# CSE 152: Computer Vision

Hao Su

## Lecture 10: Object Recognition

# How do we represent objects

- Bounding box

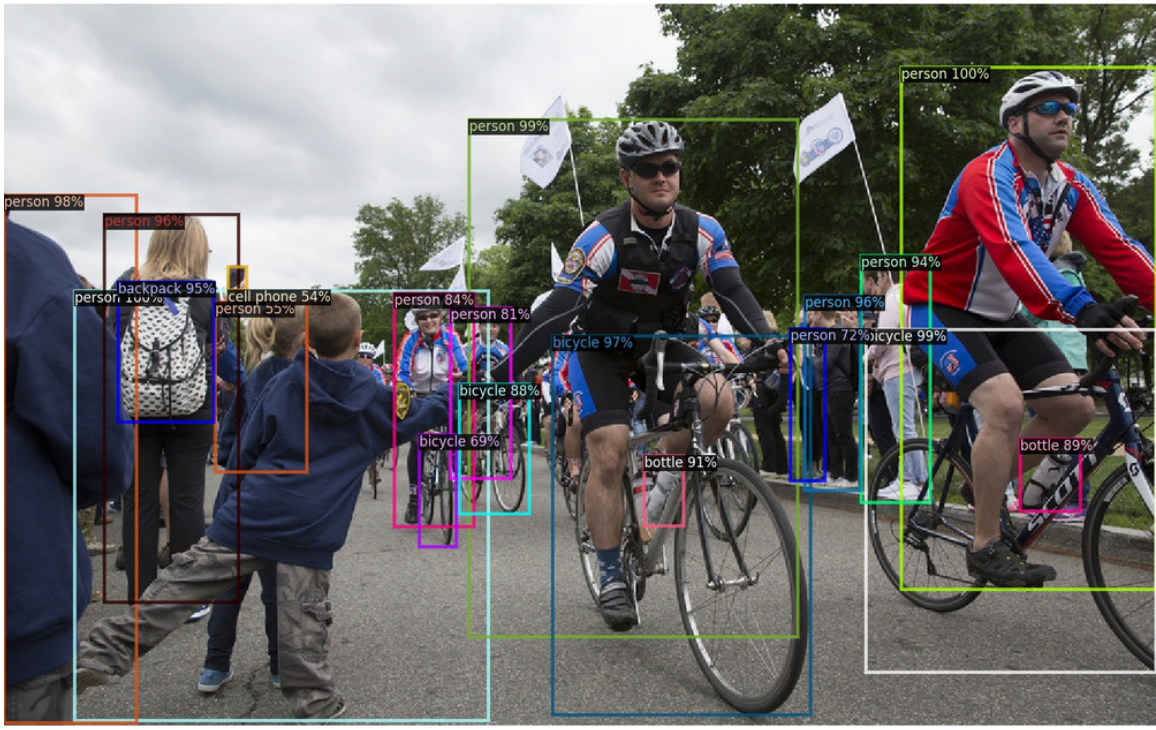


Figures from https://github.com/facebookresearch/detectron2

# How do we represent objects

- Bounding box
- Instance mask

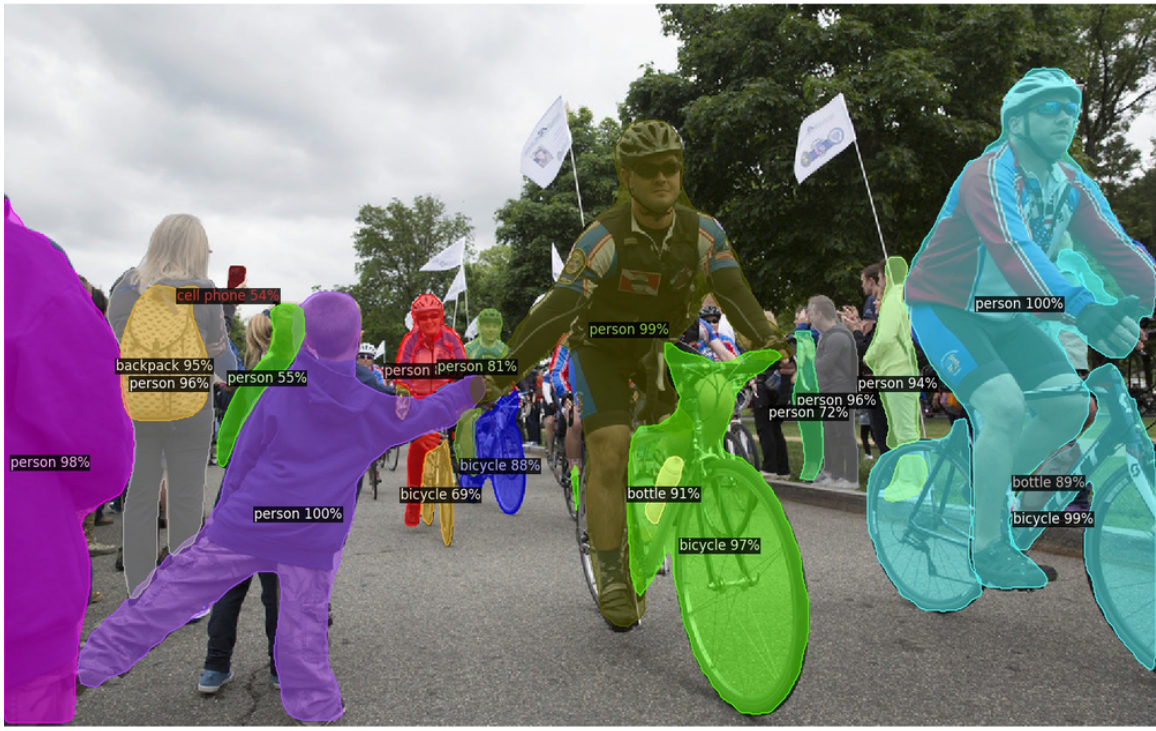


Figures from https://github.com/facebookresearch/detectron2

# How do we represent objects

- Bounding box
- Instance mask
- Keypoint

Figures from https://github.com/facebookresearch/detectron2

# How do we represent objects

- **Bounding box**
- Instance mask
- Keypoint

Figures from https://github.com/facebookresearch/detectron2

# Object Detection with Bounding Boxes

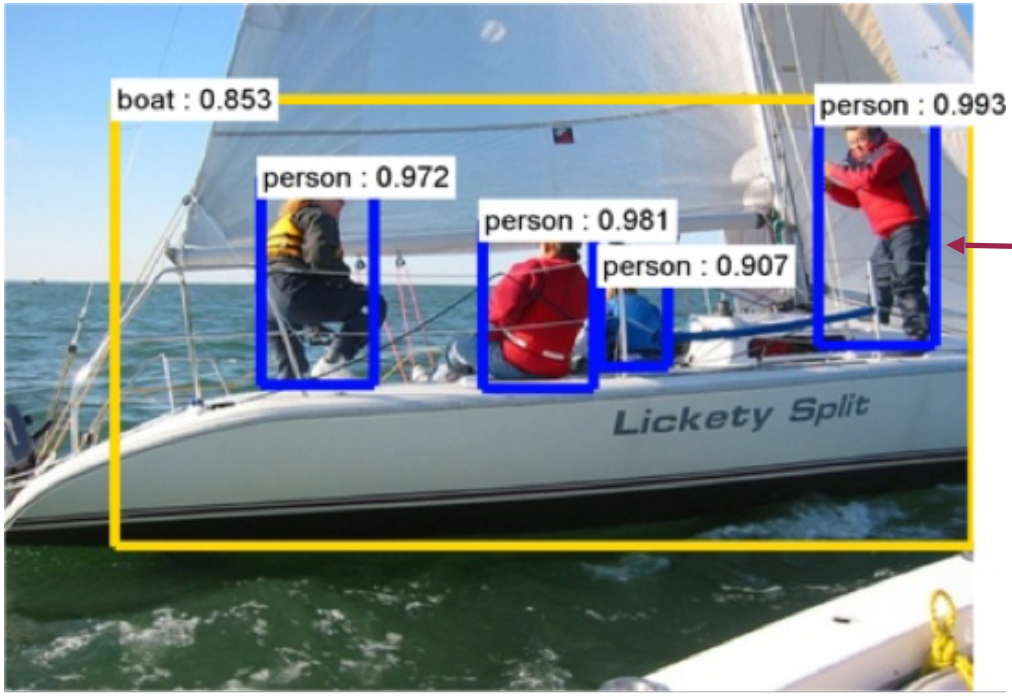


What? - Recognition/ Classification

Where? - Localization/ Regression

"Object detection"

Slides modified from Ross Girshick tutorial at CVPR 2019

# Object Detection with Segmentation Masks


"Instance segmentation"

What? - Recognition

Where? - Segmentation

Slides modified from Ross Girshick tutorial at CVPR 2019

# Semantic Segmentation

Predict a pixel-wise class label

Stuff: walls, buildings, sky, road

Things: human, cars, bikes

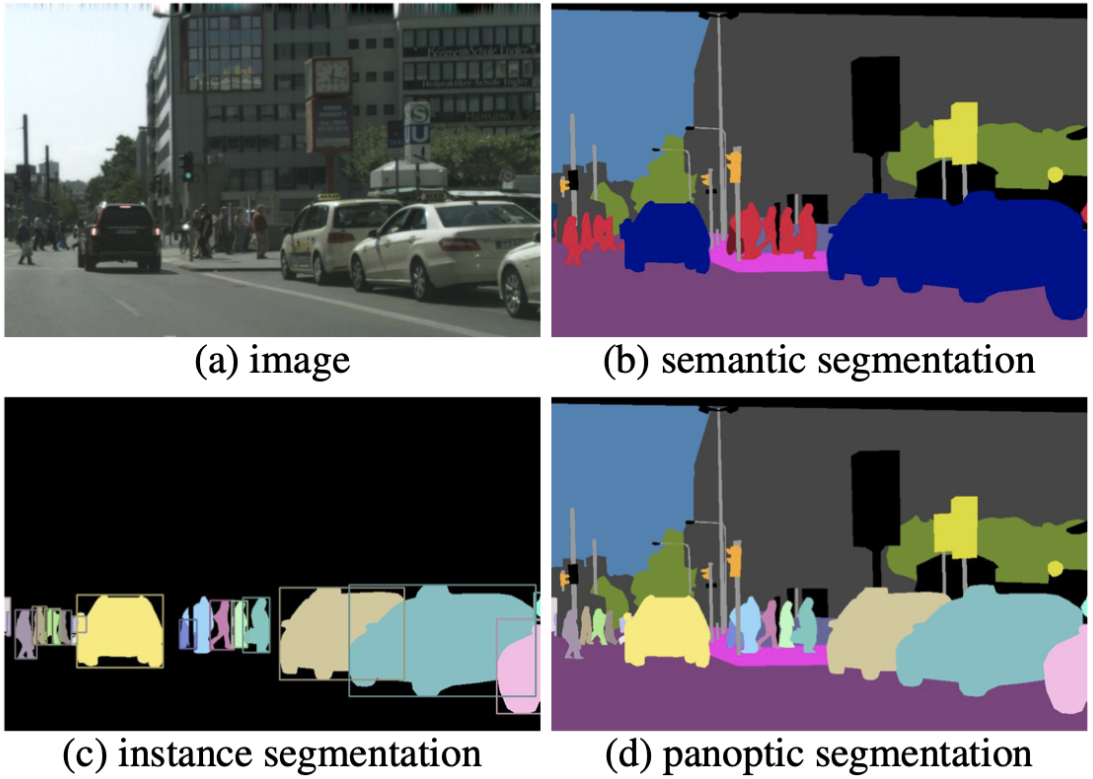

(a) image (b) semantic segmentation

(c) instance segmentation (d) panoptic segmentation

Figures from *Panoptic Segmentation*, CVPR 2019

# Datasets



Microsoft
COCO

**Visual Object Classes Challenge 2012 (VOC2012)**

# Object Detection

# Object Detection → Object Classification

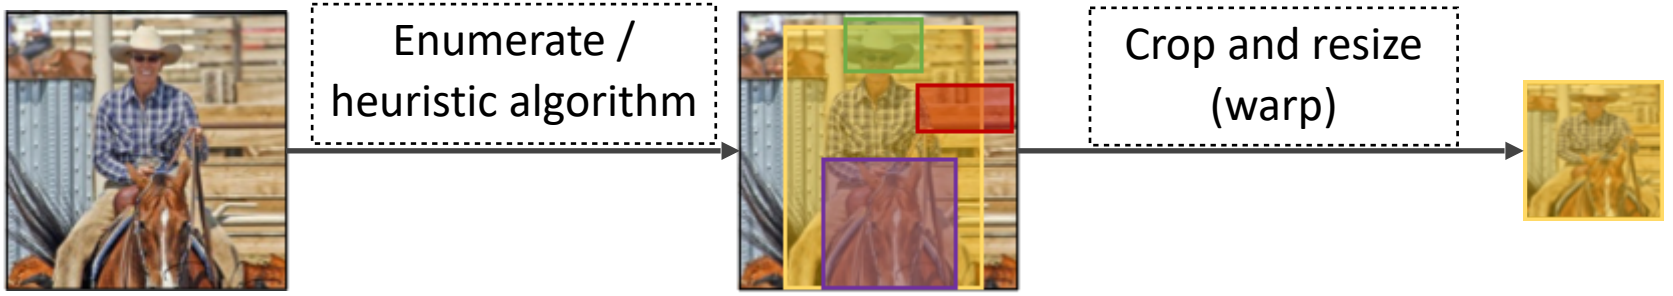


Input: an image

Proposals/Candidates

Cropped image

We've already reduced object detection to object classification!

# R-CNN (Regional ConvNet)

Computationally expensive

Input: an image

Enumerate / heuristic algorithm

Proposals/Candidates

Cropped image Region of Interests (RoI)

ConvNet

Class Probability

How probable is it a human?

BBox Regression

How can we modify this bounding box?

Slides modified from Ross Girshick tutorial at CVPR 2019

# Faster R-CNN

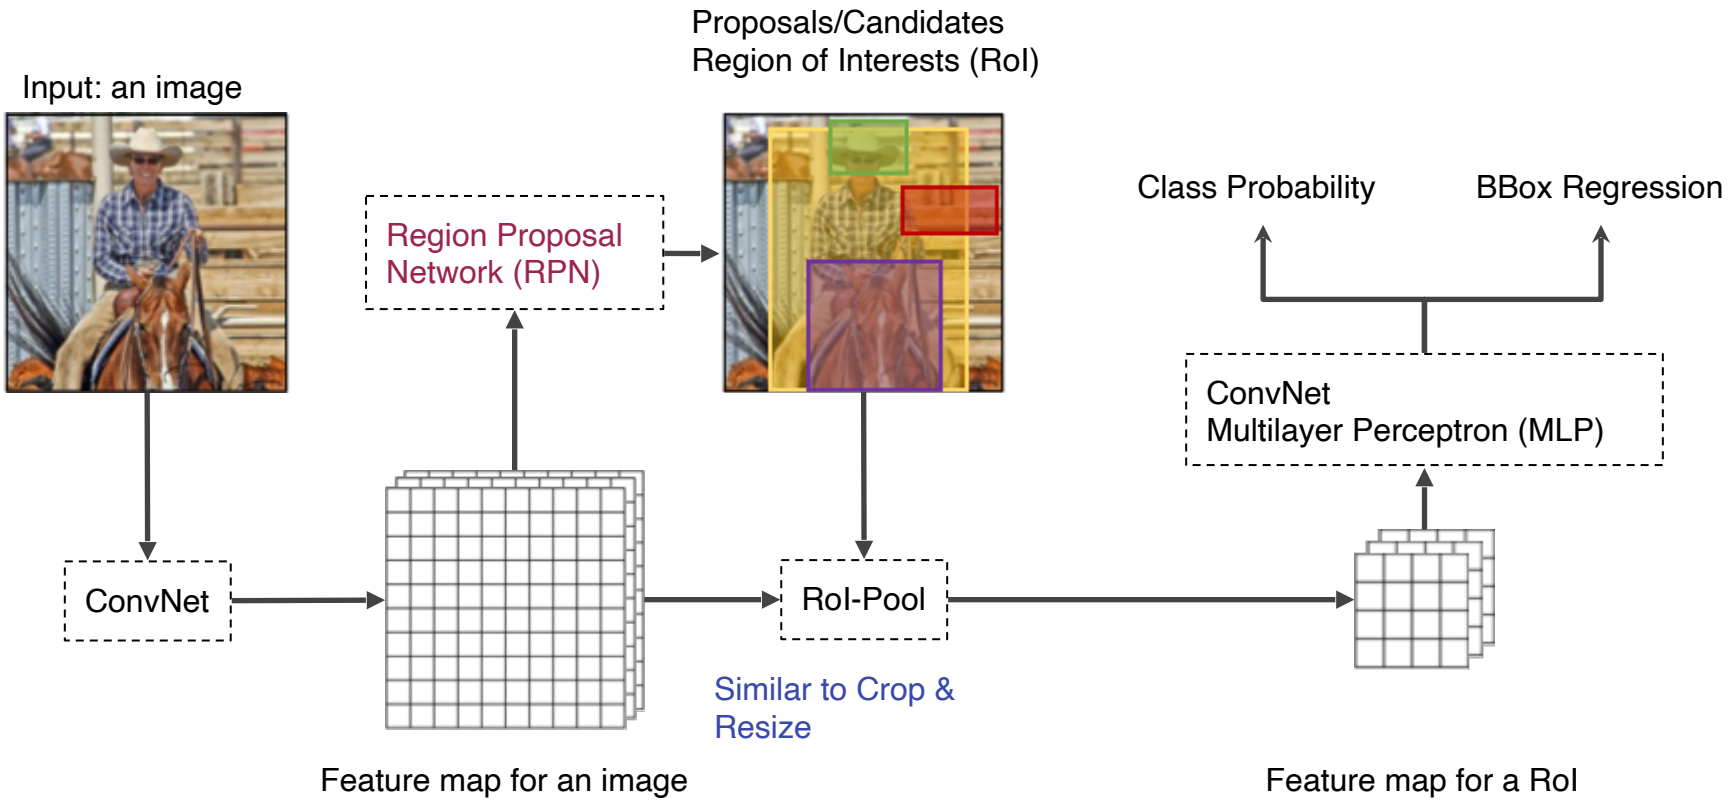
Input: an image
Proposals/Candidates
Region of Interests (RoI)
Region Proposal
Network (RPN)
Class Probability
BBox Regression
ConvNet
Multilayer Perceptron (MLP)
ConvNet
RoI-Pool
Similar to Crop &
Resize
Feature map for an image
Feature map for a RoI

# Faster R-CNN

- At each location, consider boxes of many different sizes and aspect ratios

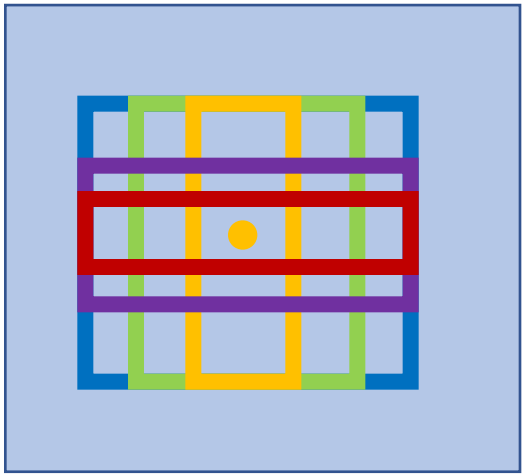

# Faster R-CNN

- At each location, consider boxes of many different sizes and aspect ratios

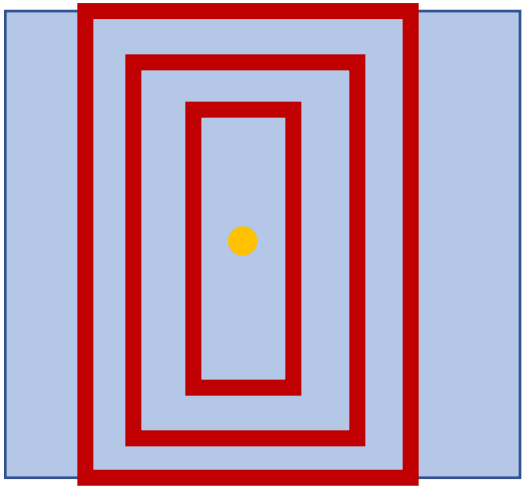

# Object Segmentation

# Semantic Segmentation Idea: Fully Convolutional

Design a network as a bunch of convolutional layers to make predictions for pixels all at once!

Input: 3 x H x W

Conv → Conv → Conv → Conv

Convolutions: D x H x W

Scores: C x H x W

argmax

Predictions: H x W

# Semantic Segmentation Idea: Fully Convolutional

Design network as a bunch of convolutional layers, with **downsampling** and **upsampling** inside the network!

Input: 3 x H x W

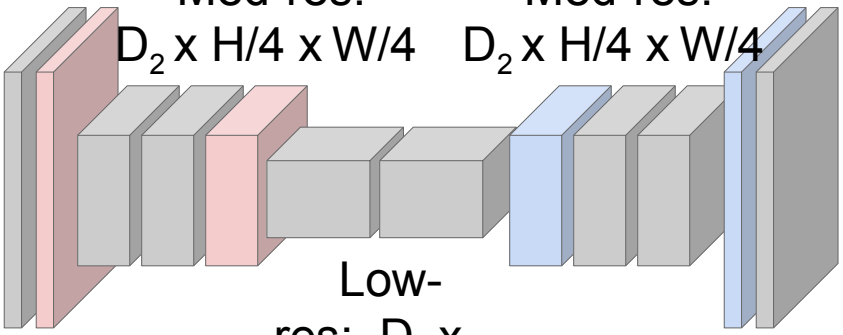


Predictions: H x W

Long, Shelhamer, and Darrell, "Fully Convolutional Networks for Semantic Segmentation", CVPR 2015
Noh et al, "Learning Deconvolution Network for Semantic Segmentation", ICCV 2015

# Semantic Segmentation Idea: Fully Convolutional

Design network as a bunch of convolutional layers, with **downsampling** and **upsampling** inside the network!

**Downsampling**:
Pooling, strided convolution

**Upsampling**:
???

Input:
3 x H x W

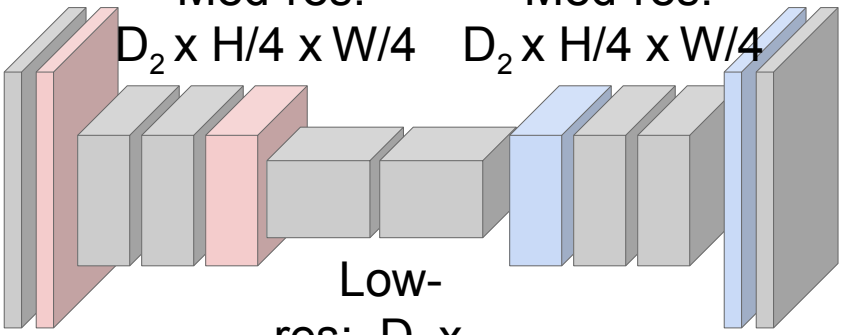


High-res: $D_1$ x H/2 x W/2

High-res: $D_1$ x H/2 x W/2

Predictions: H x W

Long, Shelhamer, and Darrell, "Fully Convolutional Networks for Semantic Segmentation", CVPR 2015
Noh et al, "Learning Deconvolution Network for Semantic Segmentation", ICCV 2015

# Learnable Upsampling: Transpose Convolution

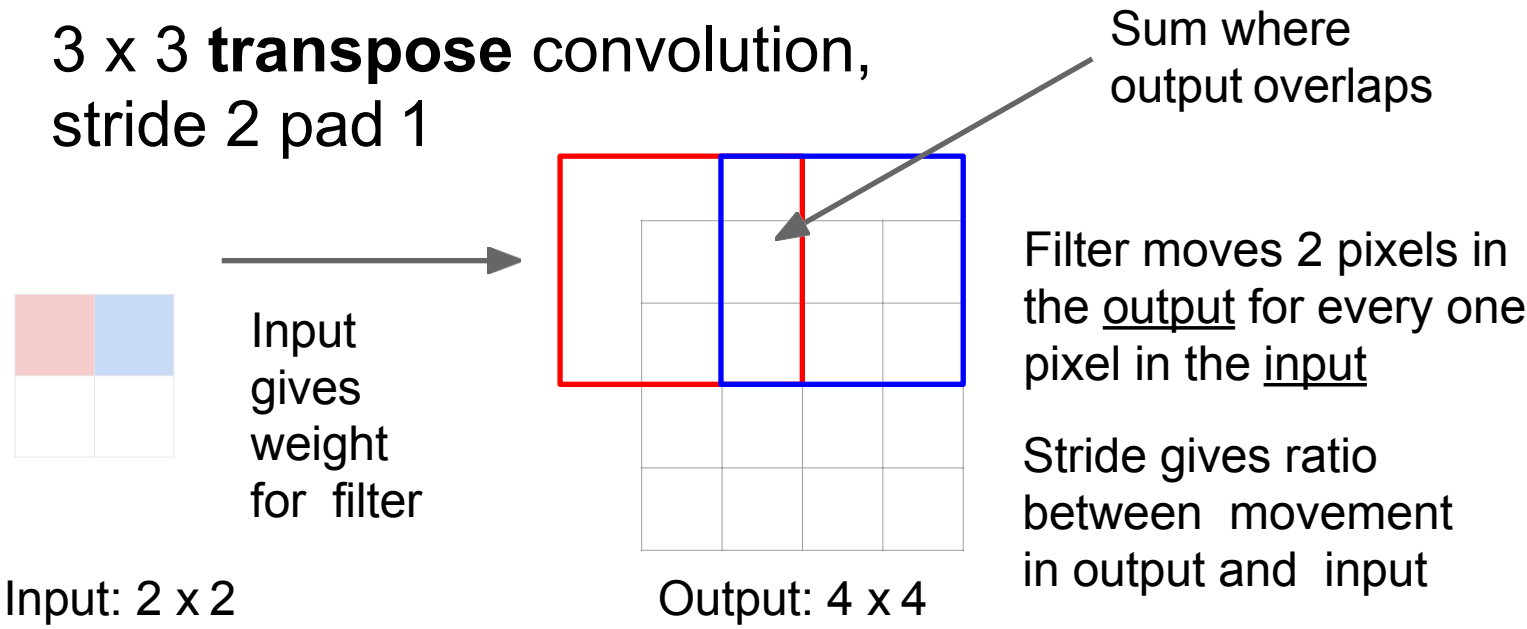

# Learnable Upsampling: Transpose Convolution

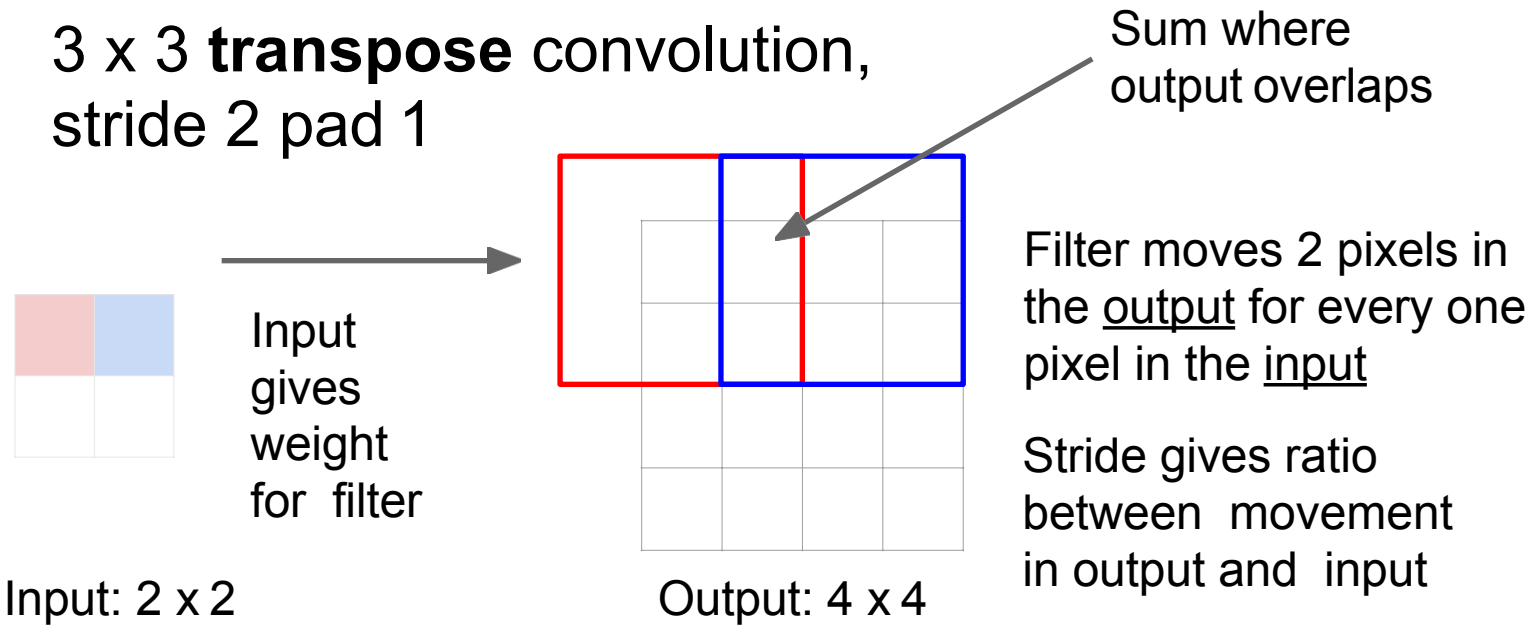


**Other names:**
-Deconvolution (bad)
-Upconvolution
-Fractionally strided convolution
-Backward strided convolution

# Semantic vs. Instance Segmentation

Object detection


Semantic segmentation

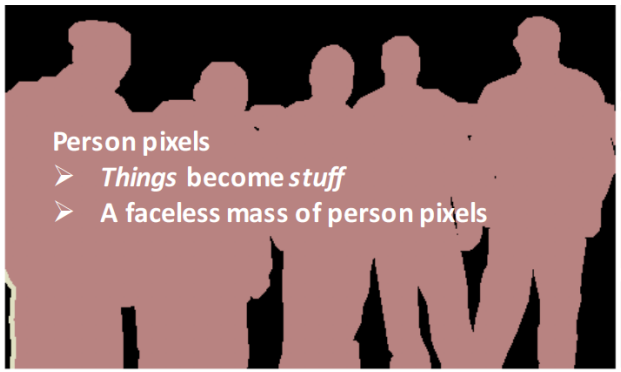


Instance segmentation

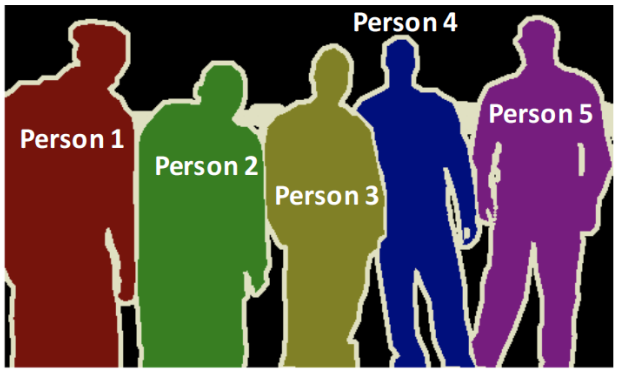


Slides modified from Ross Girshick tutorial at CVPR 2019

# Mask R-CNN

- First do object detection using the Faster R-CNN arch, and then do semantic segmentation inside the cropped region

- Share features of the first few layers for detection and segmentation