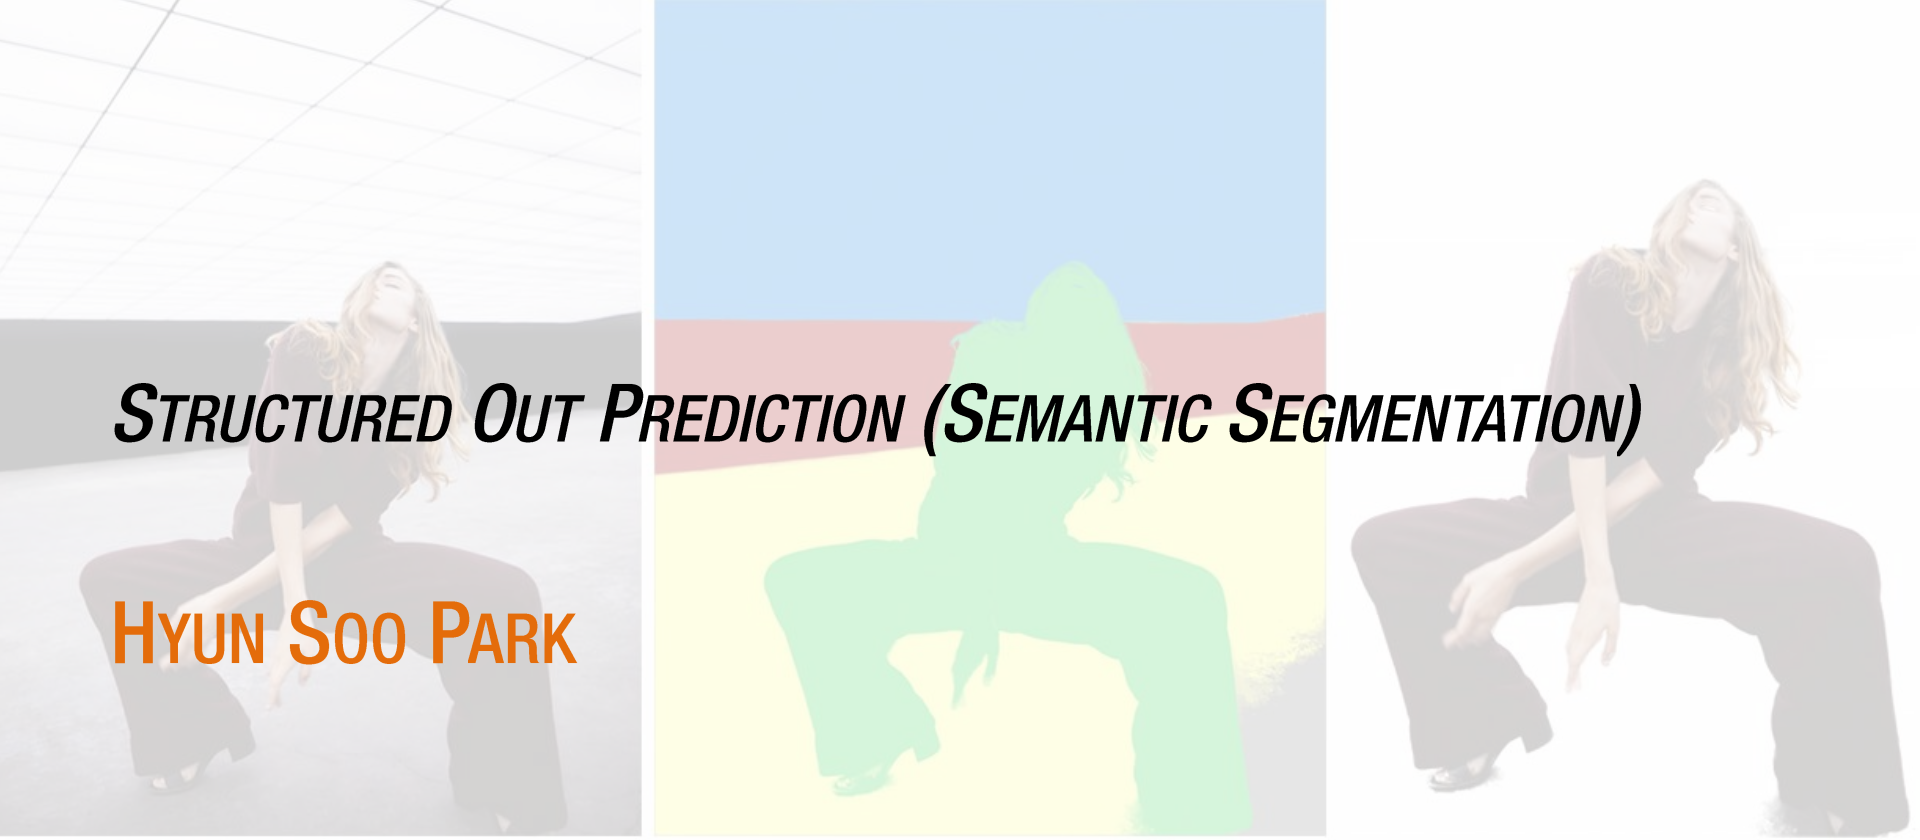

Structured Out Prediction (Semantic Segmentation)

Hyun Soo Park

Challenges of Visual Recognition

- Appearance
  - DOF: texture, illumniation, material, shading, …
- Shape
  - DOF: object category, geometric pose, viewpoint, …

Human

$$f(I) = l_{human}$$

Semantic segmentation

SEMANTIC SEGMENTATION: PIXEL CLASSIFICATION

0: Background/Unknown
1: Person
2: Purse
3: Plants/Grass
4: Sidewalk
5: Building/Structures

SEMANTIC SEGMENTATION FORMULATION

Unsupervised superpixel segmentation

x_i

Visual feature

- Color histogram
- BoW
- SIFT
- HOG

Classification

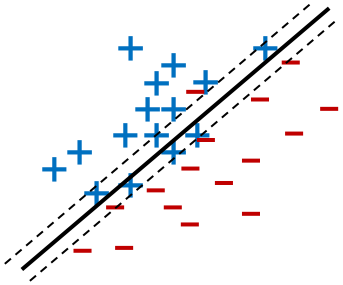

$L = \phi(x_i)$ e.g., $\phi(x_i) = w_{tree} \cdot x_i$

SEMANTIC SEGMENTATION FORMULATION

Unsupervised superpixel segmentation

x_i

x_j

Visual feature

- Color histogram
- BoW
- SIFT
- HOG

Classification

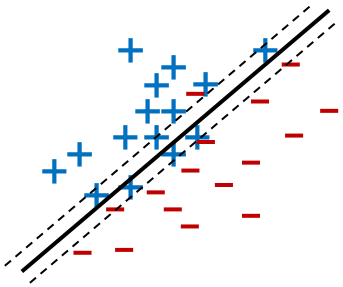

$$L = \phi(x_i) + \underbrace{\varphi(x_i, x_j)}_{\text{Context}}$$

CRF: Conditional Random Field
aka. joint classifications (structured pred.)

HOLISTIC PREDICTION VIA DEEP LEARNING

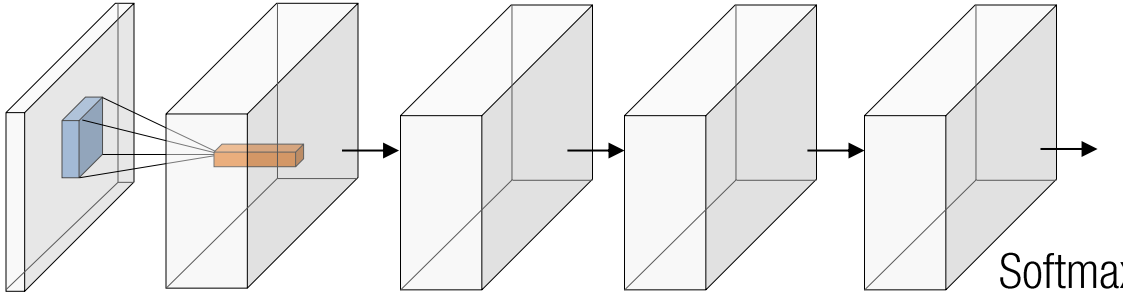

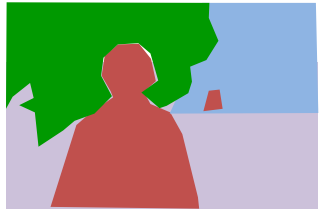

HxWx3

HxWxD

Softmax

HxWxCl

Fully convolutional layers

Computationally inefficient

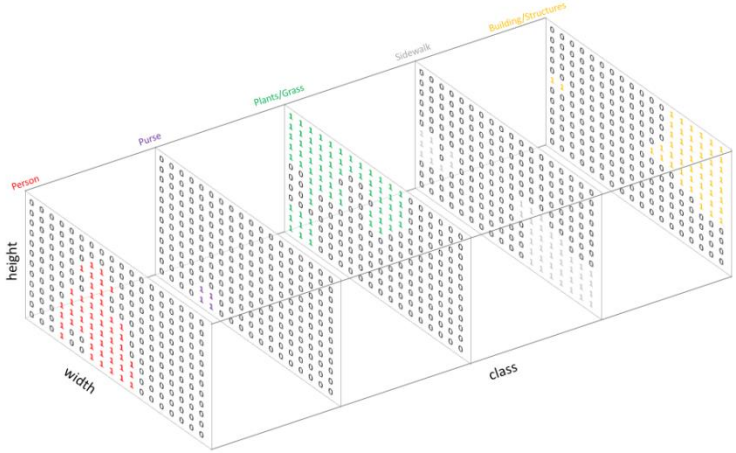

HOLISTIC PREDICTION VIA DEEP LEARNING

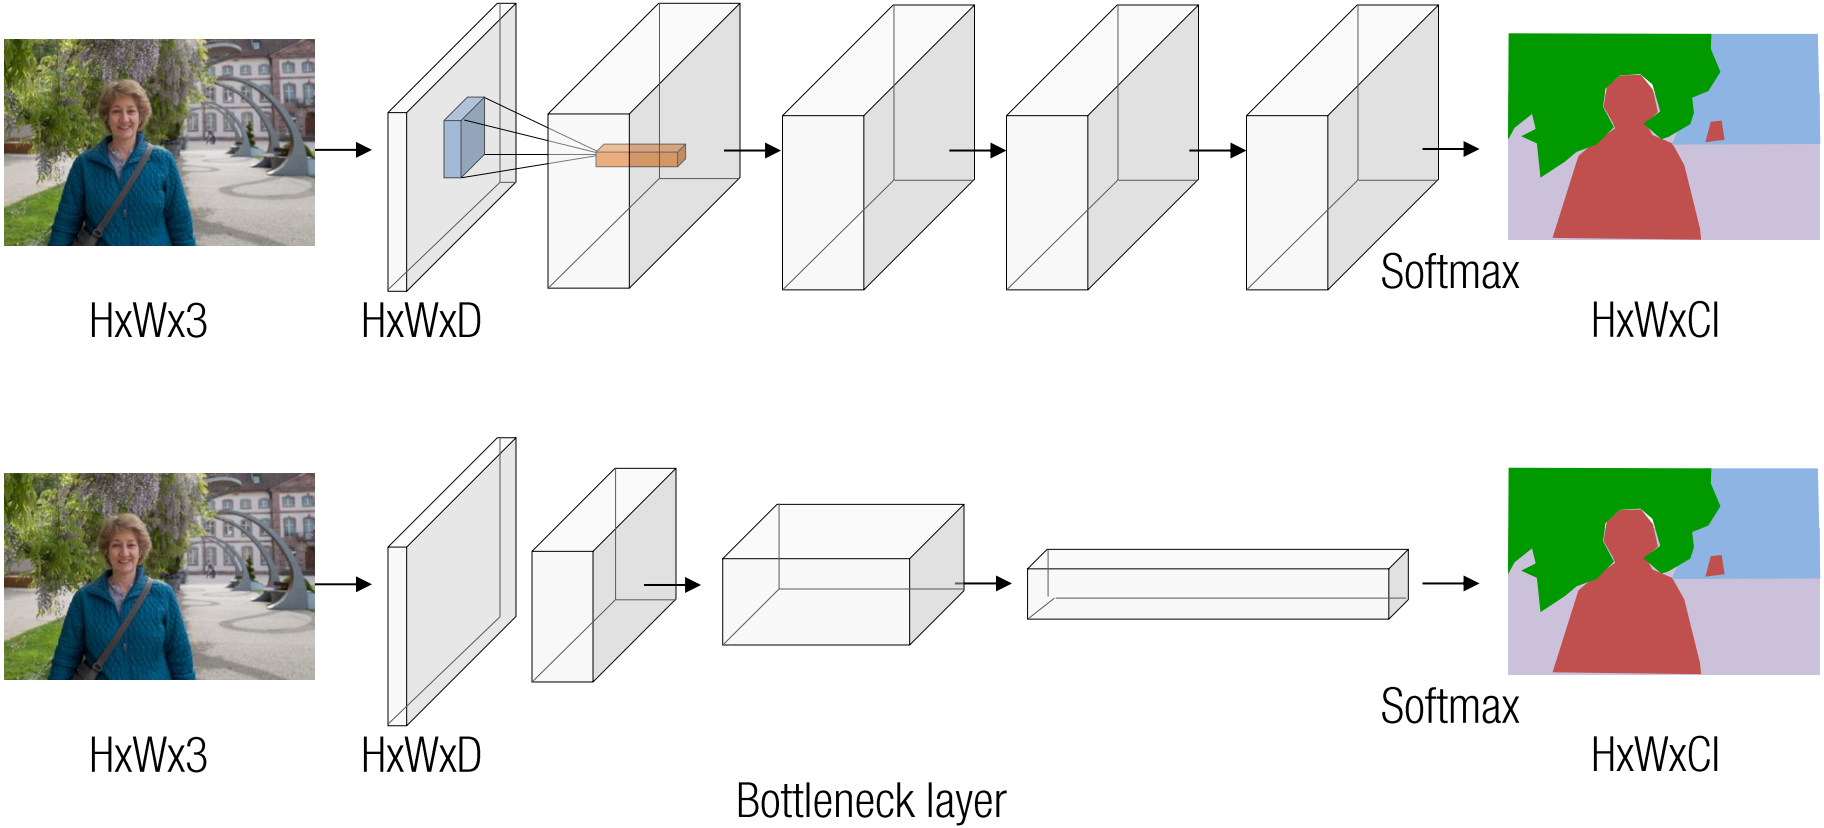

FULLY CONVOLUTIONAL NETWORK

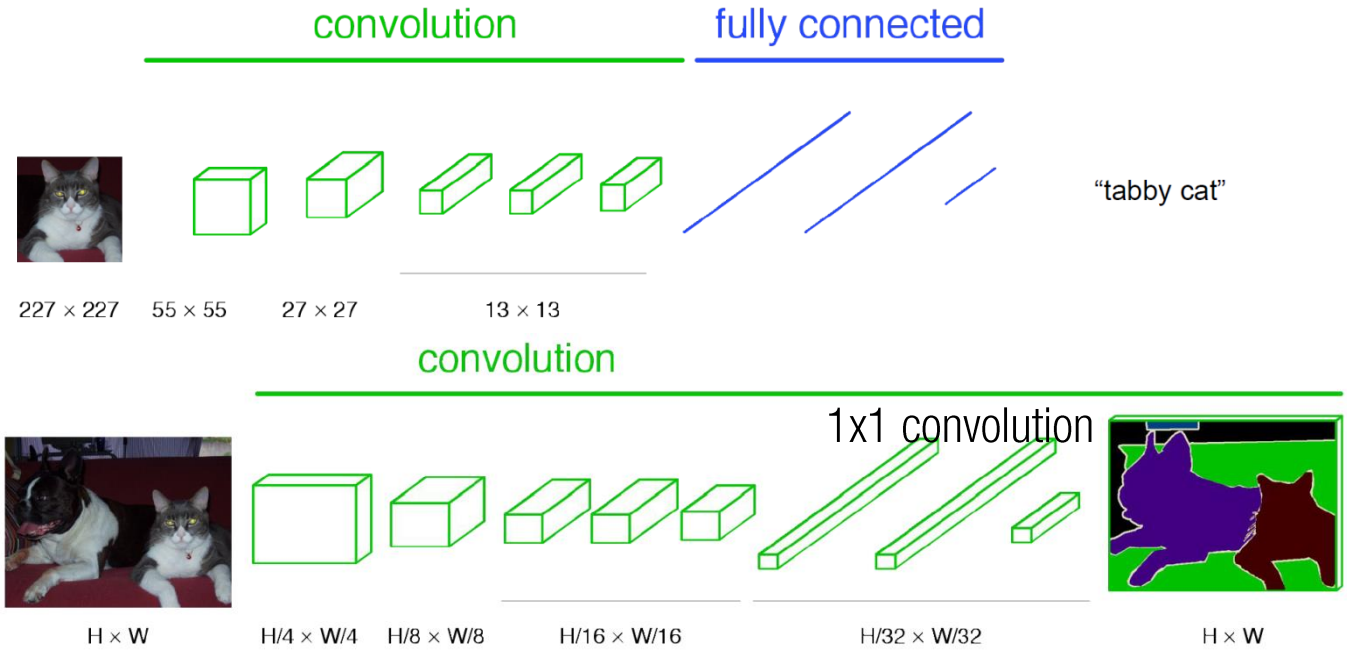

227x227

10x10

Upsampling via deconv.

Long, Shelhamer, and Darrell, "Fully Convolutional Networks for Semantic Segmentation", CVPR 2015

Fully Convolutional Network

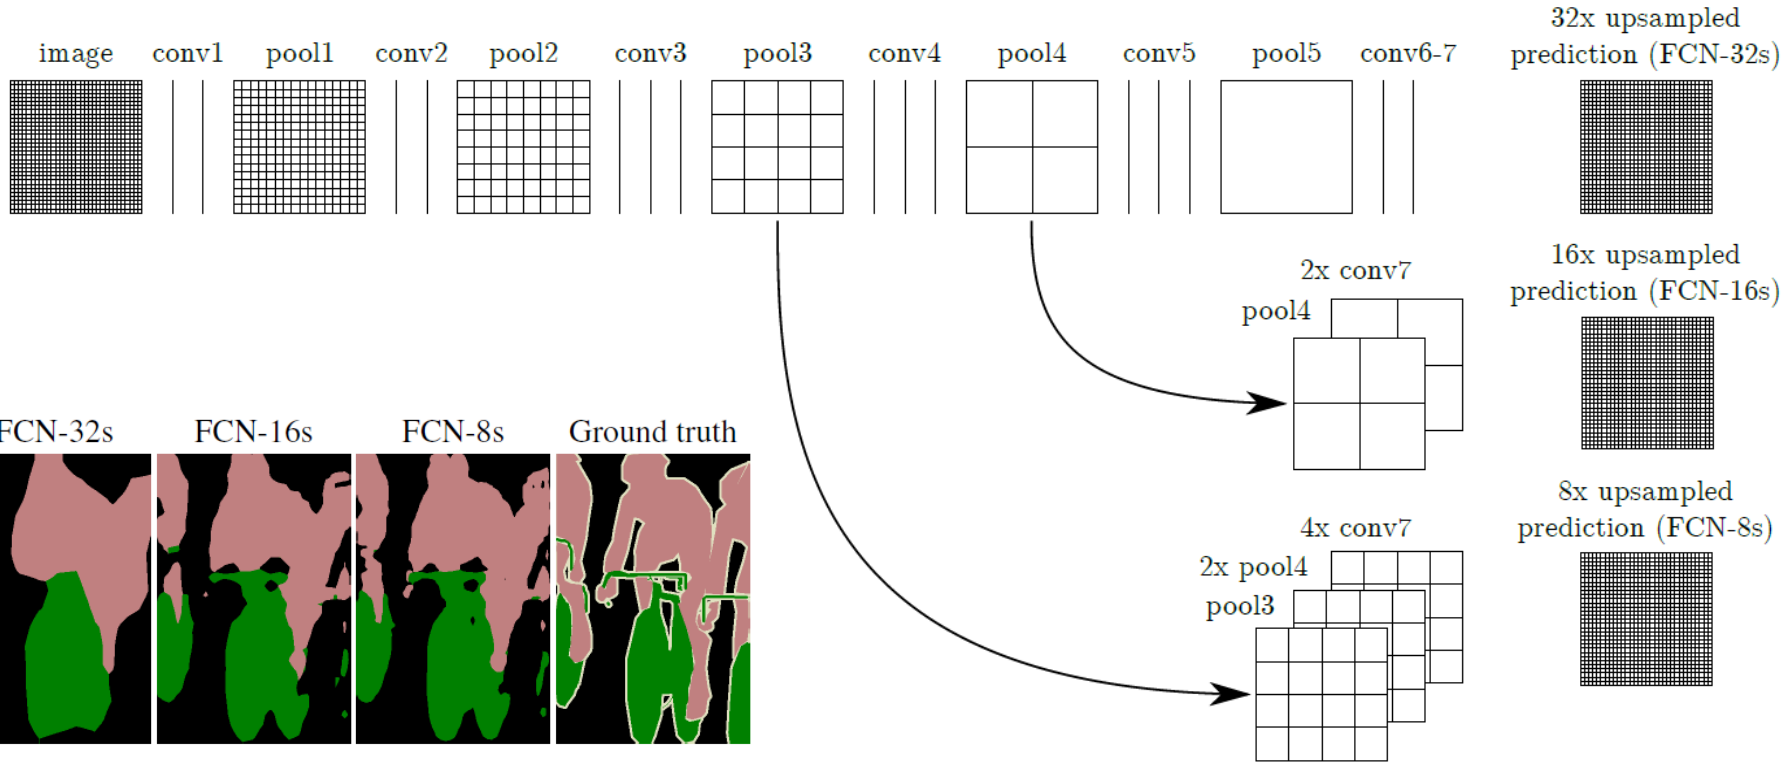

Long, Shelhamer, and Darrell, “Fully Convolutional Networks for Semantic Segmentation”, CVPR 2015

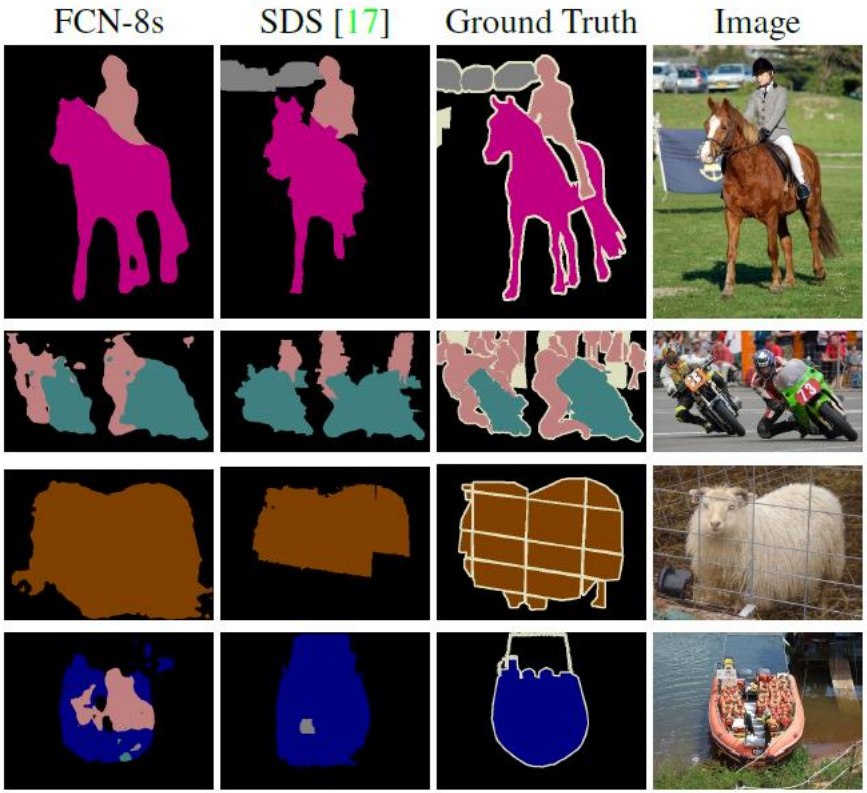
FCN-8s
SDS [17]
Ground Truth
Image

HOLISTIC PREDICTION VIA DEEP LEARNING

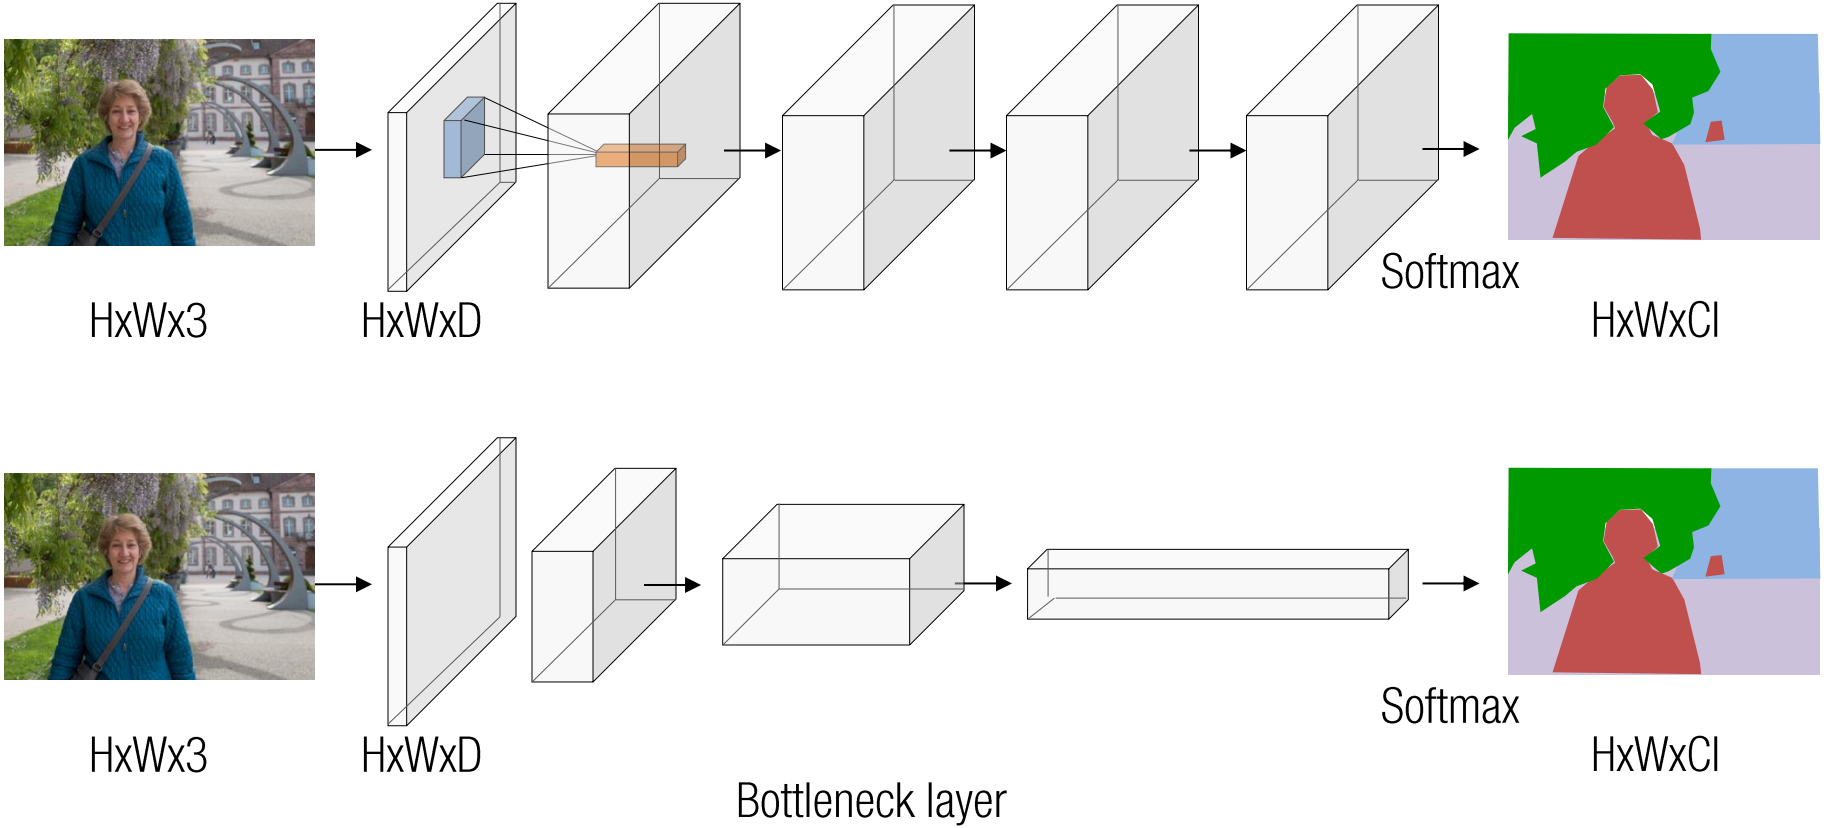

HOLISTIC PREDICTION VIA DEEP LEARNING

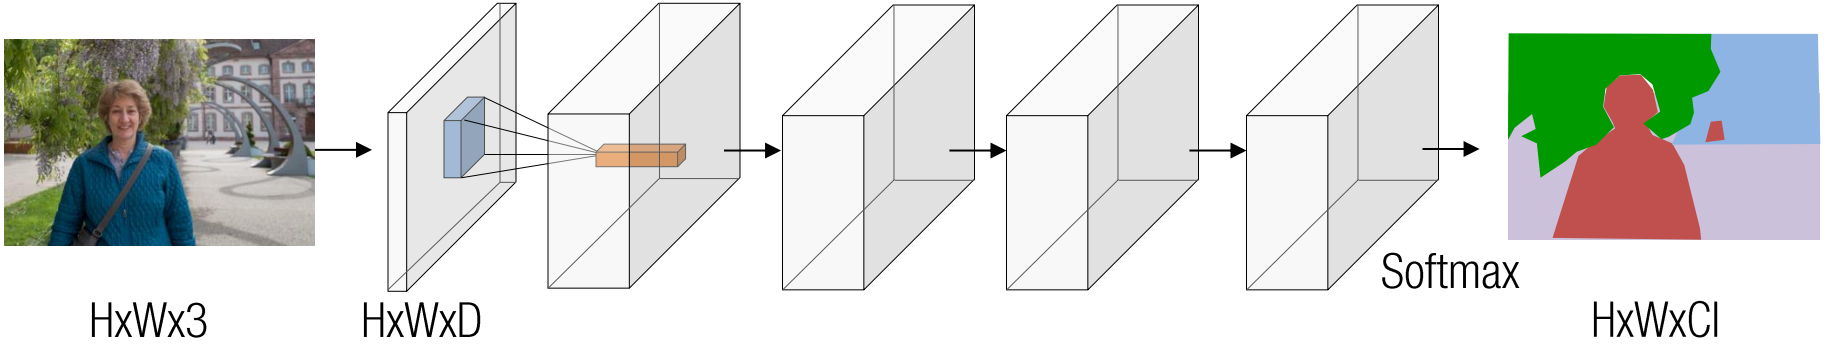

Upsampling?

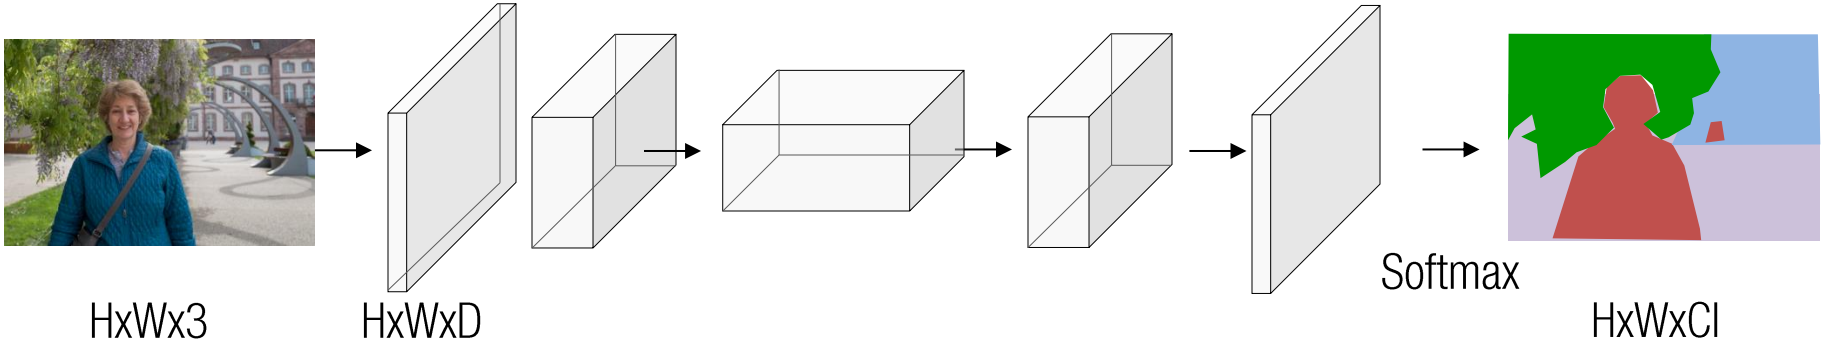

Bottleneck layer

REVISITED: SPATIAL POOLING (DOWN-SAMPLING)

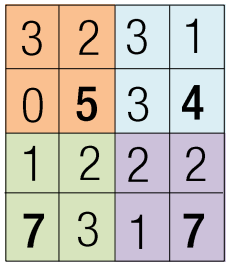

4×4

2×2

Max-pooling (window size 2x2, stride 2)

SPATIAL UNPOOLING (UP-SAMPLING)

Max-**un**pooling (window size 2x2, stride 2)

5	**4**
7	**7**

0	0	0	0
0	**5**	0	**4**
0	0	0	0
7	0	0	**7**

Learnable parameter?

Can we learn upsampling?

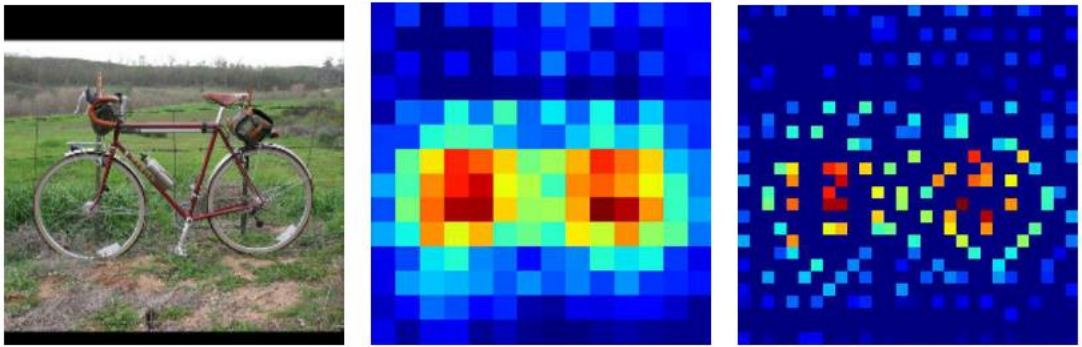

LEARNING UPSAMPLING

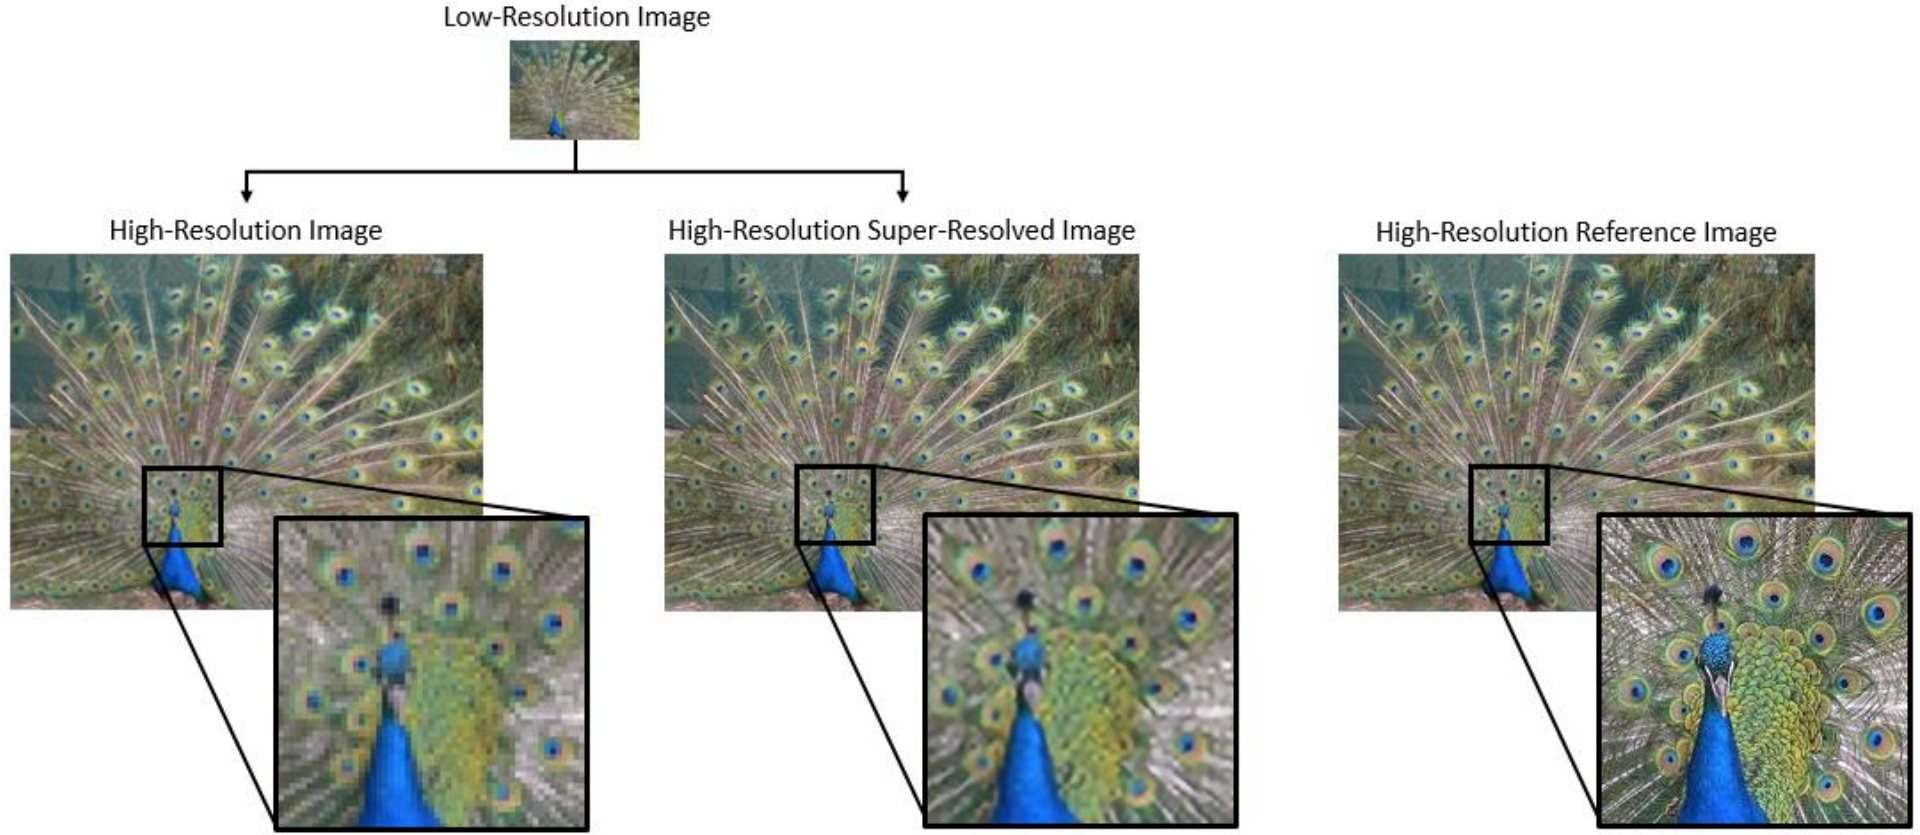

SPATIAL UNPOOLING (UP-SAMPLING)

Max-**un**pooling (window size 2x2, stride 2)

5	**4**
7	**7**

0	0	0	0
0	**5**	0	**4**
0	0	0	0
7	0	0	**7**

Learnable parameter?

Can we learn upsampling?

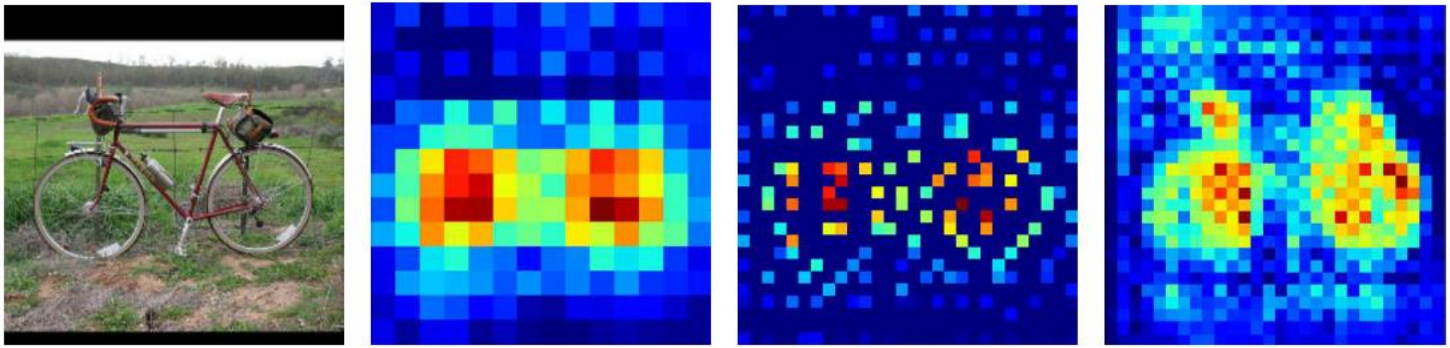

REVISITED: CONVOLUTION

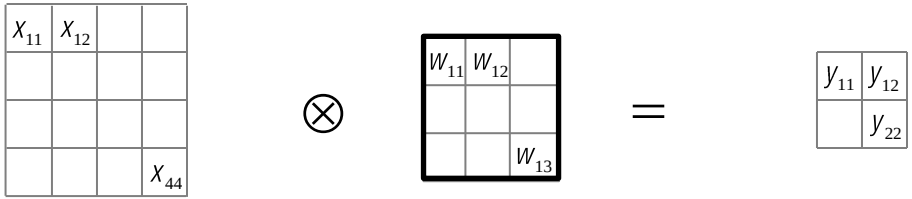

$$y_{11} = w_{11}x_{11} + w_{12}x_{12} + \cdots + w_{33}x_{33}$$

$$y_{12} = w_{11}x_{12} + w_{12}x_{13} + \cdots + w_{33}x_{34}$$

$$y_{21} = w_{11}x_{21} + w_{12}x_{21} + \cdots + w_{33}x_{43}$$

$$y_{22} = w_{11}x_{22} + w_{12}x_{23} + \cdots + w_{33}x_{44}$$

REVISITED: CONVOLUTION

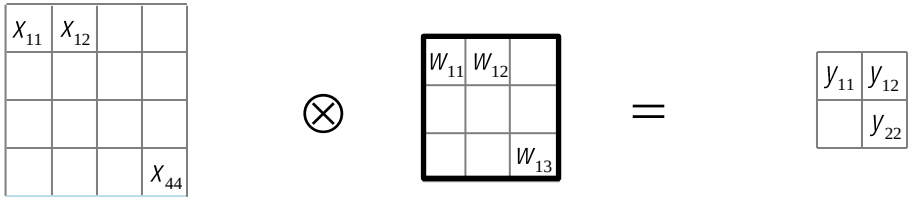

$$\begin{bmatrix} y_{11} \\ y_{12} \\ y_{21} \\ y_{22} \end{bmatrix} = \begin{bmatrix} w_{11} & w_{12} & w_{13} & 0 & w_{21} & w_{22} & w_{23} & 0 & w_{31} & w_{32} & w_{33} & 0 & 0 & 0 & 0 & 0 \\ 0 & w_{11} & w_{12} & w_{13} & 0 & w_{21} & w_{22} & w_{23} & 0 & w_{31} & w_{32} & w_{33} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & w_{11} & w_{12} & w_{13} & 0 & w_{21} & w_{22} & w_{23} & 0 & w_{31} & w_{32} & w_{33} & 0 \\ 0 & 0 & 0 & 0 & 0 & w_{11} & w_{12} & w_{13} & 0 & w_{21} & w_{22} & w_{23} & 0 & w_{31} & w_{32} & w_{33} \end{bmatrix} \begin{bmatrix} x_{11} \\ x_{12} \\ \vdots \\ x_{44} \end{bmatrix}$$

REVISITED: CONVOLUTION

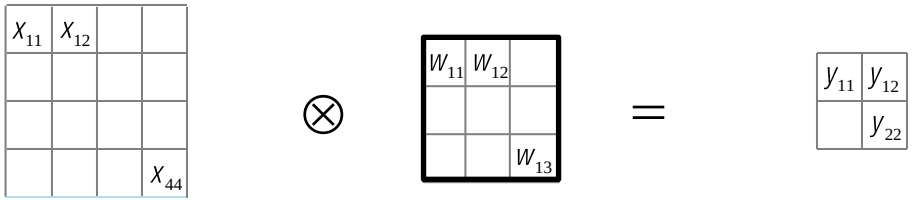

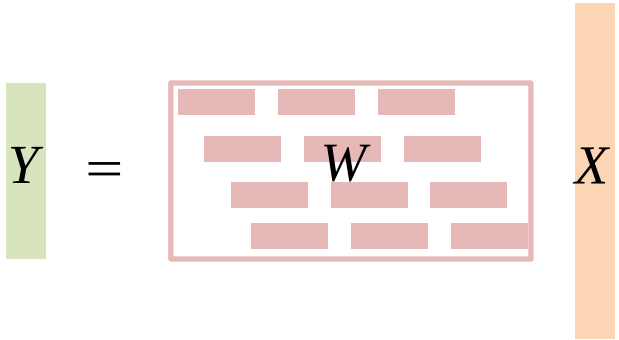

Non-zero entries of each row encode local connectivity.

INVERSE OF CONVOLUTION

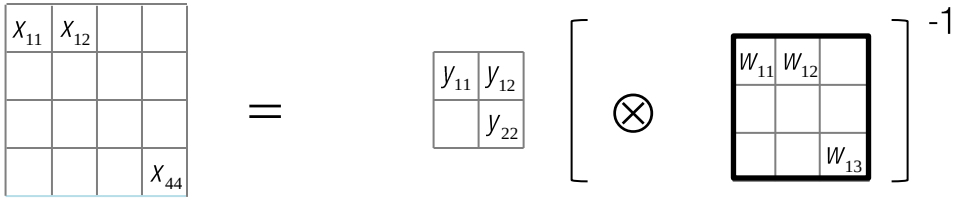

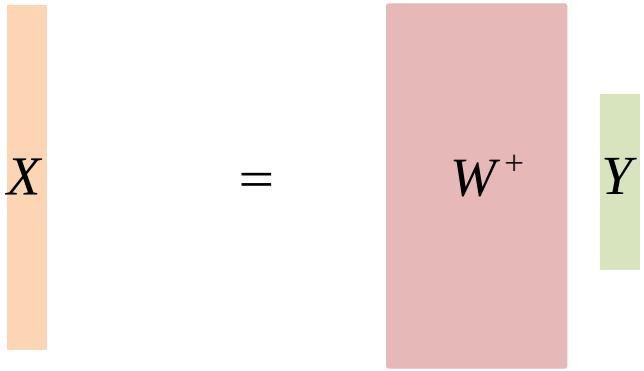

Inverse of W does not preserve local connectivity.

Upconvolution ~ Inverse of Convolution

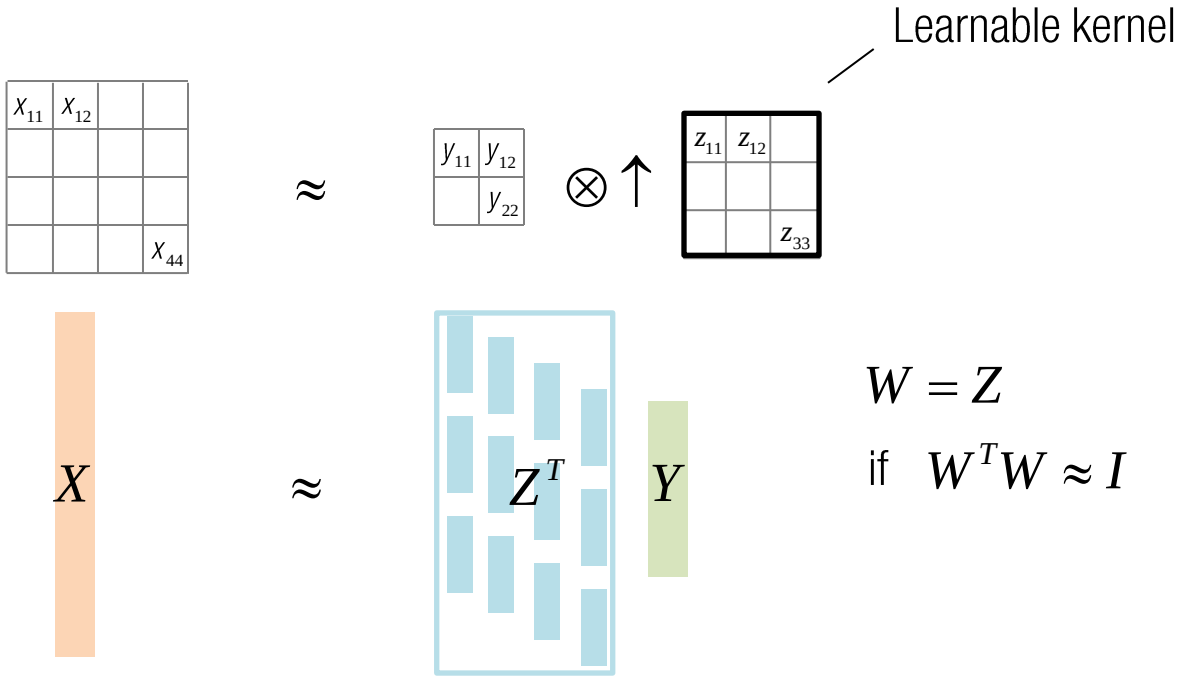

Non-zero entries of column encode local connectivity.

Upconvolution ~ Inverse of Convolution

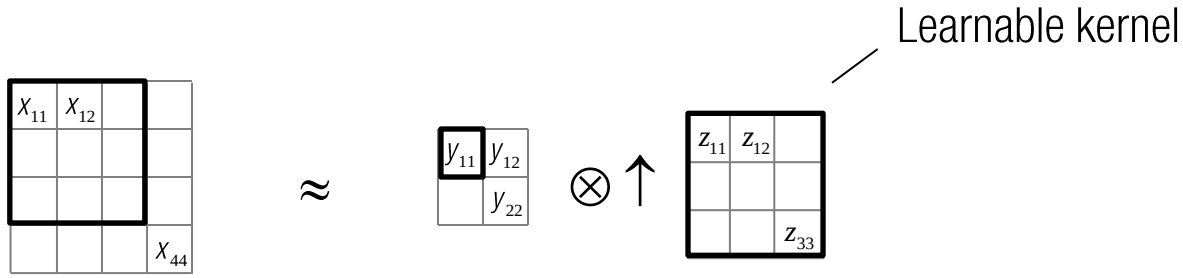

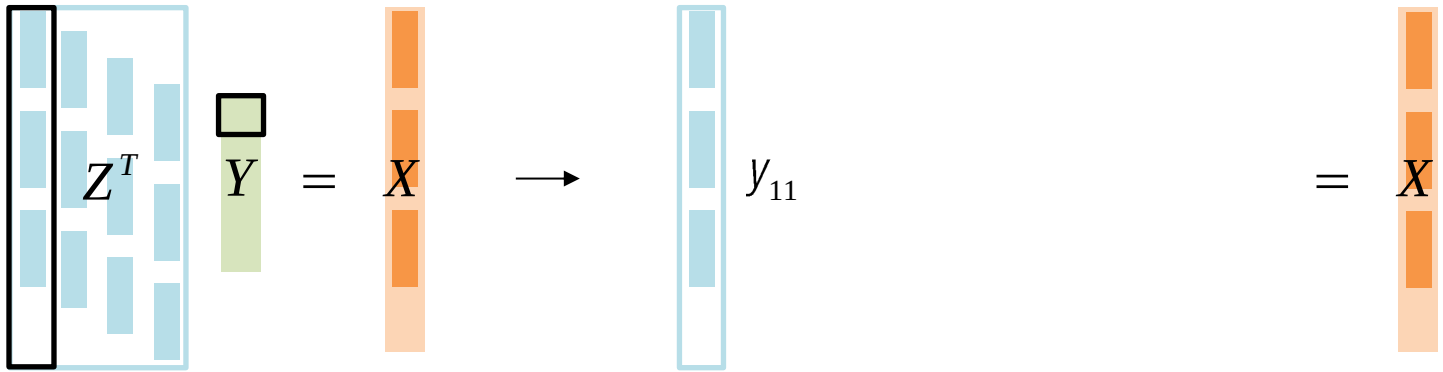

Non-zero entries of column encode local connectivity.

Transposed Convolution ~ Inverse of Convolution

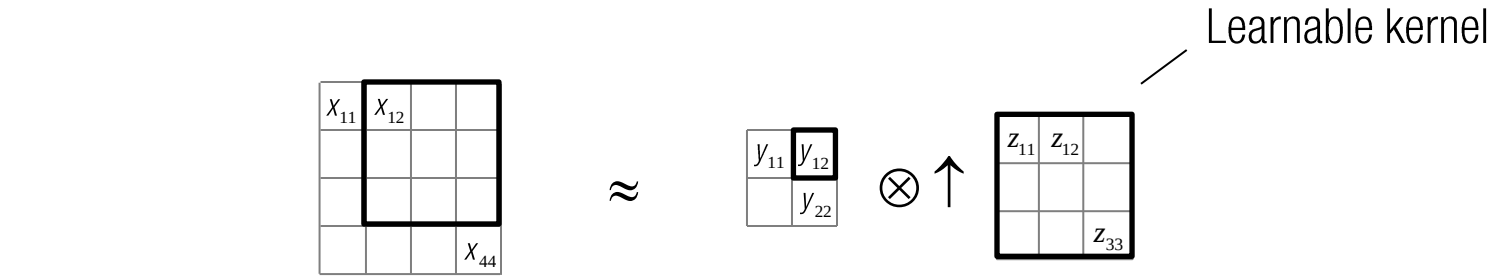

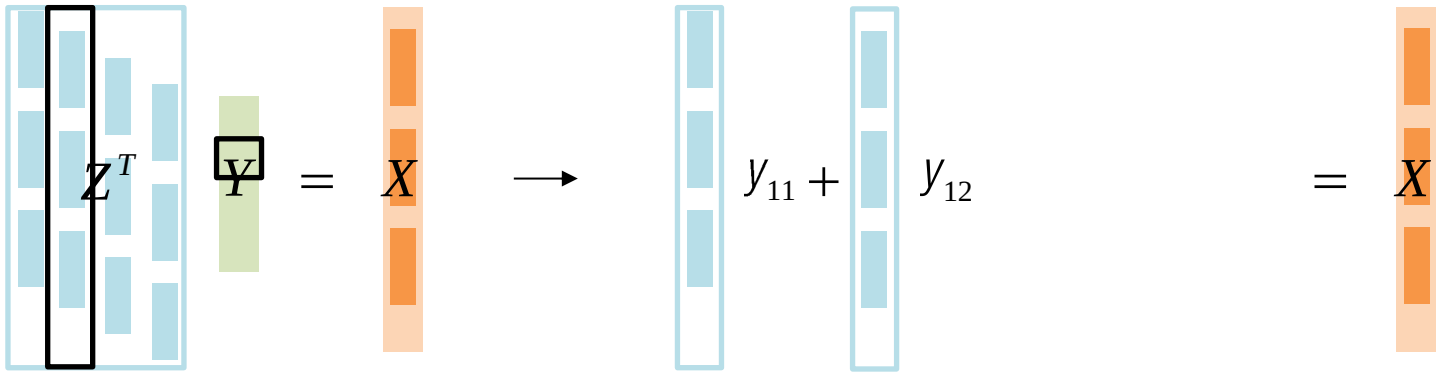

Non-zero entries of column encode local connectivity.

Transposed Convolution ~ Inverse of Convolution

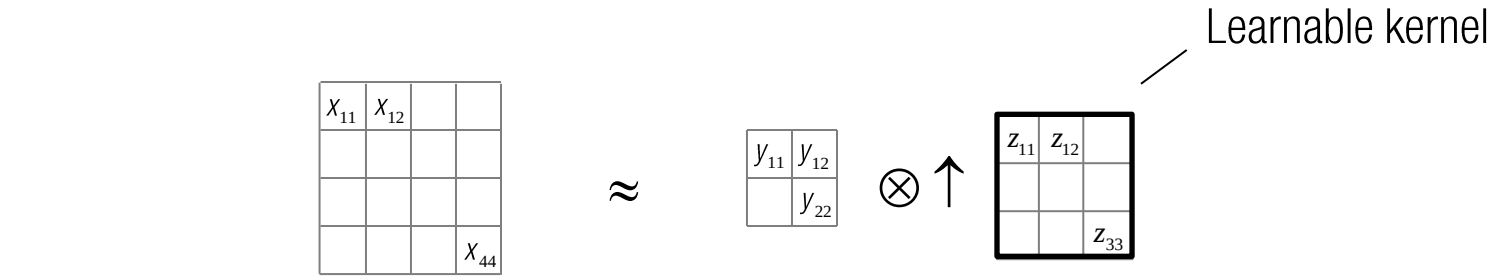

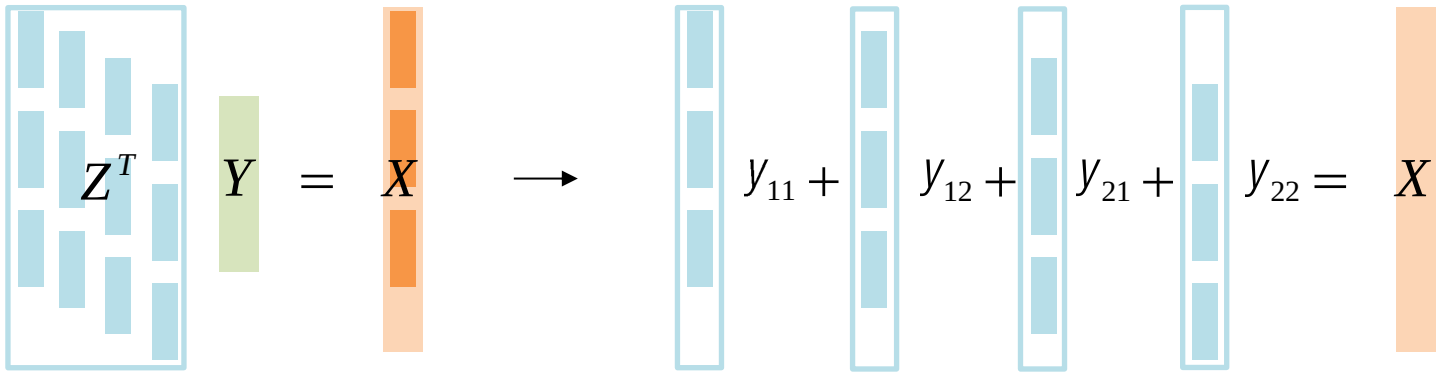

Non-zero entries of column encode local connectivity.

Transposed Convolution as Upconvolution

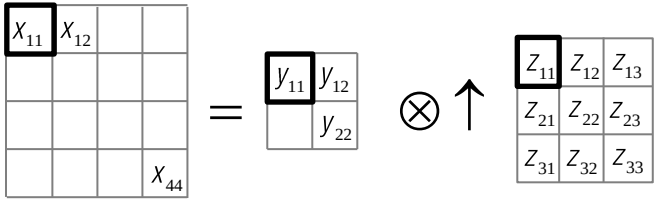

Upconvolution

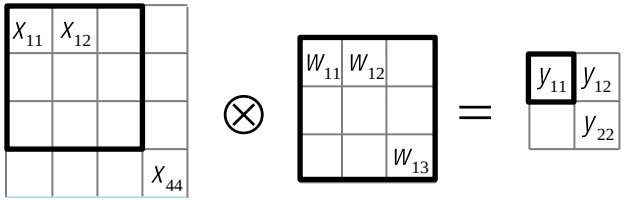

Cf) convolution

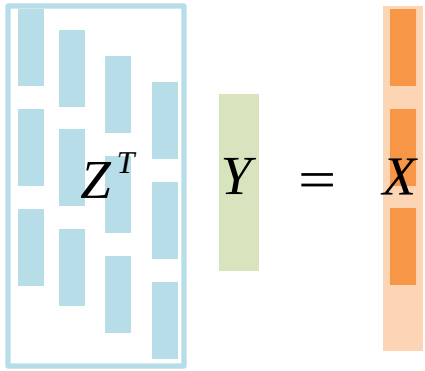

$$x_{11} = z_{11} y_{11}$$

x_{11} is dependent only on y_{11}.

x_{11} contributes only to y_{11}.

Transposed Convolution as Upconvolution

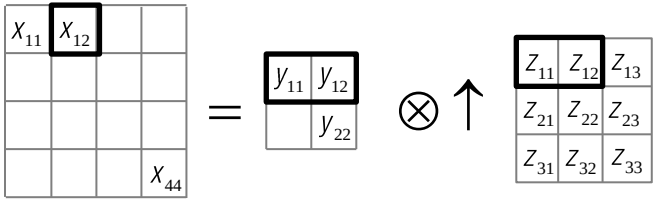

Upconvolution

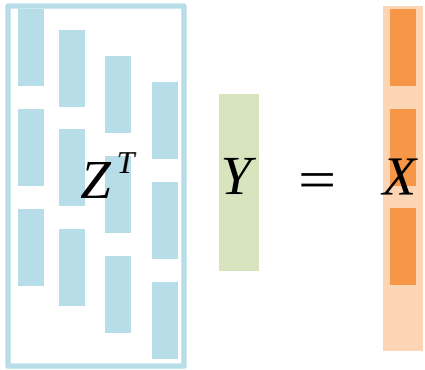

$$x_{11} = z_{11}y_{11}$$
$$x_{12} = z_{12}y_{11} + z_{11}y_{12}$$

Transposed Convolution as Upconvolution

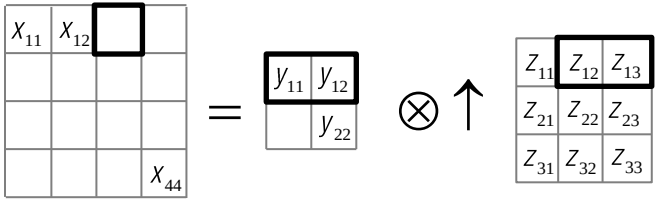

Upconvolution

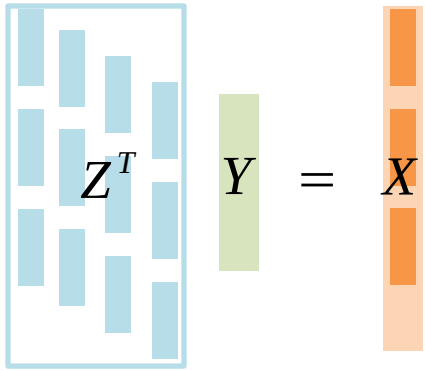

$$x_{11} = z_{11}y_{11}$$
$$x_{12} = z_{12}y_{11} + z_{11}y_{12}$$
$$x_{13} = z_{13}y_{11} + z_{12}y_{12}$$

Transposed Convolution as Upconvolution

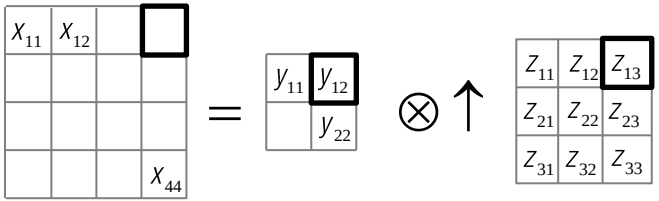

Upconvolution

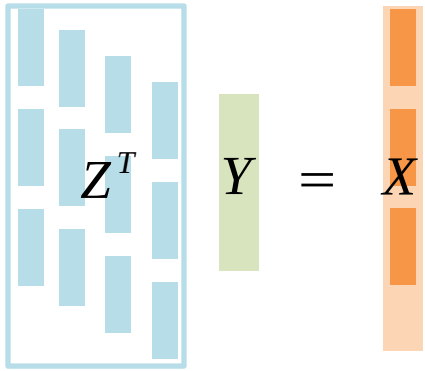

$$x_{11} = z_{11}y_{11}$$
$$x_{12} = z_{12}y_{11} + z_{11}y_{12}$$
$$x_{13} = z_{13}y_{11} + z_{12}y_{12}$$
$$x_{14} = z_{13}y_{12}$$

Transposed Convolution as Upconvolution

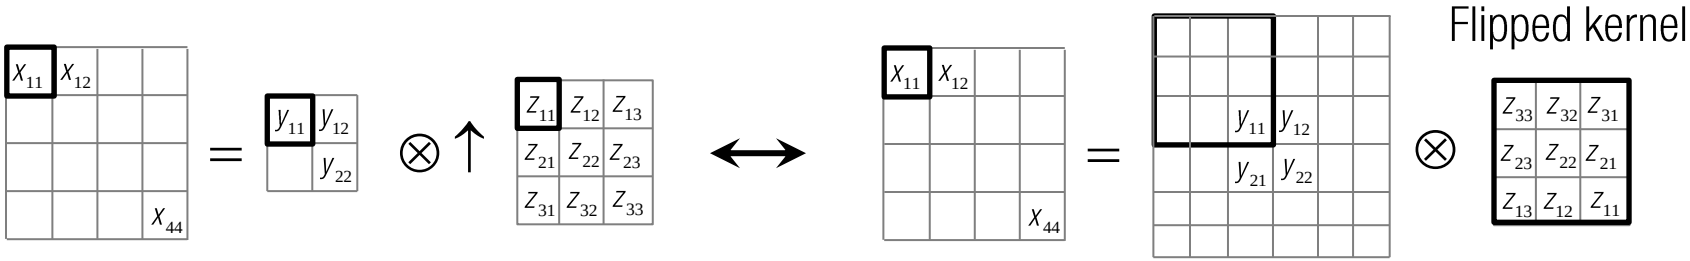

$$Z^T Y = X$$

$$x_{11} = z_{11} y_{11}$$

Transposed Convolution as Upconvolution

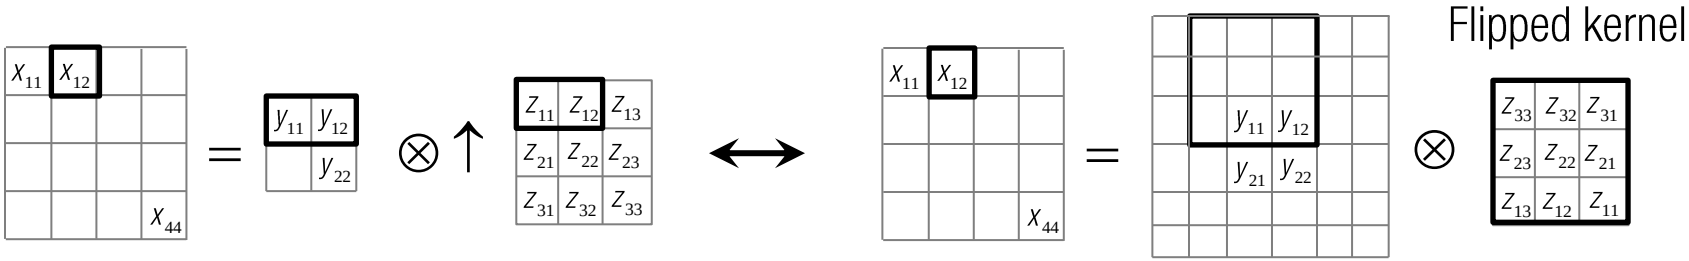

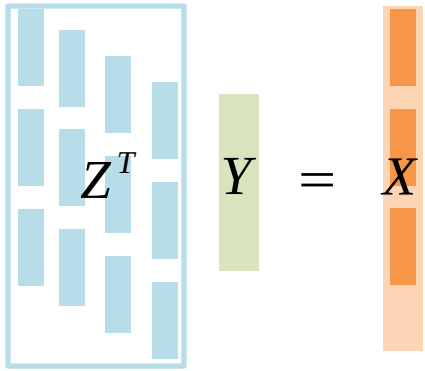

$$x_{11} = z_{11} y_{11}$$

$$x_{12} = z_{12} y_{11} + z_{11} y_{12}$$

Transposed Convolution as Upconvolution

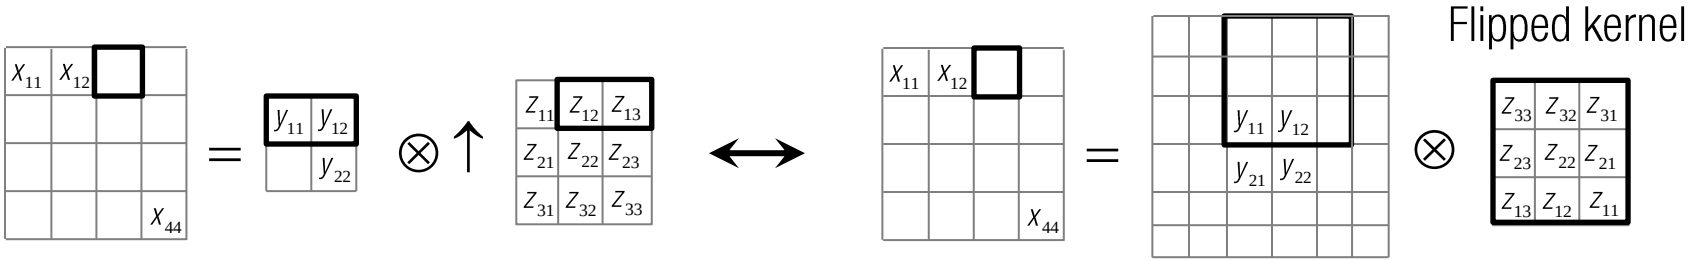

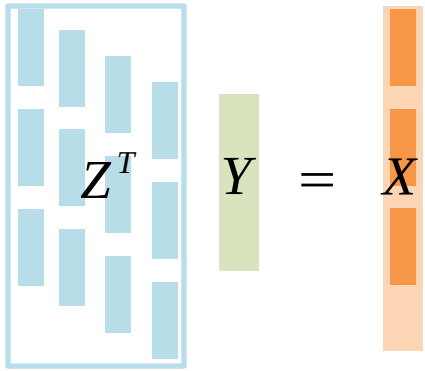

$$X_{11} = Z_{11} y_{11}$$
$$X_{12} = Z_{12} y_{11} + Z_{11} y_{12}$$
$$X_{13} = Z_{13} y_{11} + Z_{12} y_{12}$$

Transposed Convolution as Upconvolution

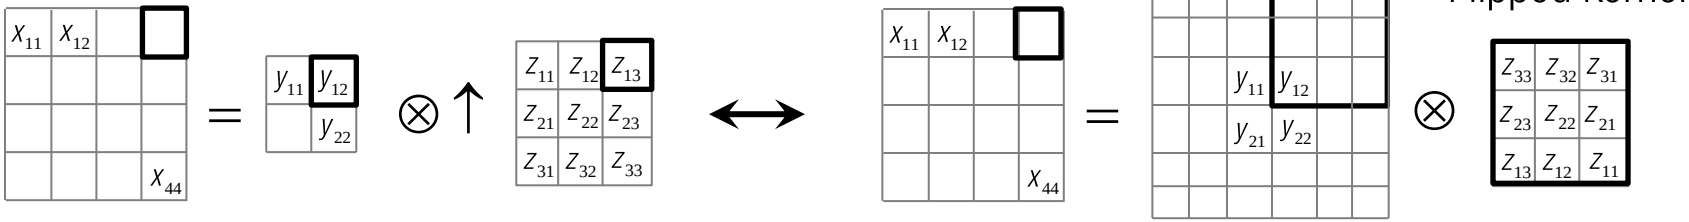

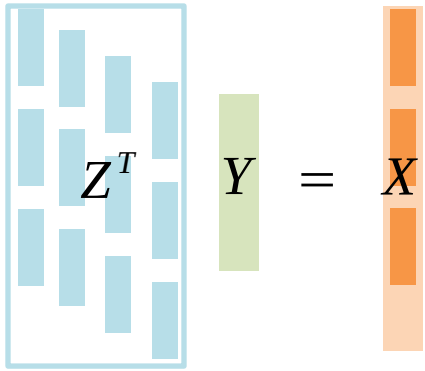

$$X_{11} = Z_{11} y_{11}$$
$$X_{12} = Z_{12} y_{11} + Z_{11} y_{12}$$
$$X_{13} = Z_{13} y_{11} + Z_{12} y_{12}$$
$$X_{14} = Z_{13} y_{12}$$

Transposed Convolution as Upconvolution

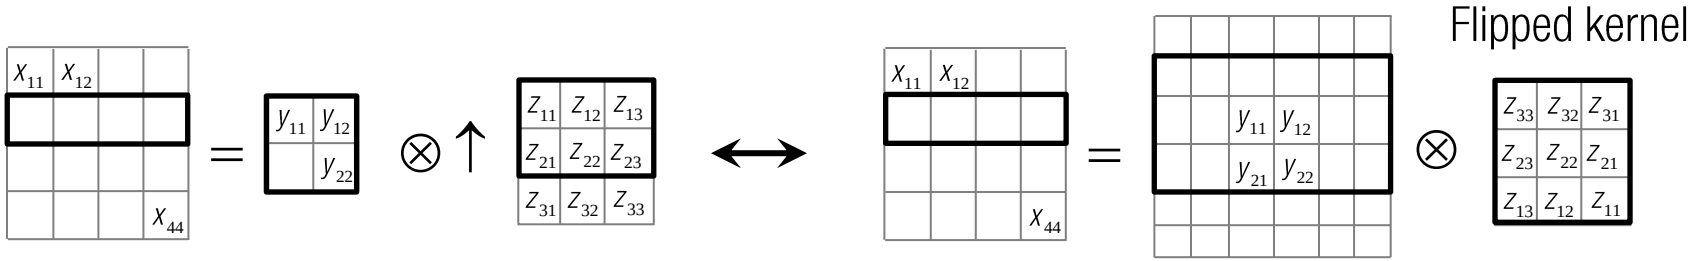

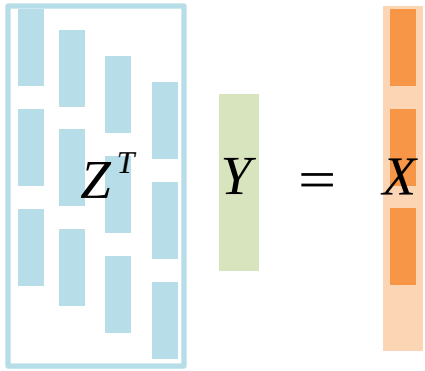

$$X_{11} = Z_{11}Y_{11}$$
$$X_{12} = Z_{12}Y_{11} + Z_{11}Y_{12}$$
$$X_{13} = Z_{13}Y_{11} + Z_{12}Y_{12}$$
$$X_{14} = Z_{13}Y_{12}$$

$$X_{21} = Z_{21}Y_{11} + Z_{11}Y_{21}$$
$$X_{22} = Z_{22}Y_{11} + Z_{21}Y_{12} + Z_{12}Y_{21} + Z_{11}Y_{22}$$
$$X_{23} = Z_{23}Y_{11} + Z_{22}Y_{12} + Z_{13}Y_{21} + Z_{12}Y_{22}$$
$$X_{24} = Z_{23}Y_{12} + Z_{13}Y_{22}$$

Transposed Convolution as Upconvolution

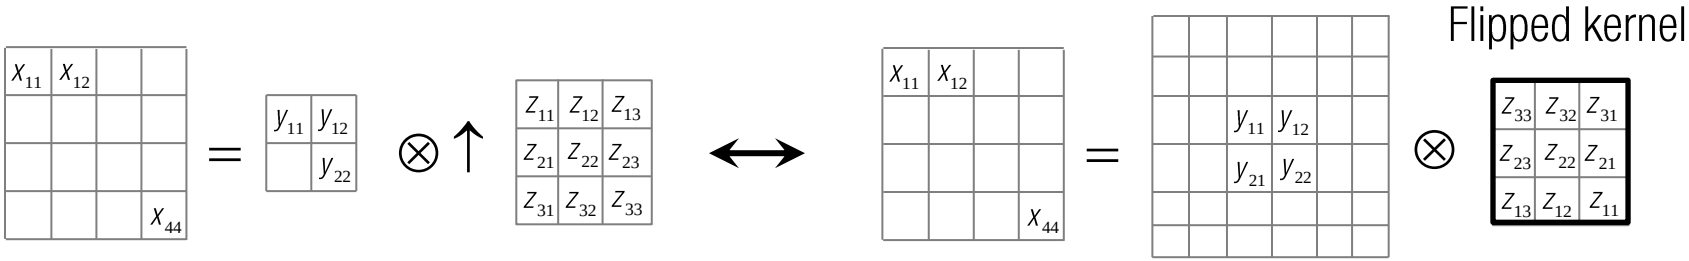

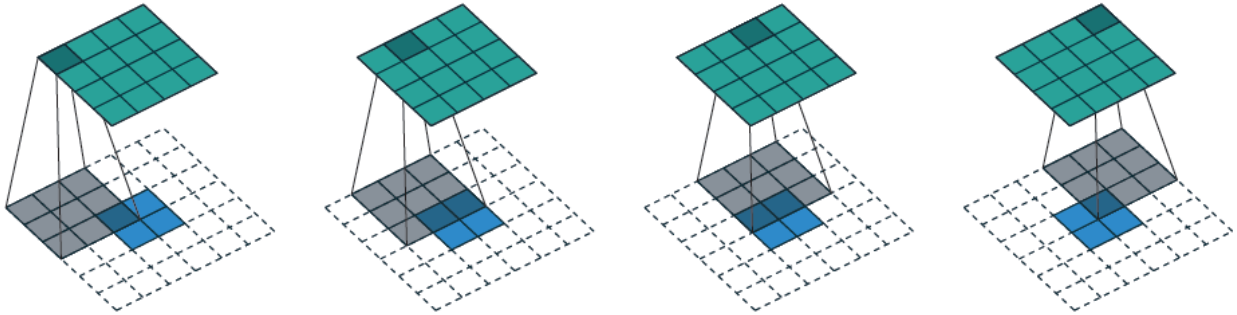

Duality: Convolution vs. Upconvolution

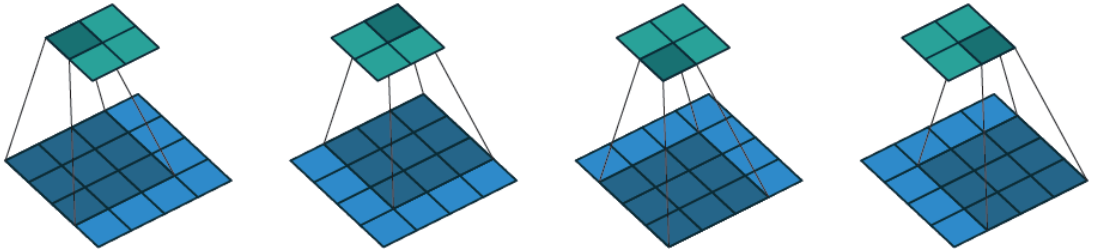

Convolution (3x3 kernel, no zero padding, 1 stride)

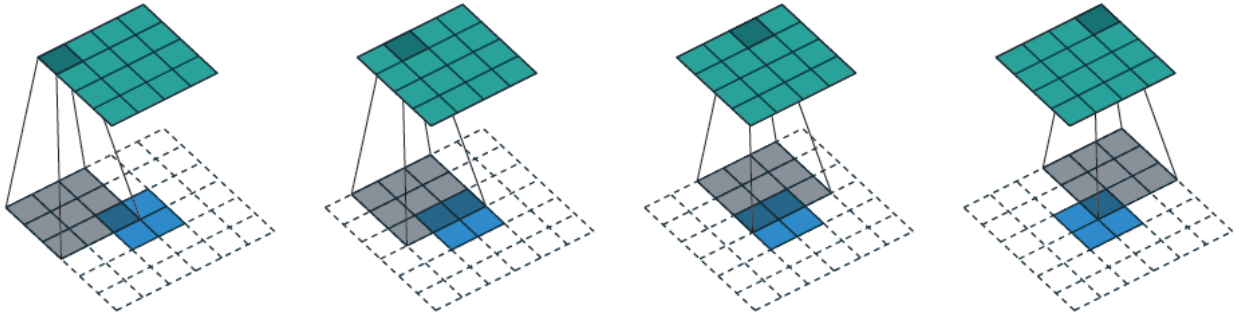

Upconvolution (3x3 kernel, 2x2 zero padding, 1 stride)

Duality: Convolution vs. Upconvolution

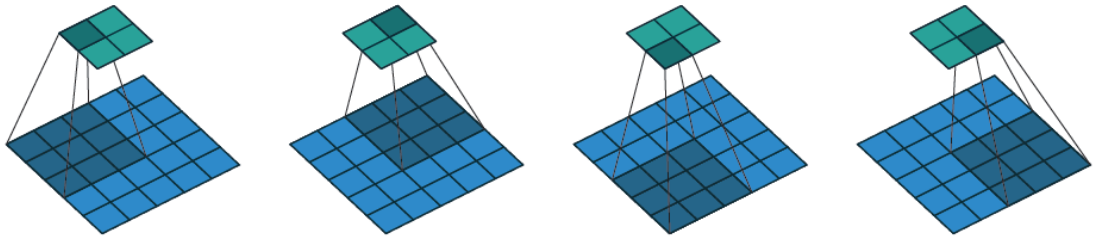

Convolution (3x3 kernel, no zero padding, 2 stride)

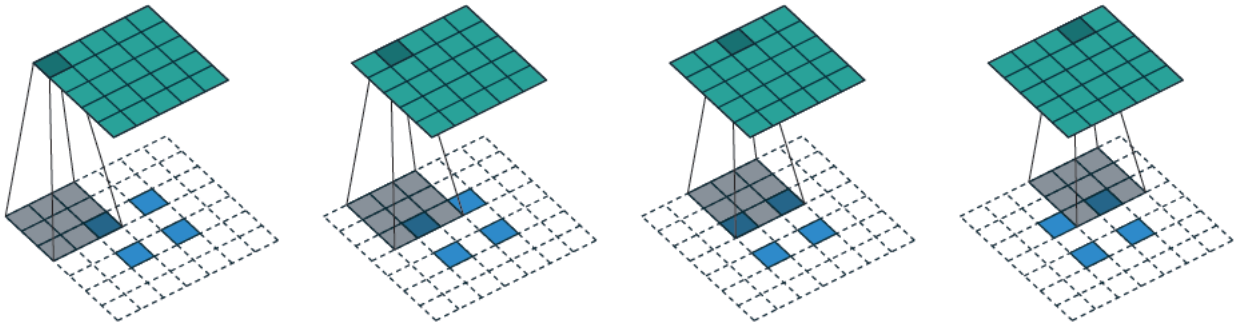

Upconvolution (3x3 kernel, 2x2 zero padding, 1 zero between input, 2 stride)

DEEPER UPCONVOLUTIONAL LAYERS

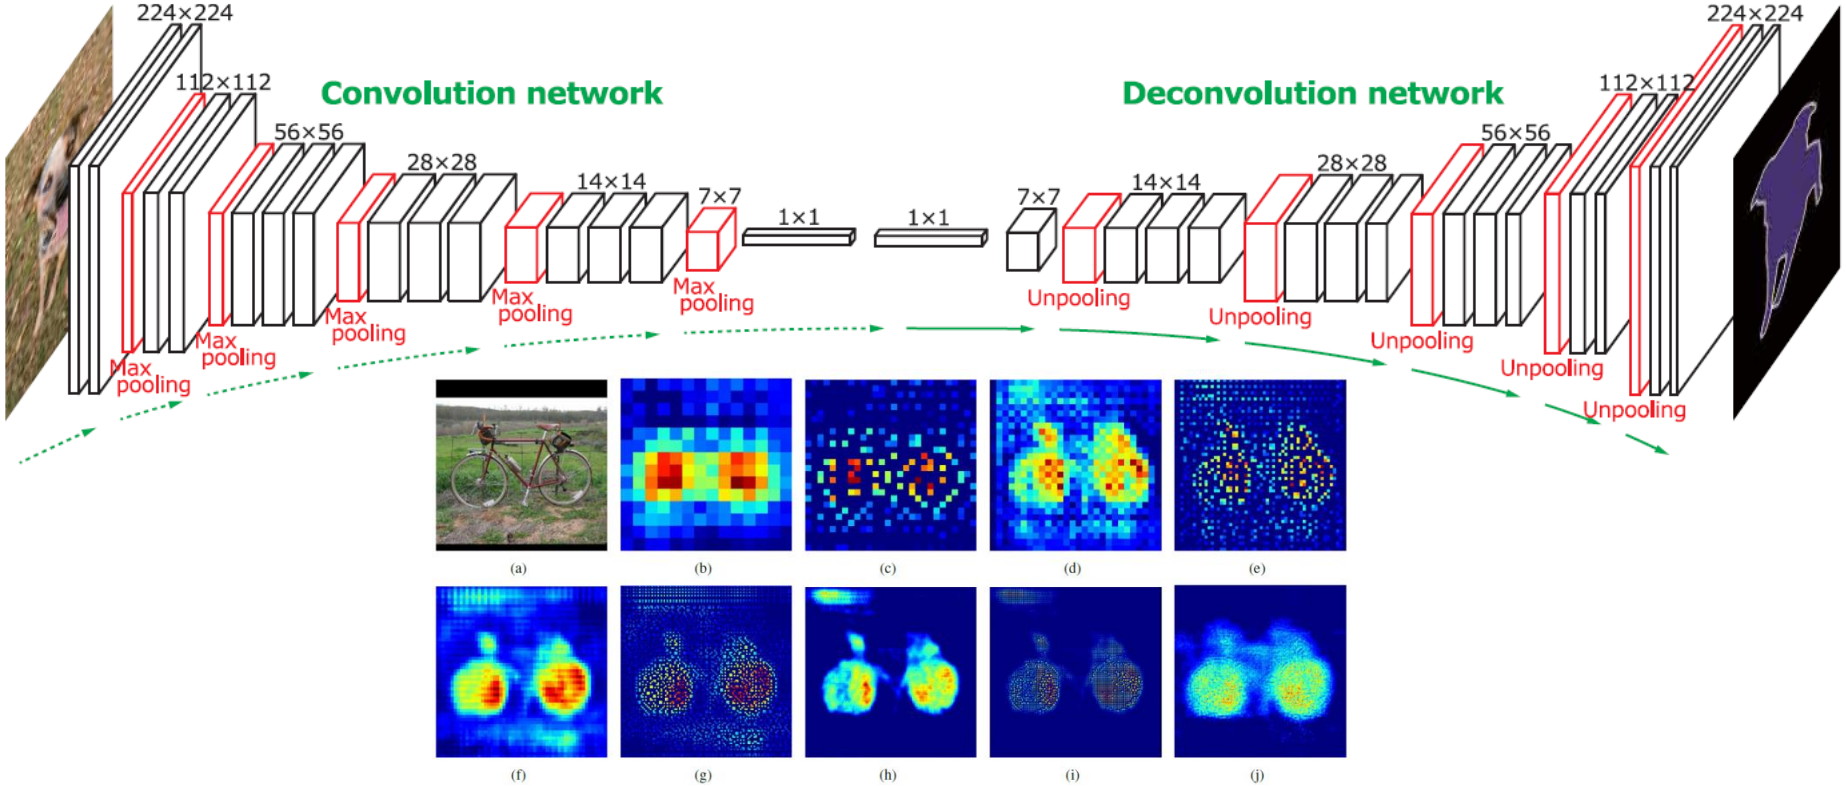

Noh, Hong, and Han, “Learning Deconvolutional Network for Semantic Segmentation”, ICCV 2015

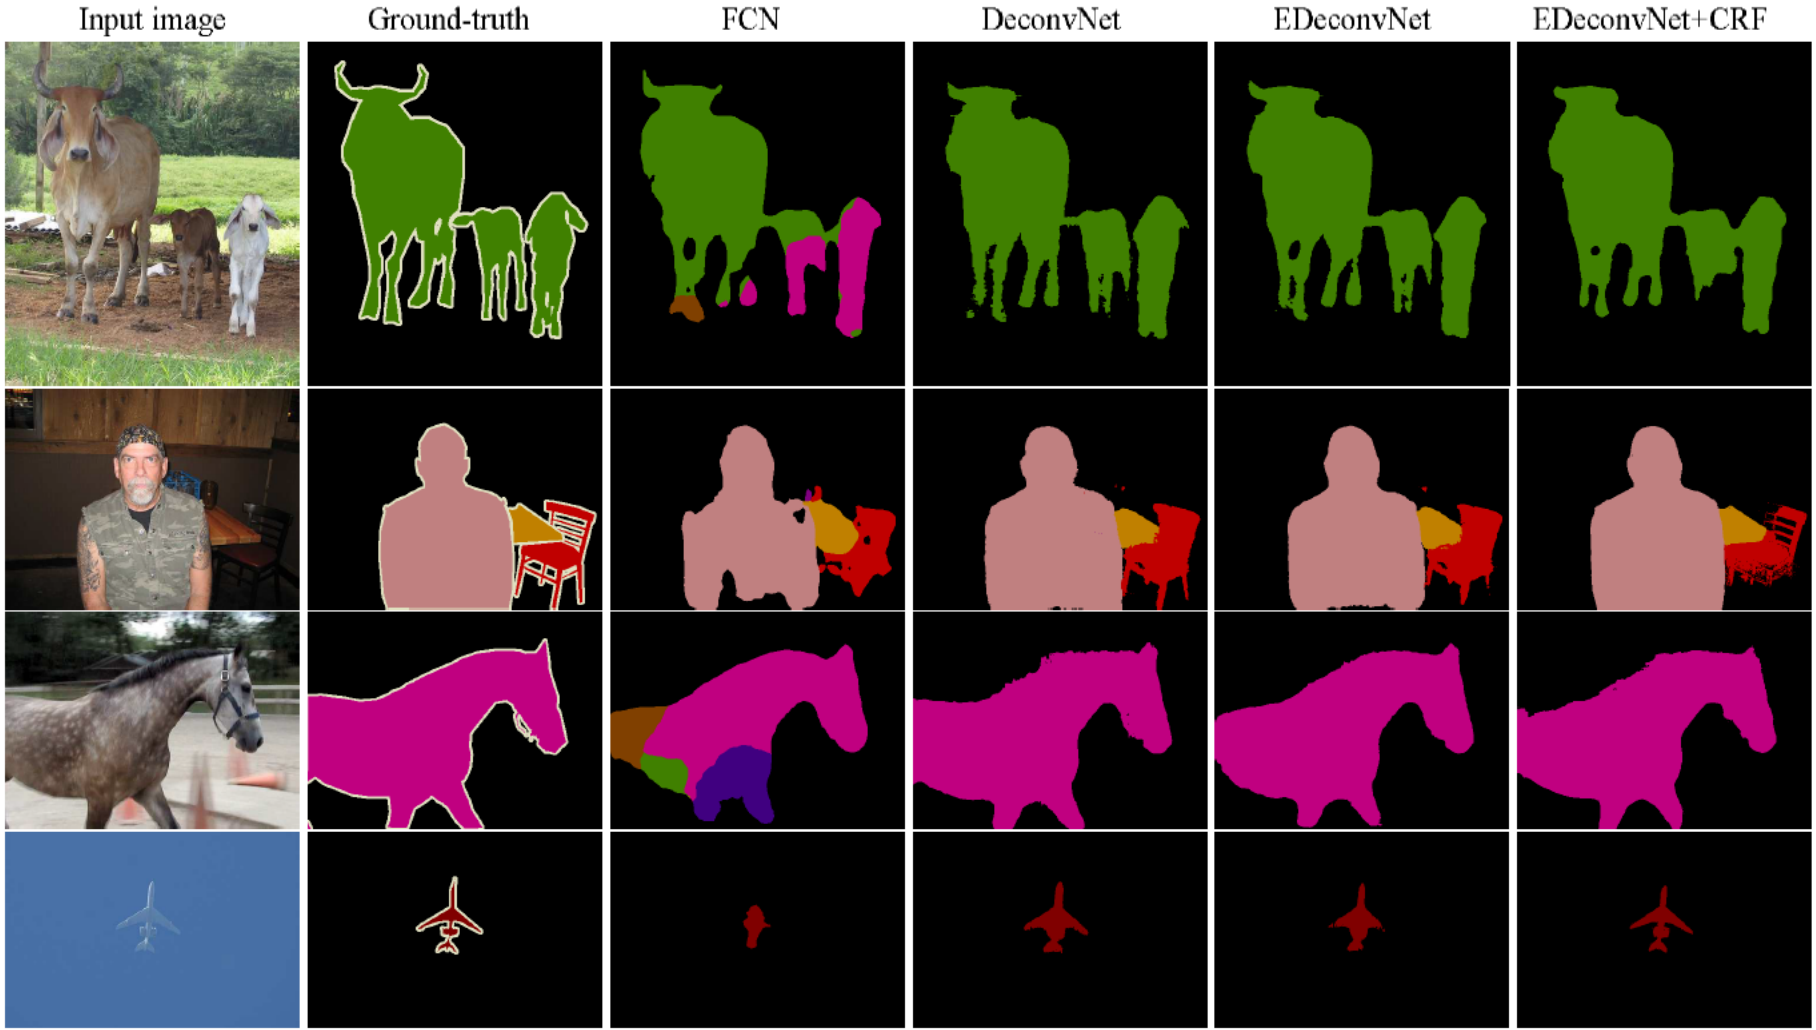
Input image
Ground-truth
FCN
DeconvNet
EDeconvNet
EDeconvNet+CRF

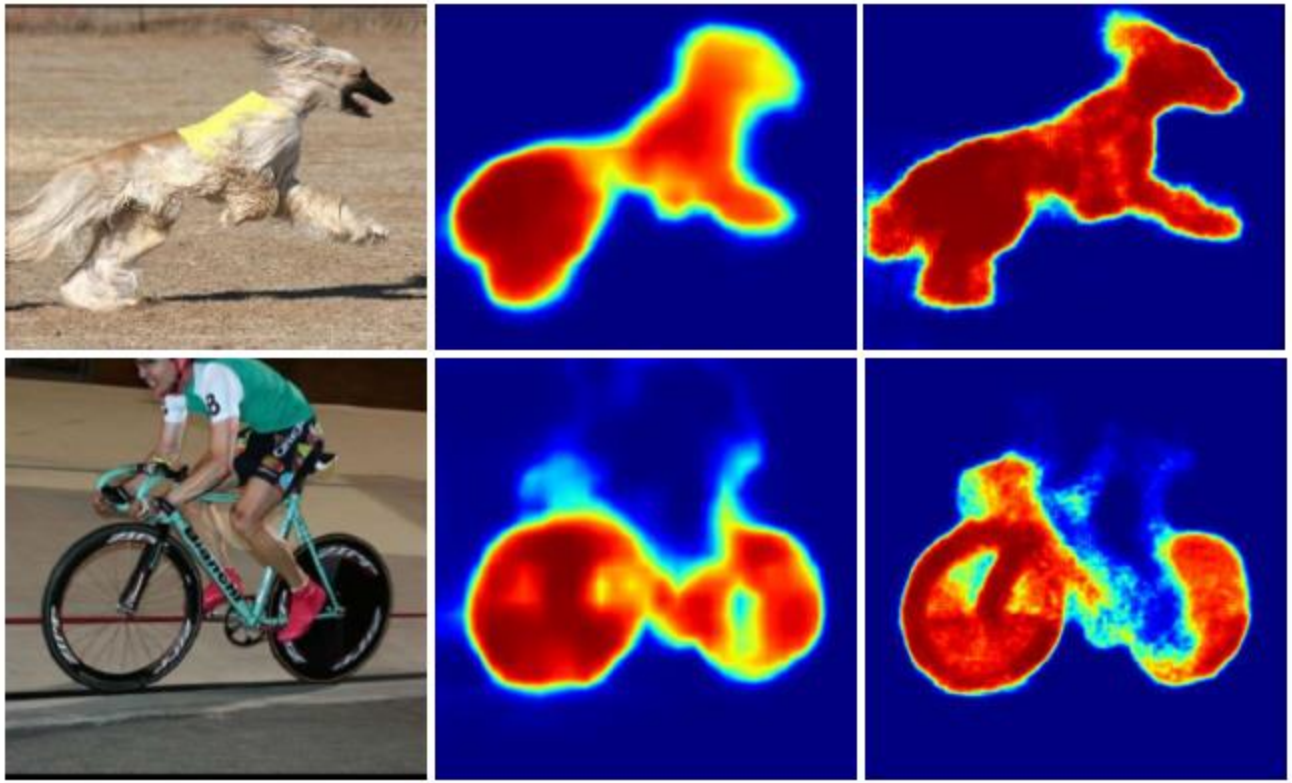

(a) Input image (b) FCN-8s (c) Ours

Receptive Field vs. Resolution

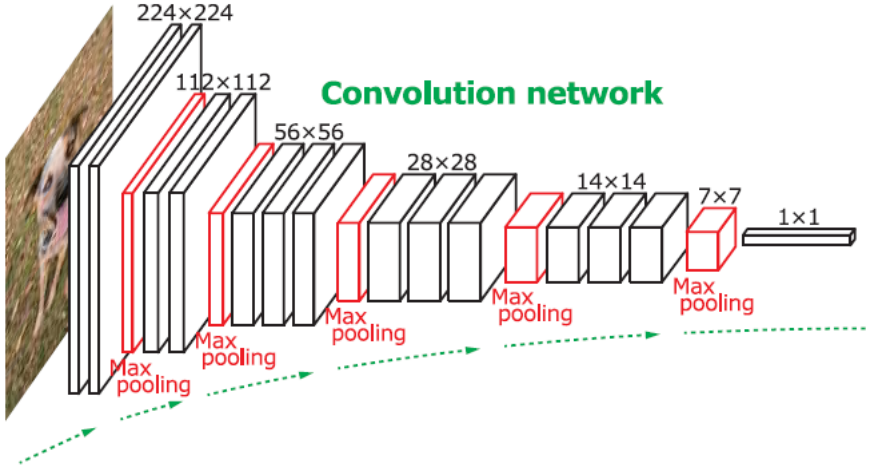

Desired properties:

- Larger receptive field (bigger spatial context) ⟶ reducing resolution
- Higher output resolution ⟶ reducing receptive field

DeepLab: Dilated Convolution

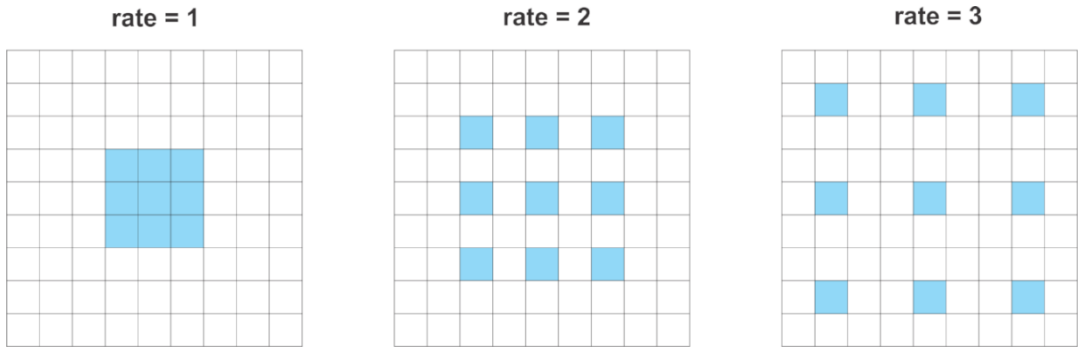

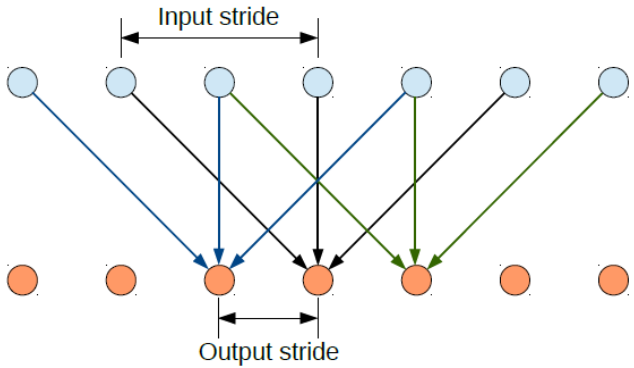

Enlarging receptive field without increasing reducing resolution

Chen, Papandreou, Kokkinos, Murphy, and Yuille, “Semantic Image Segmentation with Deep Convolutional Nets and Fully Connected CRFs”, ICLR 2015

DEEPLAB: DILATED CONVOLUTION

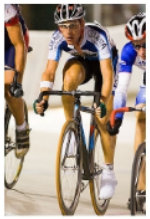

conv + pool

conv + pool

conv + pool

conv + pool

1x1x4096

FC

FC

FC

FC

1x1xN

final classification

224x224x3

28x28x1024

28x28x4096

coarse segmentation **28x28**xN

224x224xN

final segmentation

→ dilated convolutions, stride 1, no pooling
preserve spatial dims

→ 1x1 convolutions, stride 1
preserve spatial dims

bilinear upsampling (8x)
to original resolution

DeepLab: CRF Post-processing

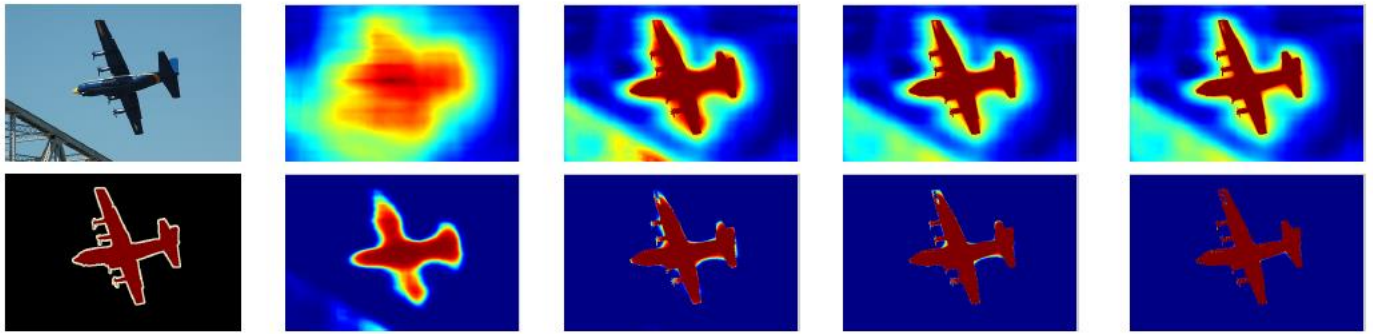

Chen, Papandreou, Kokkinos, Murphy, and Yuille, “Semantic Image Segmentation with Deep Convolutional Nets and Fully Connected CRFs”, ICLR 2015

DeepLab

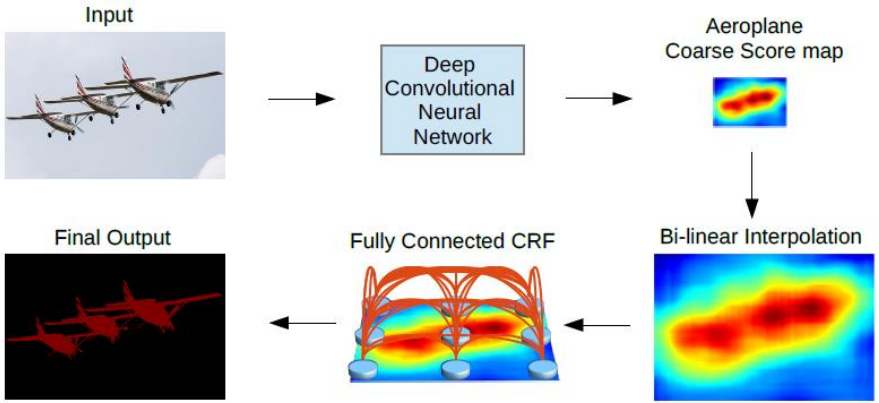

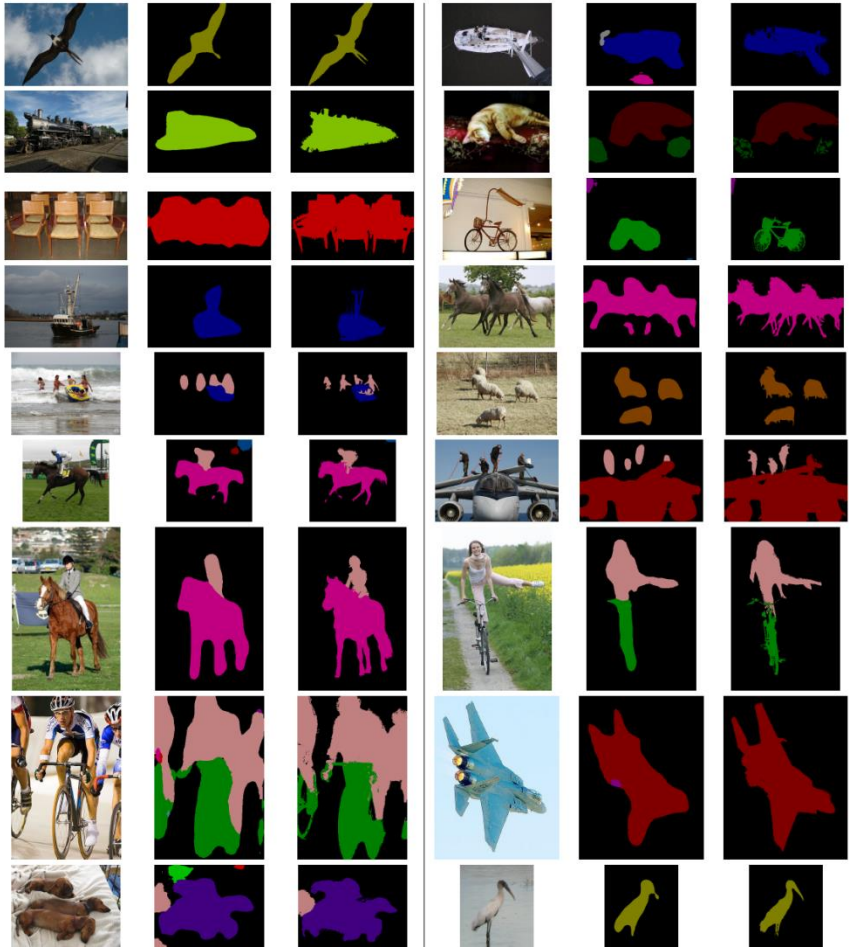

Chen, Papandreou, Kokkinos, Murphy, and Yuille, “Semantic Image Segmentation with Deep Convolutional Nets and Fully Connected CRFs”, ICLR 2015

DeepLab

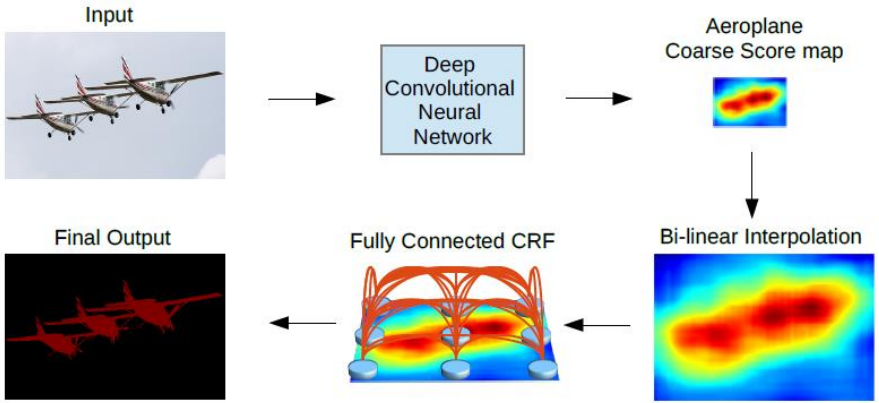

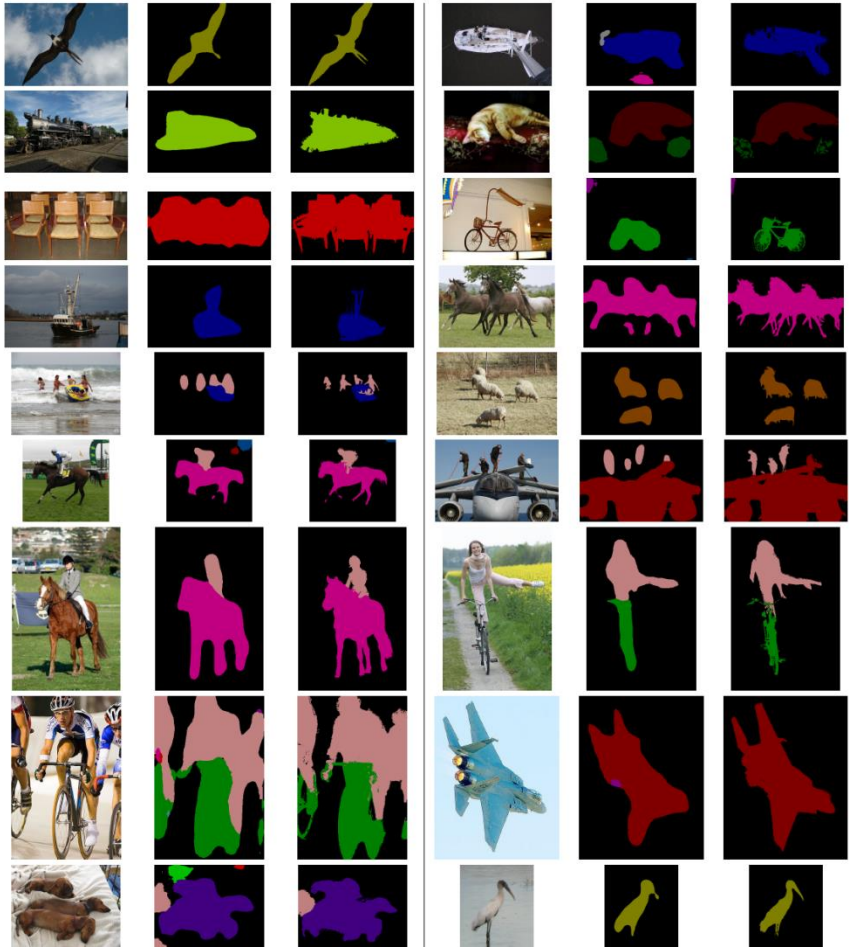

Chen, Papandreou, Kokkinos, Murphy, and Yuille, “Semantic Image Segmentation with Deep Convolutional Nets and Fully Connected CRFs”, ICLR 2015

CRF AS RECURRENT NEURAL NET (END-TO-END)

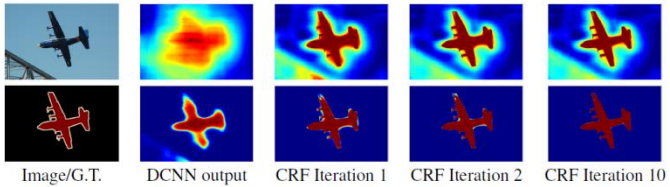

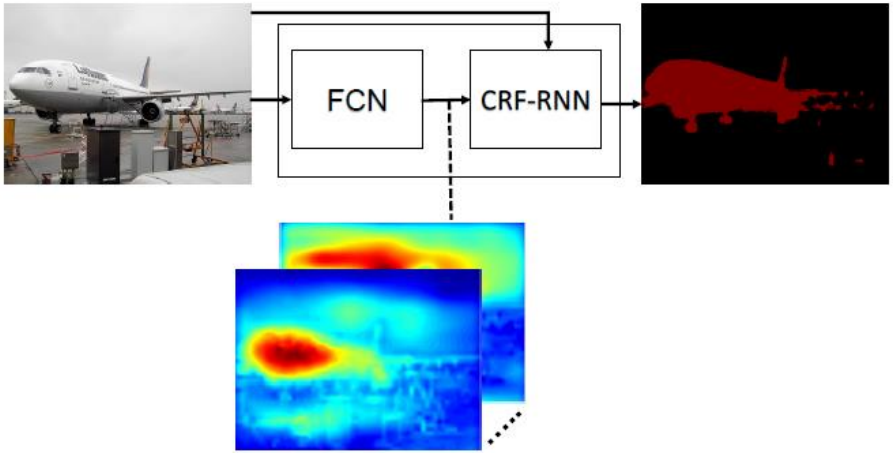

Zheng et al., “Conditional Random Fields as Recurrent Neural Networks”, ICCV 2015

CRF AS RECURRENT NEURAL NET (END-TO-END)

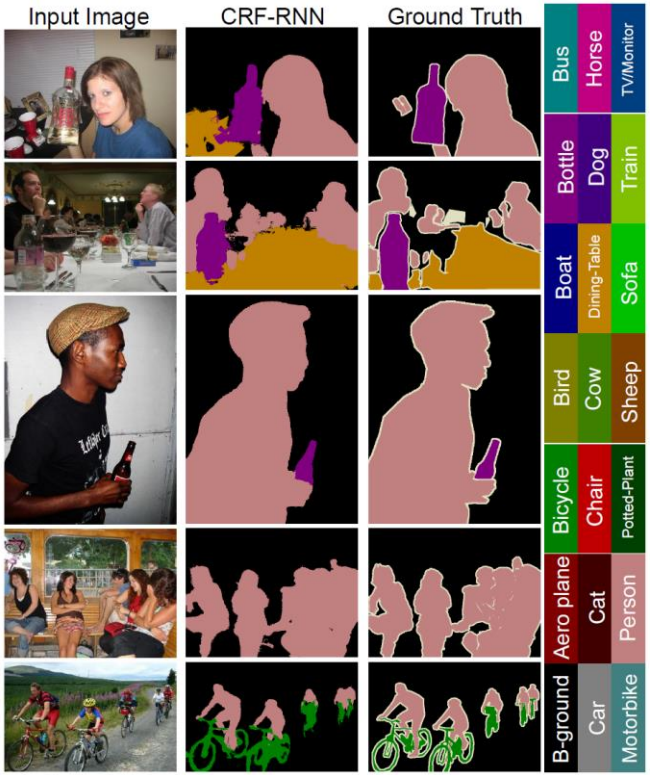

Zheng et al., “Conditional Random Fields as Recurrent Neural Networks”, ICCV 2015