus on hard, misclassified examples. As our experiments will demonstrate, the proposed focal loss enables training highly accurate dense object detectors in the presence of vast numbers of easy background examples.

Object Detection : Other architectures

Subsequent anchor-based methods:

- 2017, He et al, Mask R-CNN → perform instance and semantic segmentation
- 2020, Bochkovskiy et al, YOLOv4: Optimal Speed and Accuracy of Object Detection
- 2020, Tan et al, EfficientDet: Scalable and Efficient Object Detection

Object Detection : Other architectures

Subsequent anchor-based methods:

- 2017, He et al, Mask R-CNN → perform instance and semantic segmentation
- 2020, Bochkovskiy et al, YOLOv4: Optimal Speed and Accuracy of Object Detection
- 2020, Tan et al, EfficientDet: Scalable and Efficient Object Detection

Anchor-Free detectors : avoid cumbersome modelisation, NMS, etc..

CNN-based architecture :

- 2019, Tian et al, FCOS: Fully Convolutional One-Stage Object Detection
- 2019, Zhou et al, Objects as points.

Object Detection : Other architectures

Subsequent anchor-based methods:

- 2017, He et al, Mask R-CNN → perform instance and semantic segmentation
- 2020, Bochkovskiy et al, YOLOv4: Optimal Speed and Accuracy of Object Detection
- 2020, Tan et al, EfficientDet: Scalable and Efficient Object Detection

Anchor-Free detectors : avoid cumbersome modelisation, NMS, etc..

CNN-based architecture :

- 2019, Tian et al, FCOS: Fully Convolutional One-Stage Object Detection
- 2019, Zhou et al, Objects as points.

Transformers-based:

- 2020, Carion et al, DETR: End-to-End Object Detection with Transformers
- 2023, Zong et al, DETRs with Collaborative Hybrid Assignments Training

Thank you for your
attention



*slides adapted from [CS231n](#)