



GluonCV: Segmentation

Tong He

Applied Scientist, Amazon Web Services.


2018.12.18



Introduction

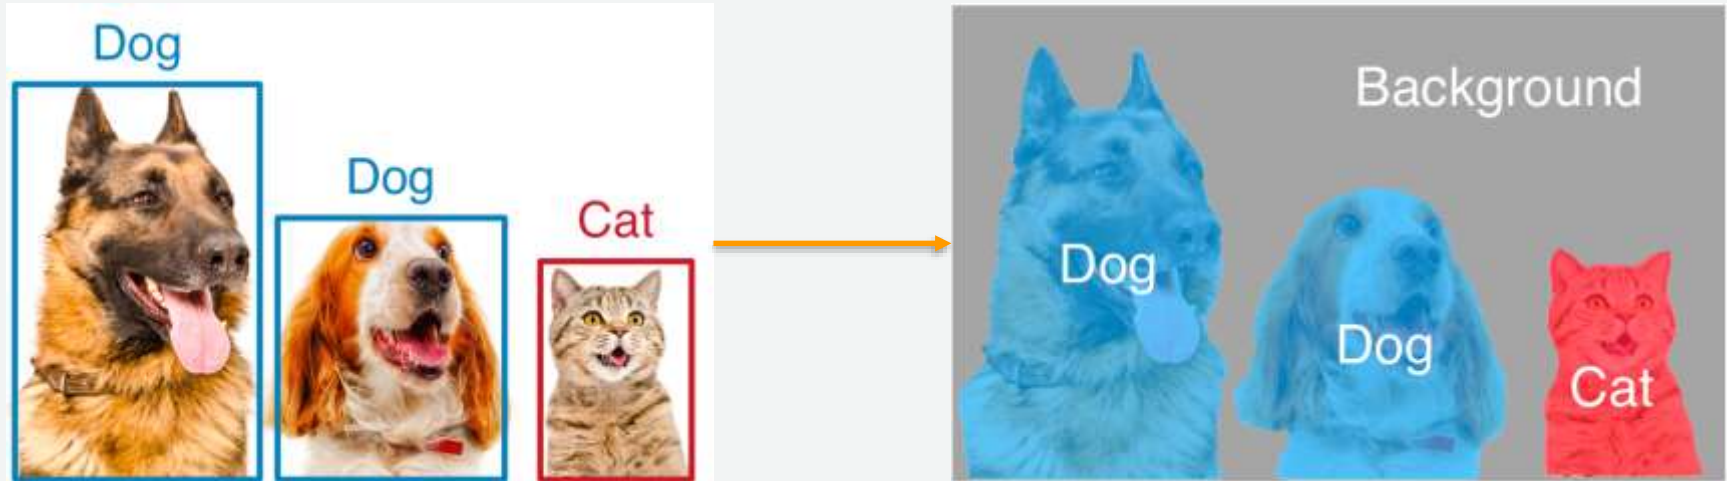
What is segmentation?

Classify and mask the pixels of the object




Why Segmentation?

Precise information

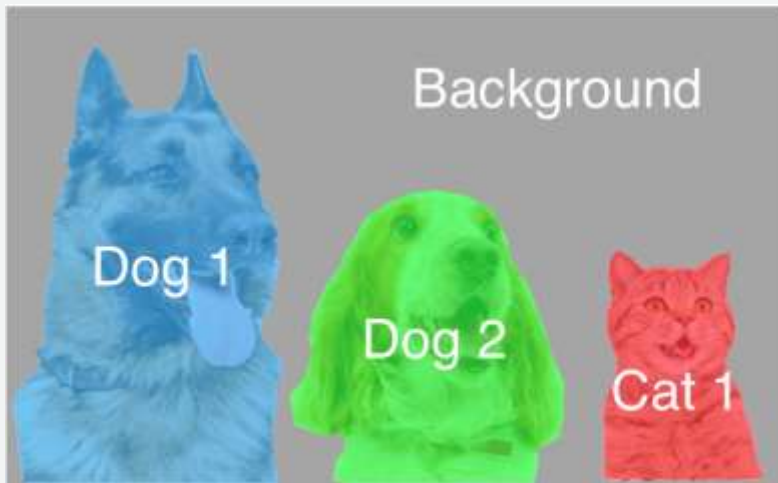


Segmentation

Semantic Segmentation



Instance Segmentation



Segmentation with Deep Learning

Segmentation with Deep Learning

Comparison

Classification:

Input: image

Output: class label

Evaluation: accuracy

Detection:

Input: image

Output: boxes with labels

Evaluation: IoU

Segmentation:


Input: image

Output: mask with label

Evaluation: mIoU, pixAcc

Segmentation with Deep Learning


Mask



Segmentation with Deep Learning

IoU

- Similar to Detection



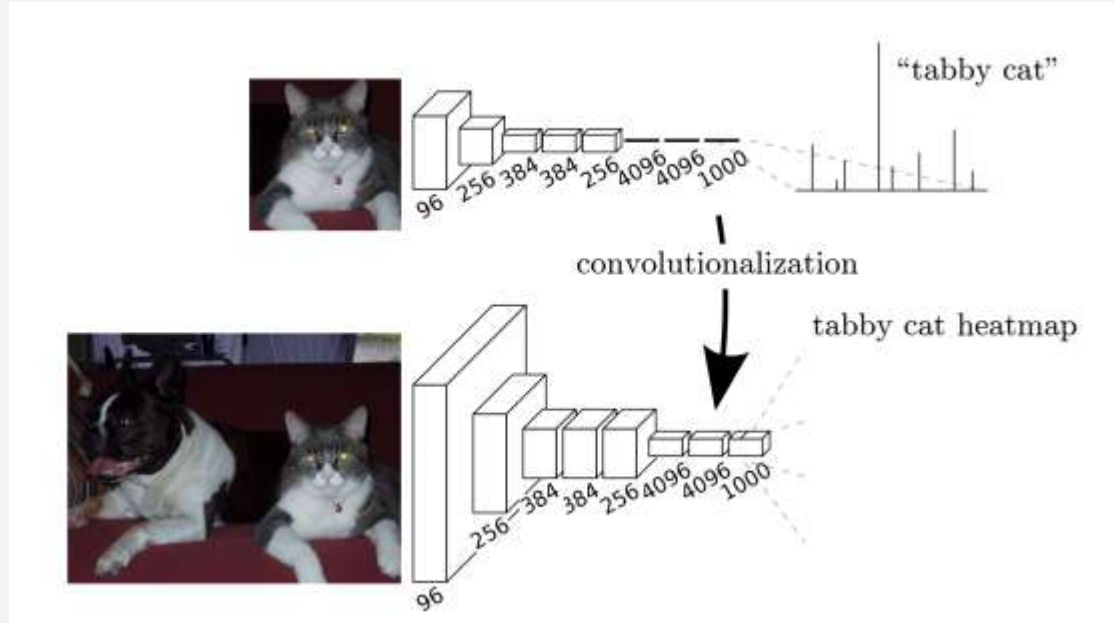
Segmentation with Deep Learning

Popular models

- Semantic Segmentation
 - FCN
 - PSPNet
 - DeepLab v3
- Instance Segmentation
 - Mask R-CNN

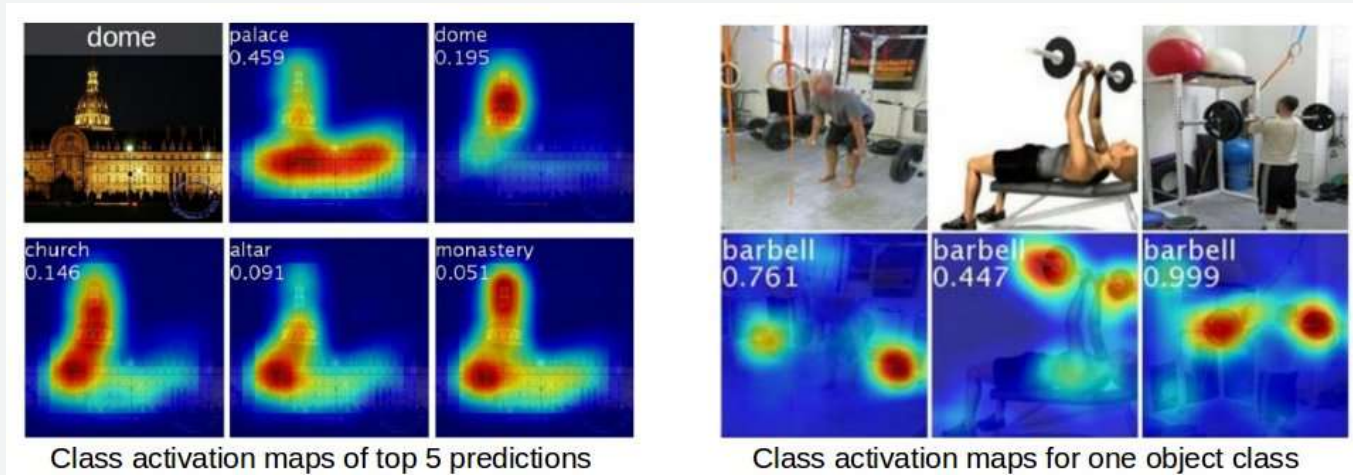
Segmentation with Deep Learning

FCN: Fully Convolutional Network



Segmentation with Deep Learning


FCN: Fully Convolutional Network



Segmentation with Deep Learning

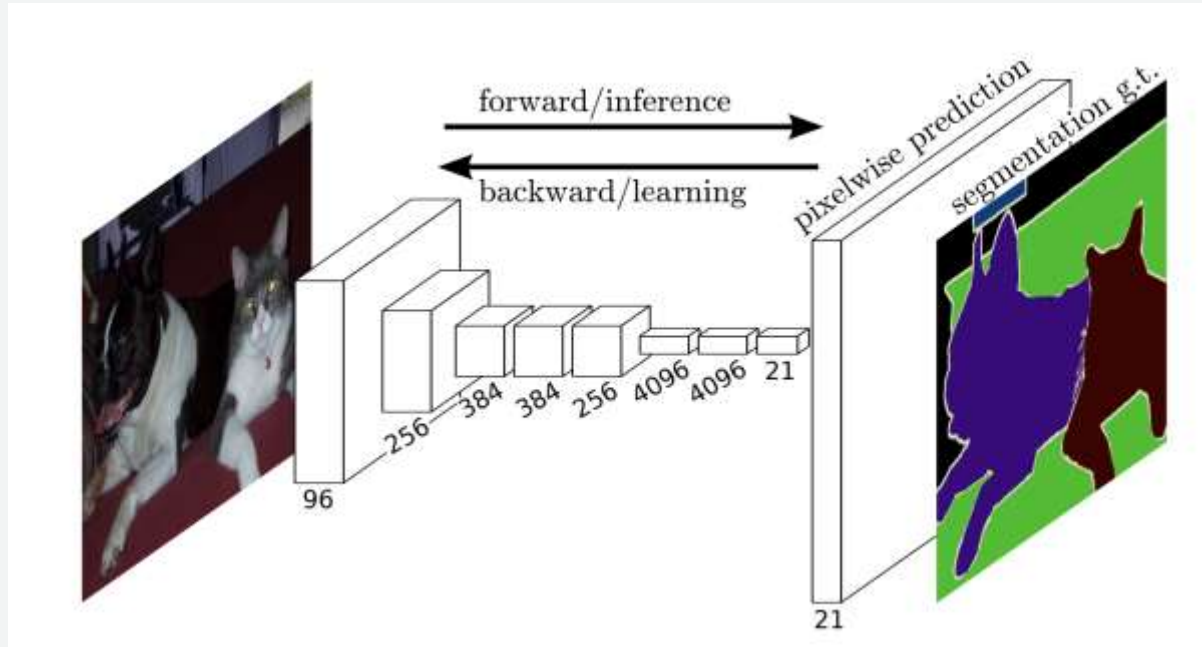
Upsampling

- Small to large feature map




Segmentation with Deep Learning

FCN: Fully Convolutional Network




Segmentation with Deep Learning

Dilated Convolution



Normal



Dilated


Segmentation with Deep Learning

Dilated Convolution


- Preserve information
- Keep feature map size
- Larger reception field

Segmentation with Deep Learning


Dilated Residual Network



(a) ResNet

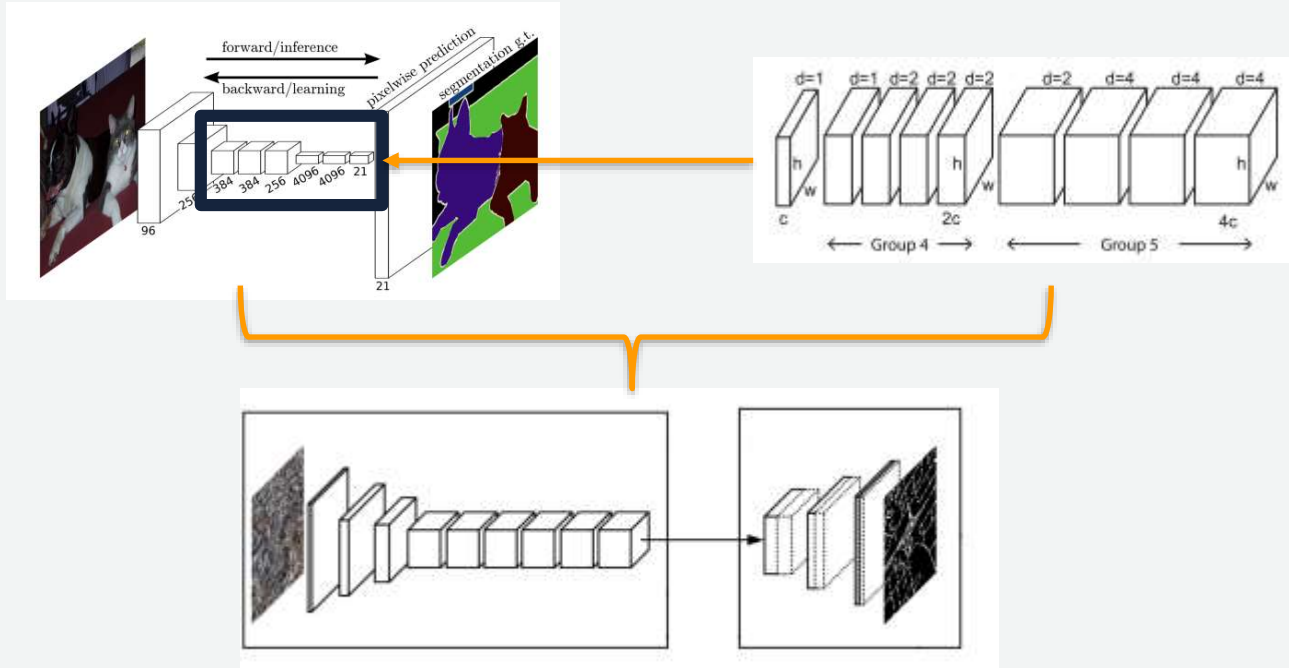


(b) DRN



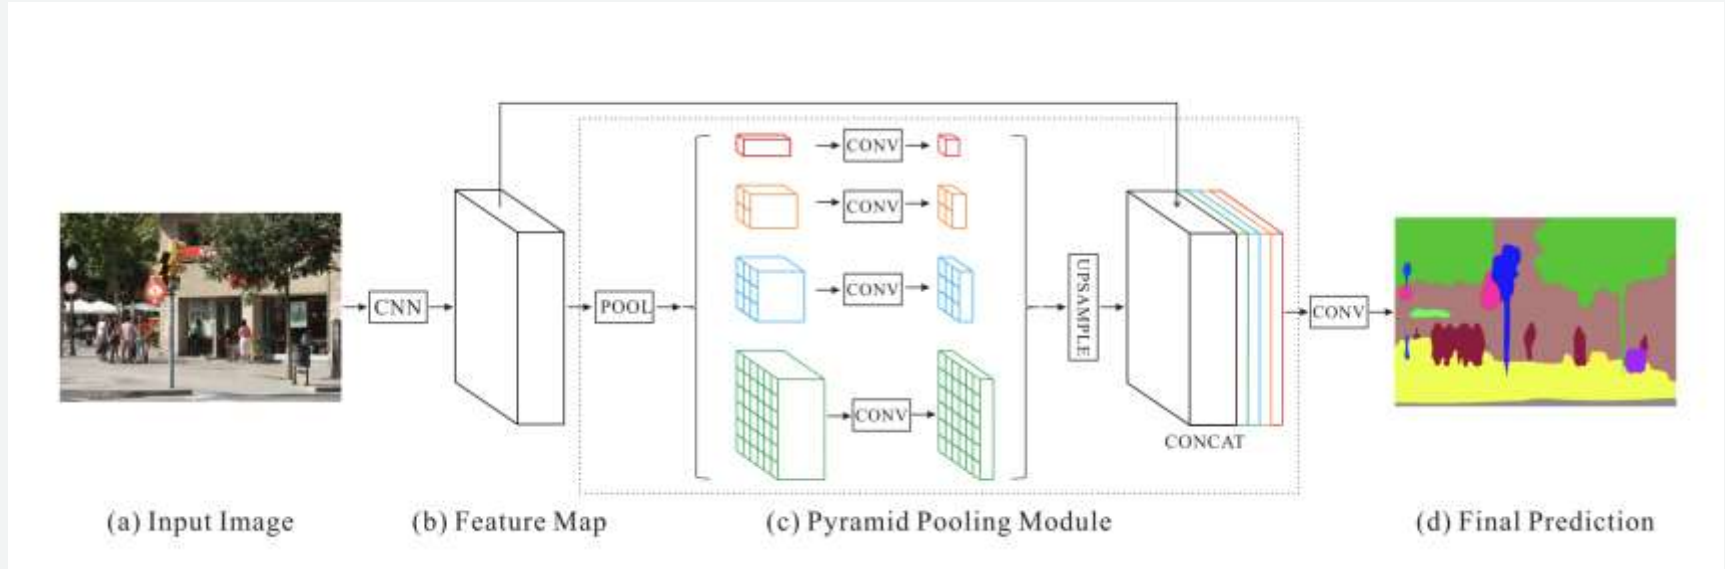
Segmentation with Deep Learning

FCN + Dilation



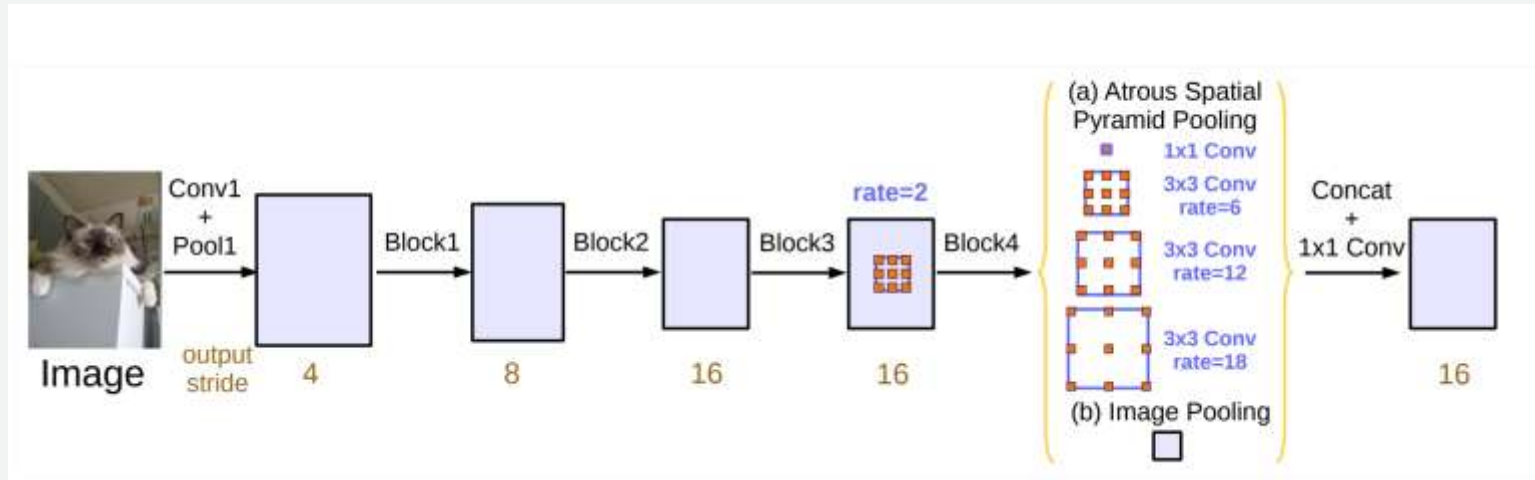
Segmentation with Deep Learning

PSP: Pyramid Scene Parsing Network



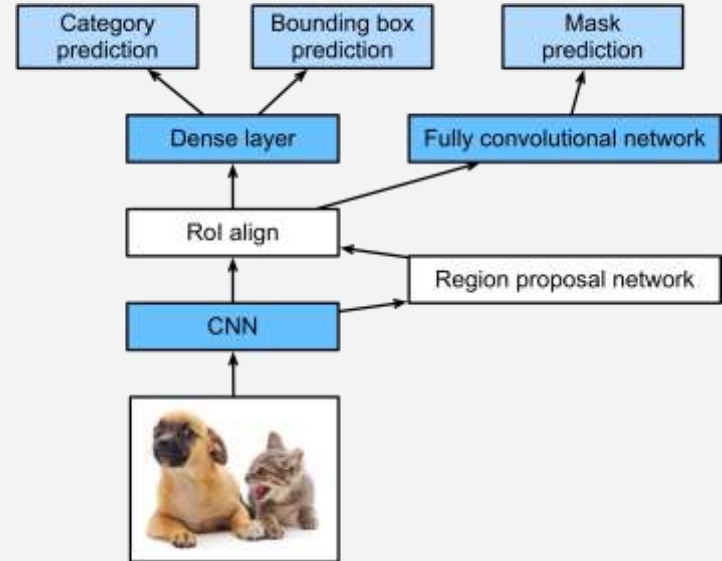
Segmentation with Deep Learning

DeepLab V3



Segmentation with Deep Learning

Mask R-CNN



Segmentation with Deep Learning

Summary

- FCN
 - easy to implement, less accurate
- PSPNet
 - Slower, more accurate
- DeepLab
 - Even slower, very accurate

Segmentation with GluonCV

Segmentation with GluonCV

Our Office in Palo Alto, California



Segmentation with GluonCV

Model Zoo

DeepLab-v3	mIoU on VOC	86.7%	85.7% (paper)
Mask-RCNN	mask AP on COCO	33.1%	32.8% (Detectron)

Segmentation with GluonCV

Advanced Training

- Learning Rate Schedule
- Multi-Transfer Learning
 - MS COCO -> Pascal VOC Augmented -> Pascal VOC

Segmentation with GluonCV

Multi-Transfer Learning

- DeepLab V3
 - MS COCO -> Pascal VOC Augmented -> Pascal VOC

Segmentation with GluonCV

Resources:

- Model Zoo: https://gluon-cv.mxnet.io/model_zoo/segmentation.html
- Tutorials: https://gluon-cv.mxnet.io/build/examples_segmentation/index.html
- Deep Learning Book: <http://en.diveintodeeplearning.org/>

Segmentation with GluonCV

Hands On!