

CS 4803 / 7643: Deep Learning

Topics:

- Convolutional Neural Networks
- What is a convolution?
- FC vs Conv Layers

Dhruv Batra
Georgia Tech

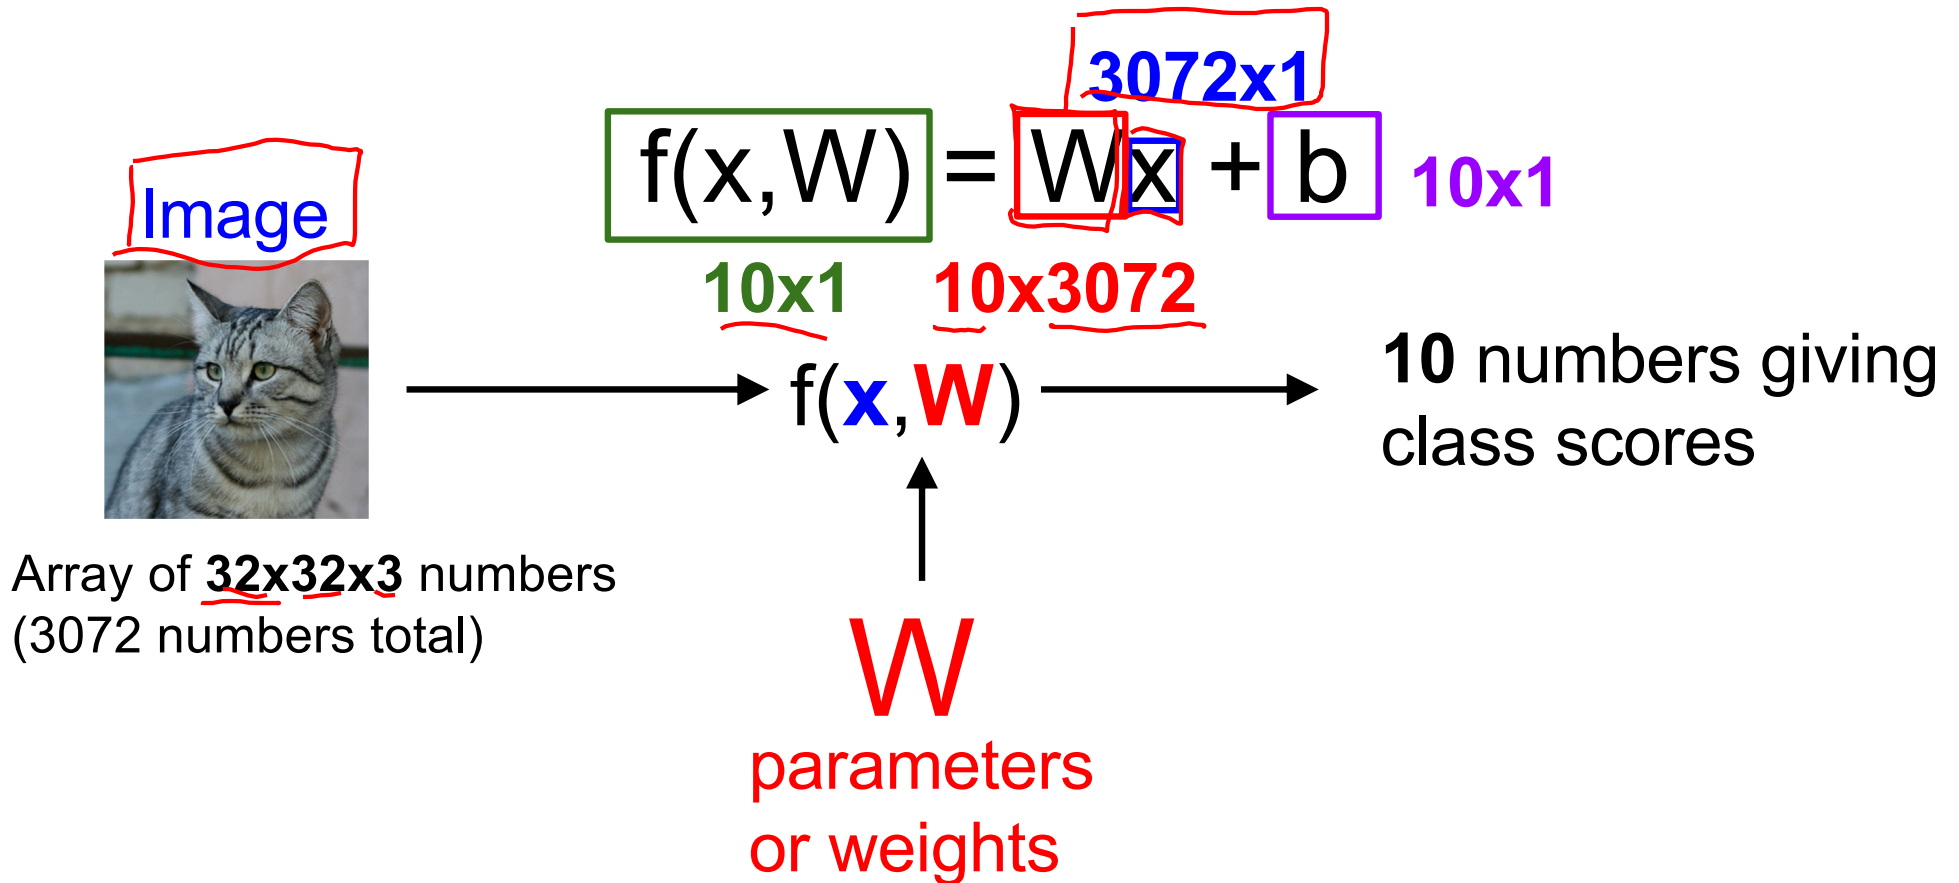
Administrativa

- HW1 Reminder
 - Due: 09/26, 11:55pm
 - <https://evalai.cloudcv.org/web/challenges/challenge-page/431/leaderboard/1200>
- Project Teams Google Doc
 - https://docs.google.com/spreadsheets/d/1ouD6ctaemV_3nb2MQHs7rUOAaW9DFLu8I5Zd3yOFs7E/edit?usp=sharing
 - Project Title
 - 1-3 sentence project summary TL;DR
 - Team member names

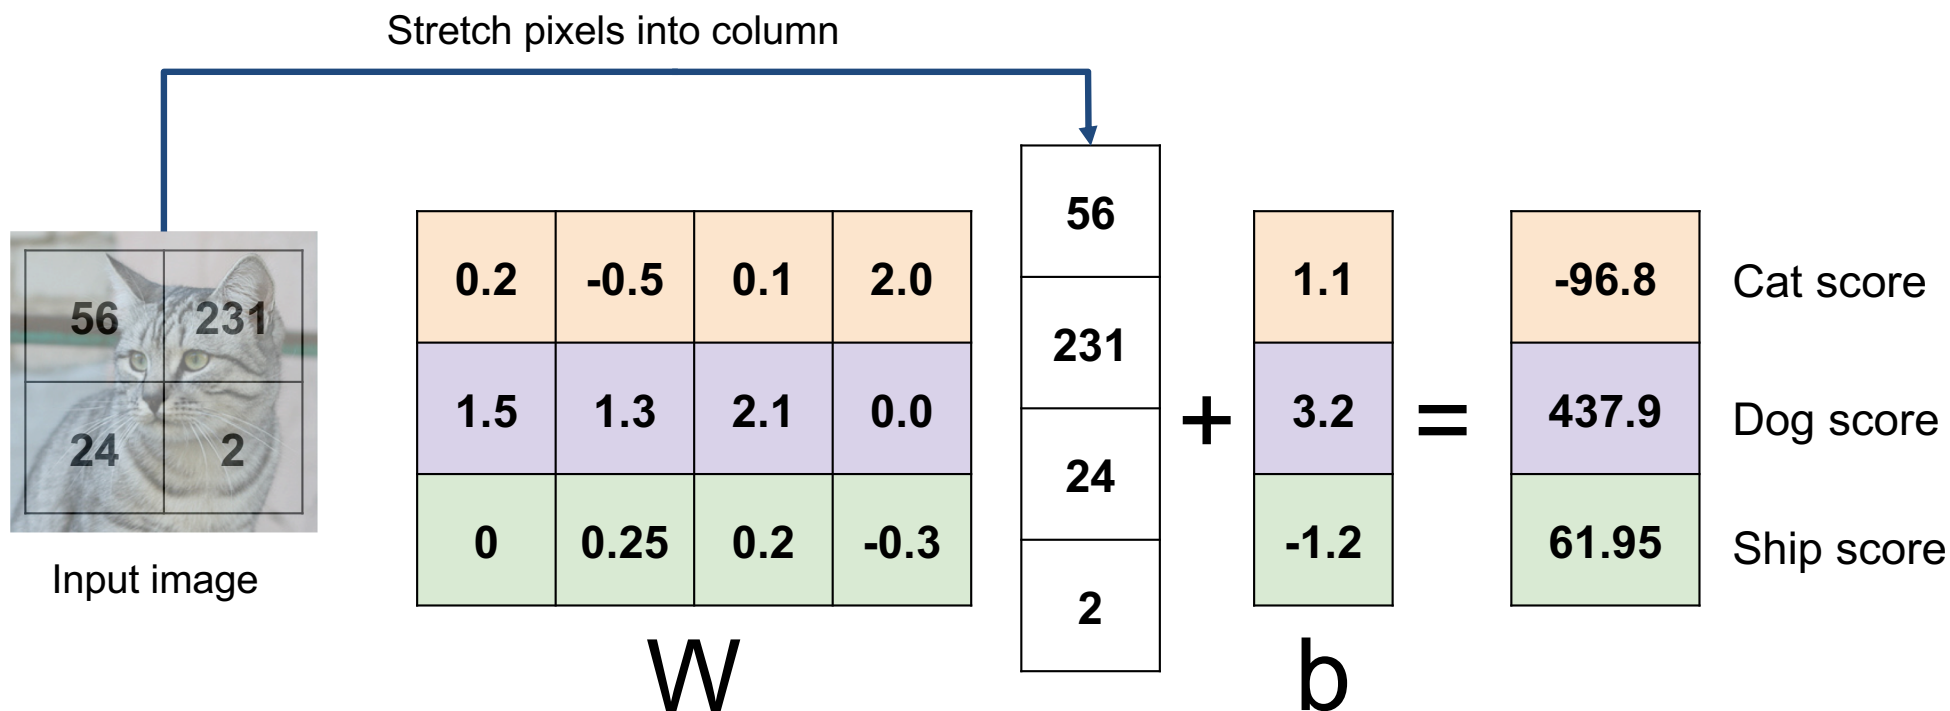
Plan for Today

- Convolutional Neural Networks
 - What is a convolution?
 - FC vs Conv Layers

Recall: Linear Classifier





Example with an image with 4 pixels, and 3 classes (cat/dog/ship)



Recall: (Fully-Connected) Neural networks

(**Before**) Linear score function: $f = \underline{W}x$

(**Now**) 2-layer Neural Network $f = W_2 \underline{\max(0, \underline{W_1}x)}$



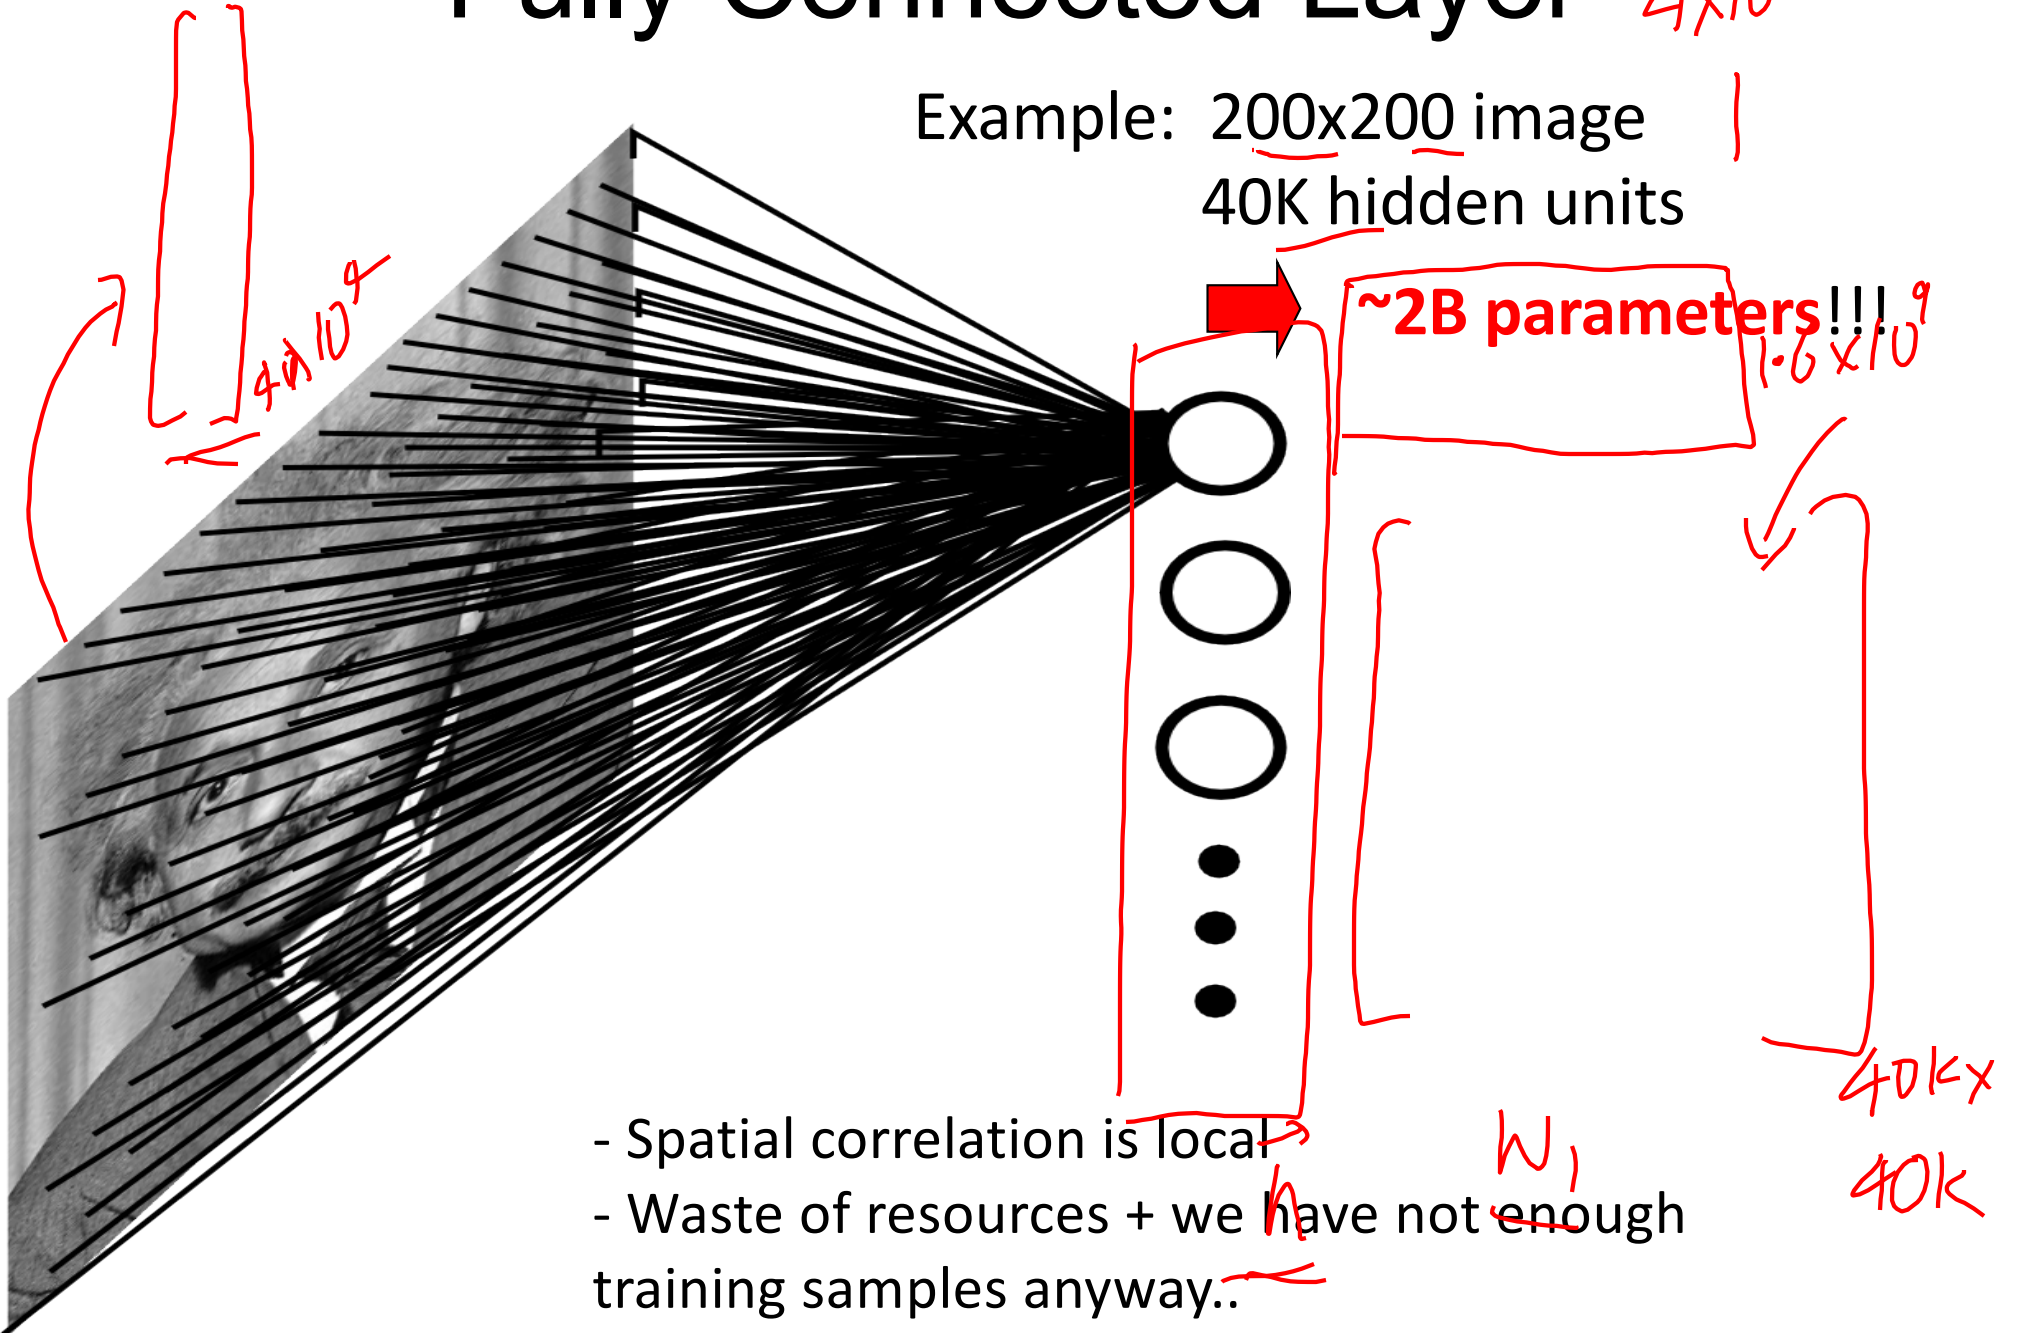


Convolutional Neural Networks

(without the brain stuff)


Fully Connected Layer 4×10^4

Example: 200x200 image
40K hidden units




Locally Connected Layer

Assumption #1



Example: 200x200 image
40K hidden units
"Filter" size: 10x10
4M parameters



Note:

This parameterization is good when input image is registered (e.g., face recognition). *40k x 100*

Locally Connected Layer


STATIONARITY?

Statistics similar at all locations

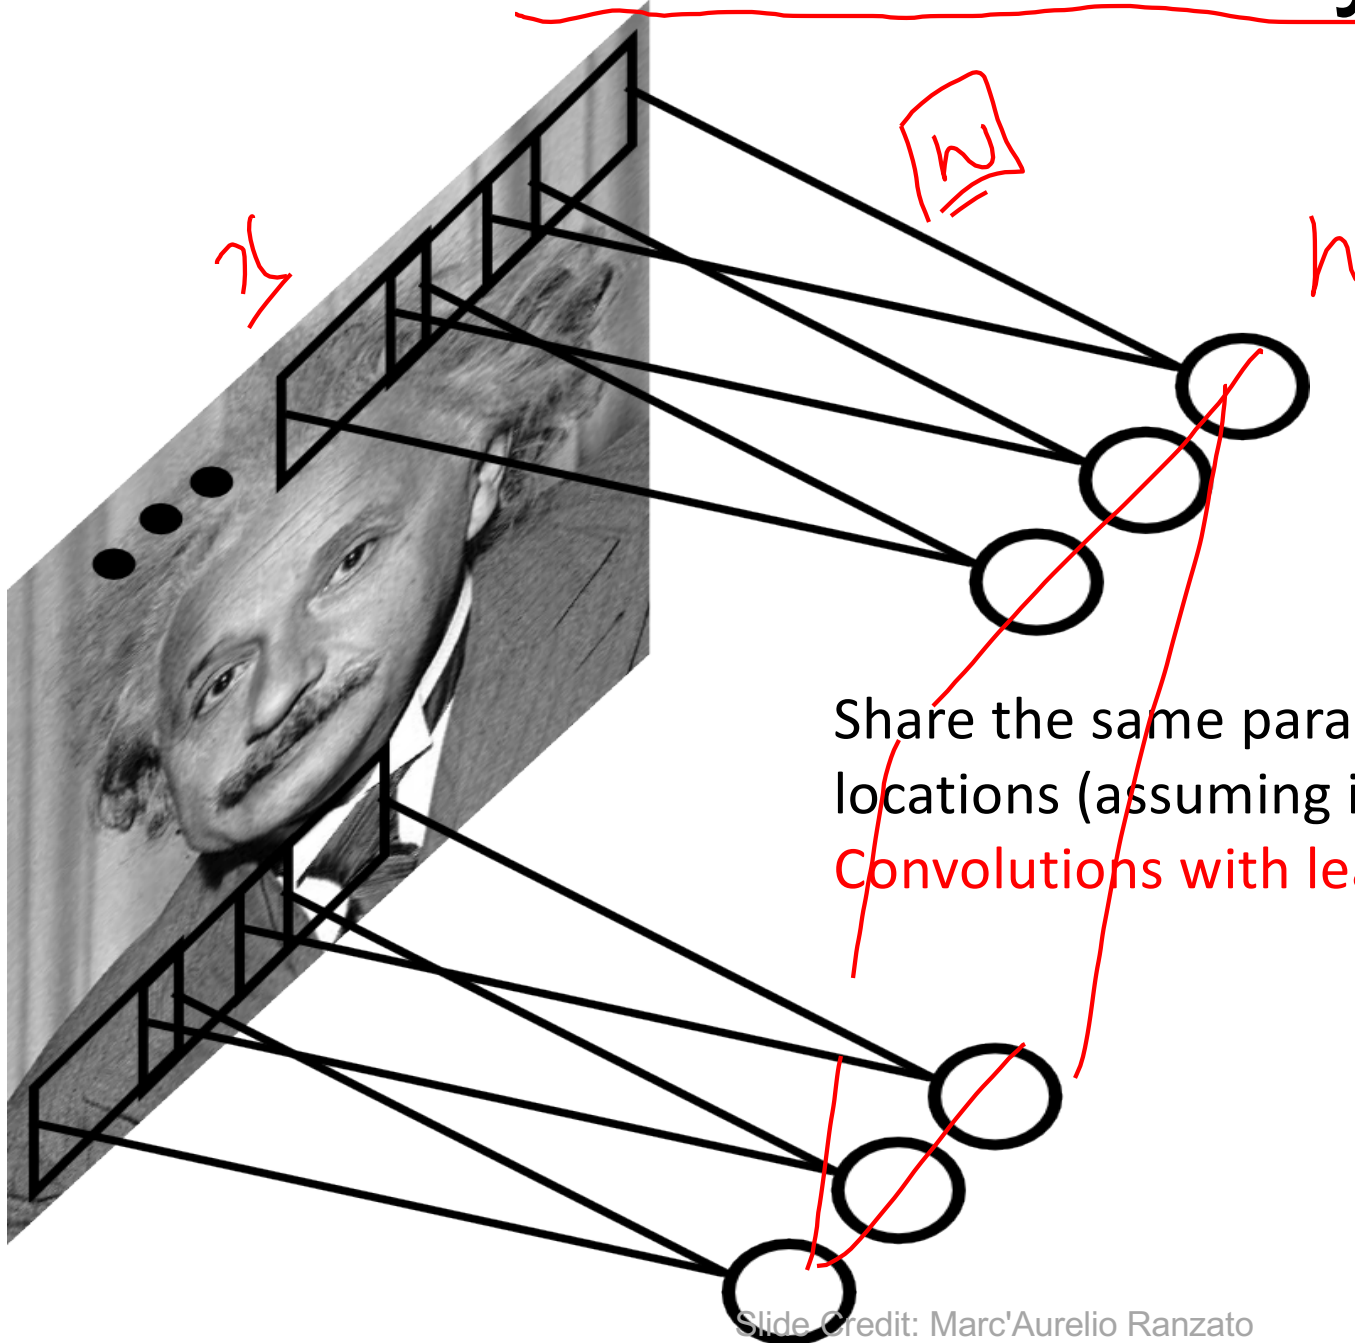
$2B \rightarrow 4M \rightarrow$

100

$w_i = w_j$



Convolutional Layer



Share the same parameters across different locations (assuming input is stationary):

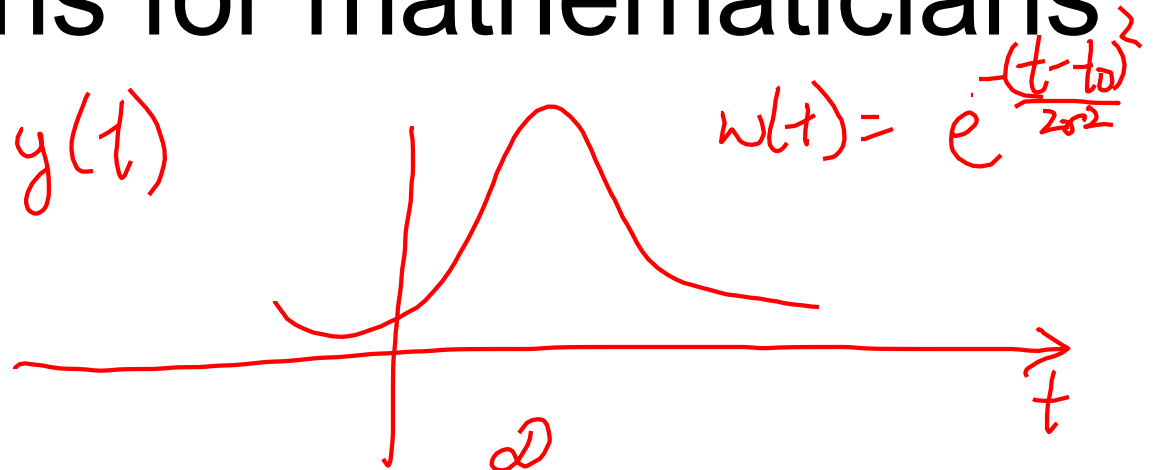
Convolutions with learned kernels

Convolutions!

math \rightarrow CS \rightarrow programming

Convolutions for mathematicians

$x(t)$ $w(t)$ $y(t)$



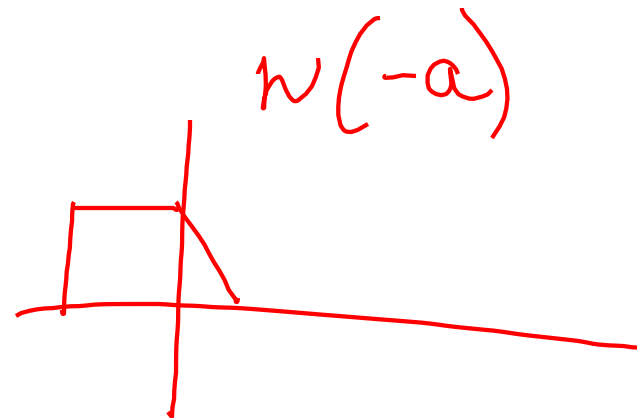
$$\begin{aligned} y(t) &= (x * w)(t) = \int_{-\infty}^{\infty} \underbrace{x(t-a)} \underbrace{w(a)} da \\ &= (w * x)(t) = \int_{-\infty}^{\infty} \underbrace{x(a)} w(t-a) da \end{aligned}$$

Convolutions for mathematicians


$$y(t) = \int_{-\infty}^{\infty} x(a) w(t-a) da$$





flip
about
y-axis





translate by
 t





"Convolution of box signal with itself2" by Convolution_of_box_signal_with_itself.gif: Brian Amberg derivative work: Tinos (talk) - Convolution_of_box_signal_with_itself.gif. Licensed under CC BY-SA 3.0 via Commons - https://commons.wikimedia.org/wiki/File:Convolution_of_box_signal_with_itself2.gif#/media/File:Convolution_of_box_signal_with_itself2.gif



"Convolution of box signal with itself2" by Convolution_of_box_signal_with_itself.gif: Brian Amberg derivative work: Tinos (talk) - Convolution_of_box_signal_with_itself.gif. Licensed under CC BY-SA 3.0 via Commons - https://commons.wikimedia.org/wiki/File:Convolution_of_box_signal_with_itself2.gif#/media/File:Convolution_of_box_signal_with_itself2.gif

Convolutions for mathematicians

- One dimension

$$y(t) = \int_{-\infty}^{\infty} x(t-a)w(a)da$$


- Two dimensions

$$y(t_1, t_2) = \iint_{-\infty}^{\infty} x(t_1-a, t_2-b)w(a,b)dadb$$

$a = -\infty$ $b = -\infty$

Convolutions for computer scientists


$$y[t_1, t_2] = \sum_{a=-\infty}^{\infty} \sum_{b=-\infty}^{\infty} x[t_1 - a, t_2 - b] w[a, b]$$




Convolutions for programmers

$$y[r, c] = \sum_{a=0}^{k_1-1} \sum_{b=0}^{k_2-1} x[r+a, c+b] w[a, b]$$


$\boxed{\text{rows}}$ $\boxed{\text{cols}}$




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer




Convolutional Layer



Convolutional Layer




Convolutional Layer



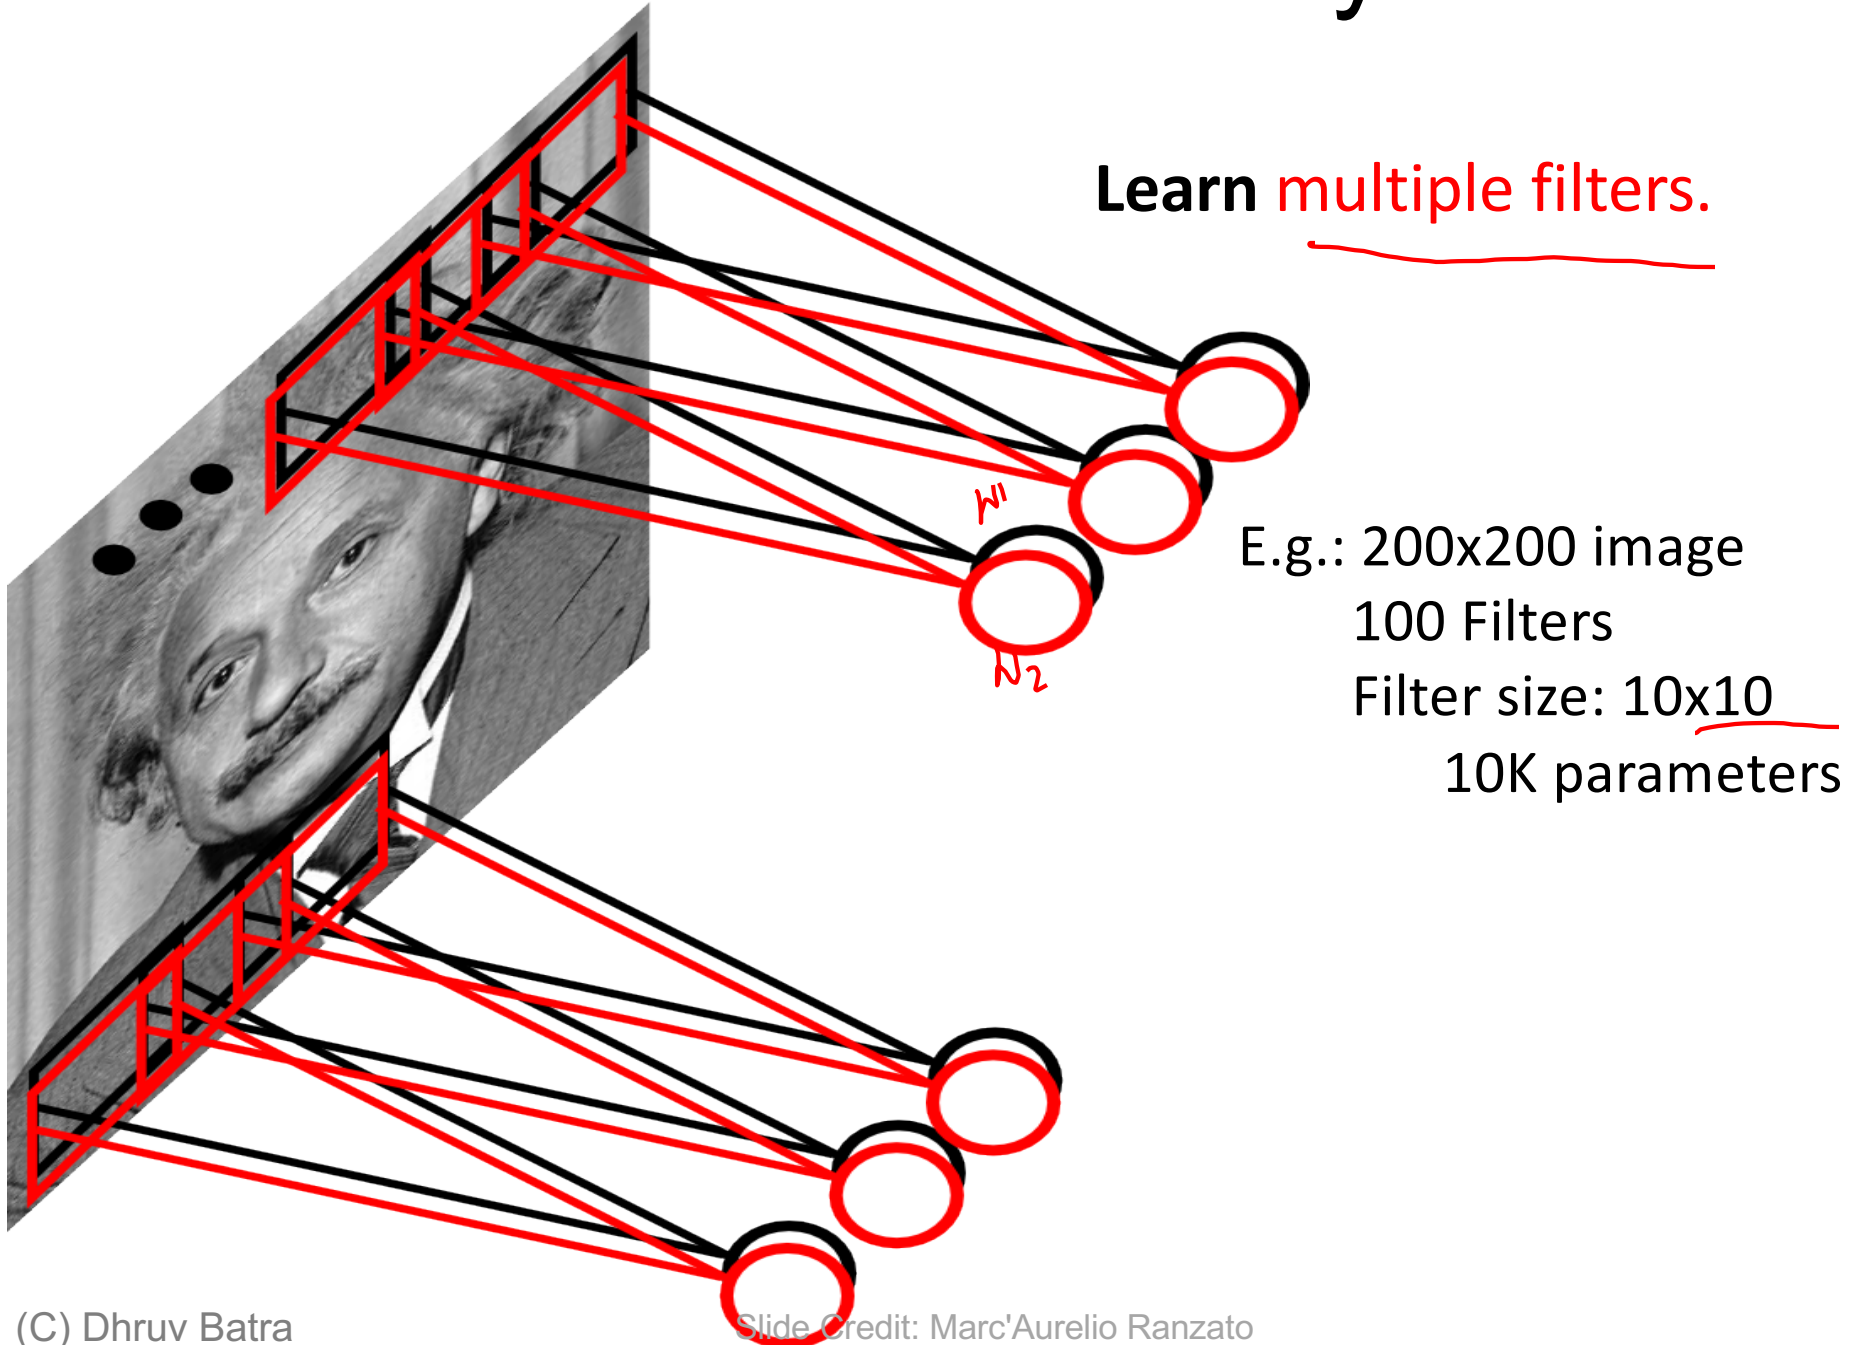
Convolution Explained

- <http://setosa.io/ev/image-kernels/>
- <https://github.com/bruckner/deepViz>


Convolutional Layer



Convolutional Layer



FC vs Conv Layer




$$h \in \mathbb{R}^{C_{l-1}}$$

$$h_i^l = \sum_{j=0}^{C_{l-1}} h_j^{l-1} W_{ij}$$

$$+ \text{bias}(i)$$

scalar product

$$(W_1^T \dots W_{C_2}^T) \quad C_1 \times C_2$$



$$H_i^l = \sum_{j=0}^{C_{l-1}} H_j^{l-1} * W_{ij} + \text{bias}(i)$$

$$+ \text{bias}(i)$$

$$H_i^l [a, c]$$


$$= \sum_{j=0}^{C_{l-1}} \sum_{a=0}^{k_1-1} \sum_{b=0}^{k_2-1} H_j^{l-1} [a+a, c+b] W_{ij} [a, b]$$

$$H_j^{l-1} [a+a, c+b] W_{ij} [a, b]$$

$$C_2 = (k_1 \times k_2 \times (C_{l-1} - k_1 + 1) - 1) + 1$$


Convolution Layer

32x32x3 image -> preserve spatial structure




Convolution Layer

32x32x3 image




$k_1 \times k_2 \times C_1$
5x5x3 filter



Convolve the filter with the image
i.e. “slide over the image spatially,
computing dot products”

Convolution Layer

32x32x3 image



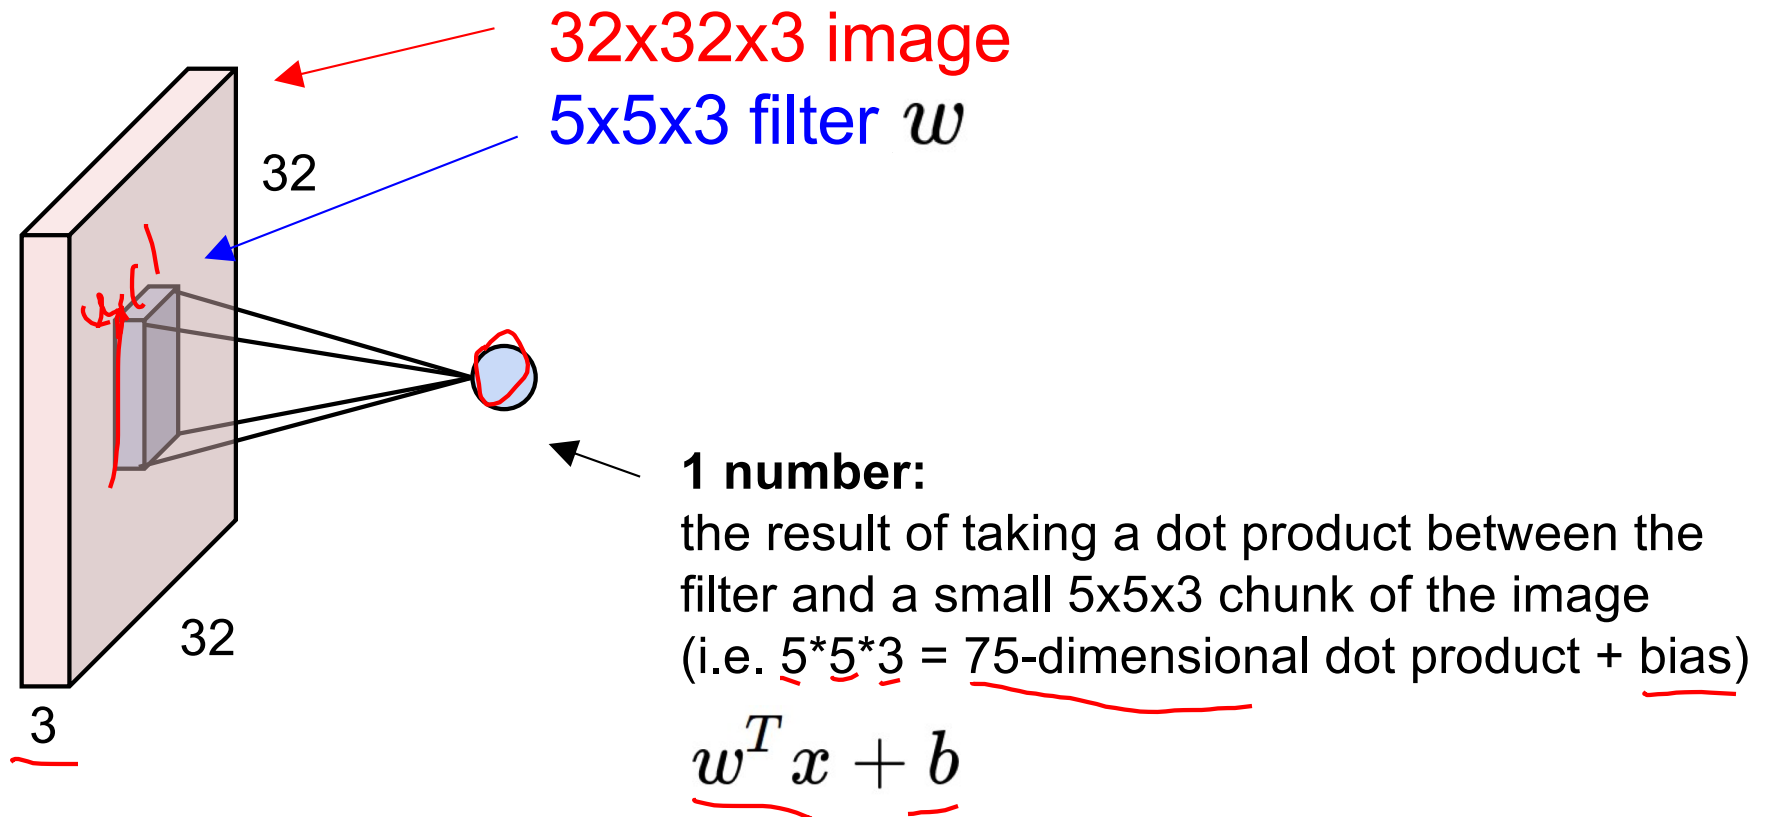
Filters always extend the full depth of the input volume

5x5x3 filter




Convolve the filter with the image
i.e. “slide over the image spatially,
computing dot products”

Convolution Layer




Convolution Layer




Convolution Layer

consider a second, **green** filter

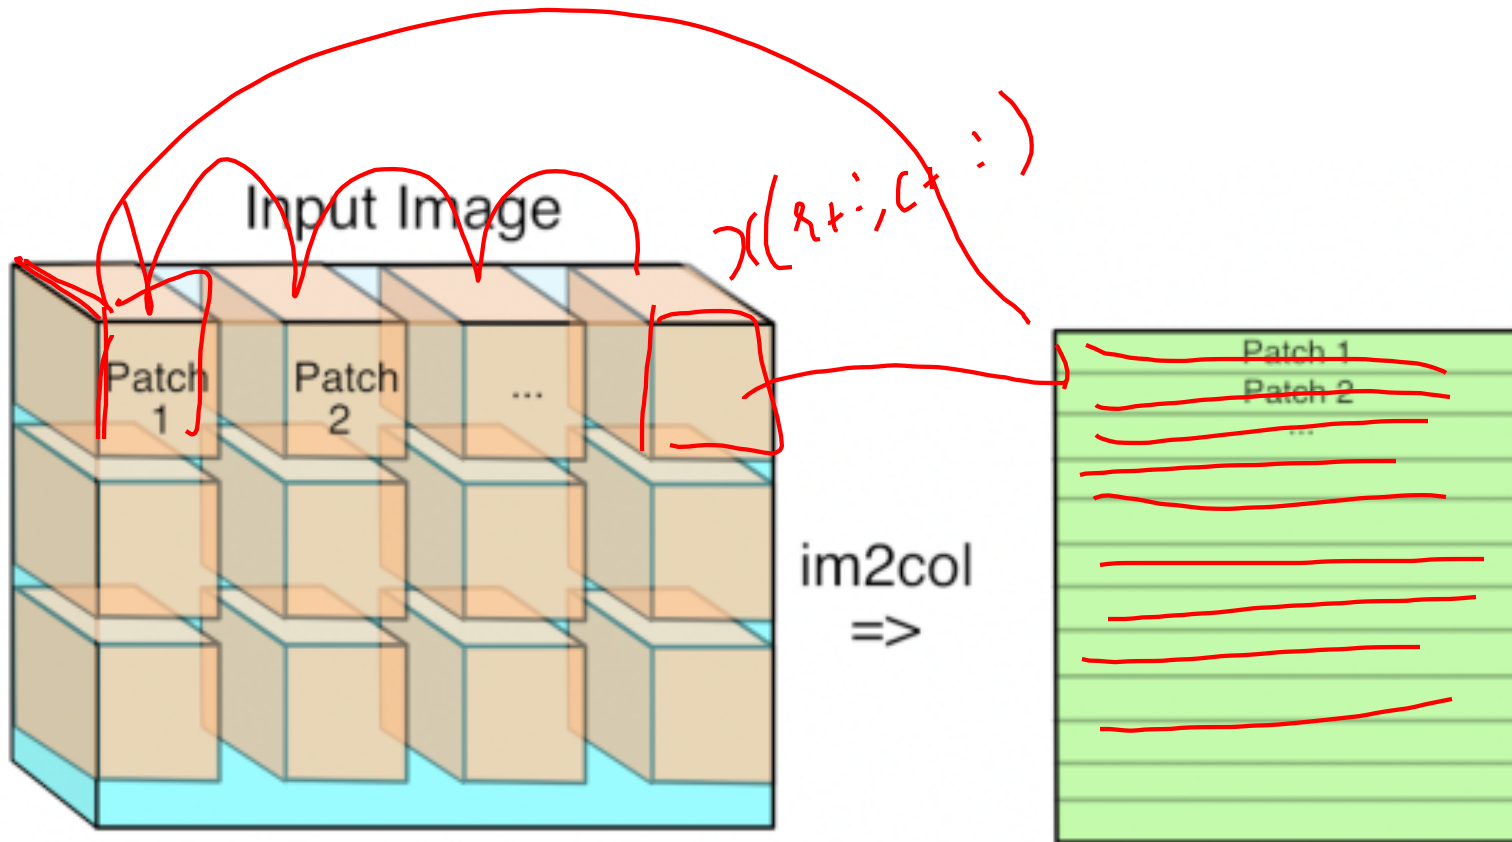


For example, if we had 6 5x5 filters, we'll get 6 separate activation maps:




We stack these up to get a "new image" of size 28x28x6!

Im2Col




GEMM




Time Distribution of AlexNet


GPU Forward Time Distribution



CPU Forward Time Distribution



Preview: ConvNet is a sequence of Convolution Layers, interspersed with activation functions



Preview: ConvNet is a sequence of Convolutional Layers, interspersed with activation functions

