Deep Networks

Vittal Premachandran

Supervised Learning $y=f(x)$

- Binary Classification:
  - Given $\mathbf{x}$ find $y \in \{-1, 1\}$
- Multiclass Classification
  - Given $\mathbf{x}$ find $y \in \{1,...,k\}$
- Regression
  - Given $\mathbf{x}$ find $y \in R$ (or, R^d)

Traditional Recognition Approach

Input data (pixels) → feature representation (hand-crafted) → Learning Algorithm (e.g., SVM)

Image

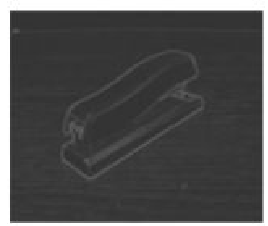

Low-level vision features (edges, SIFT, HOG, etc.)

Object detection / classification

Computer vision features

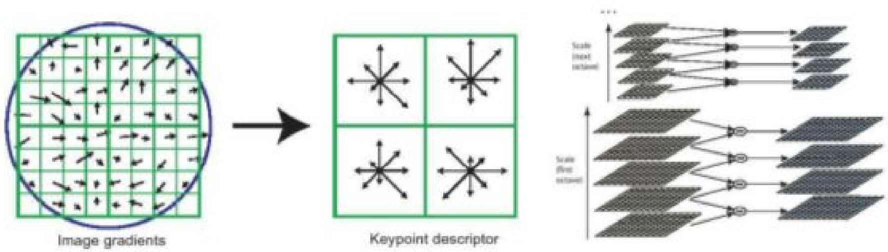

SIFT

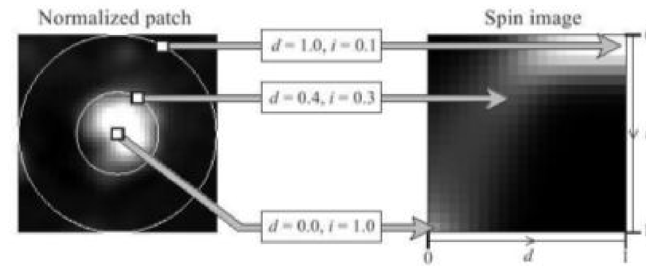

Spin image

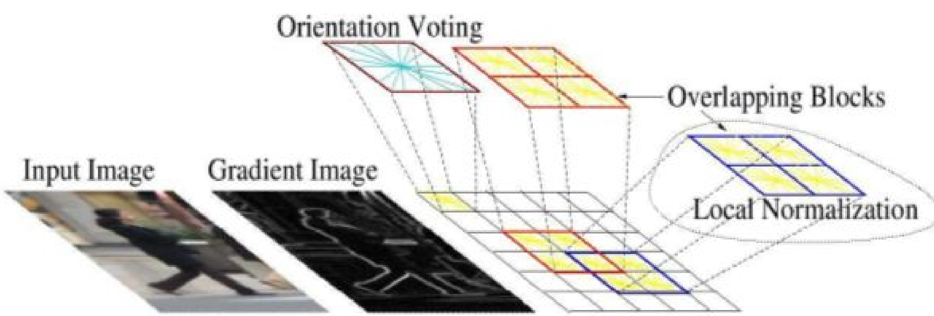

HoG

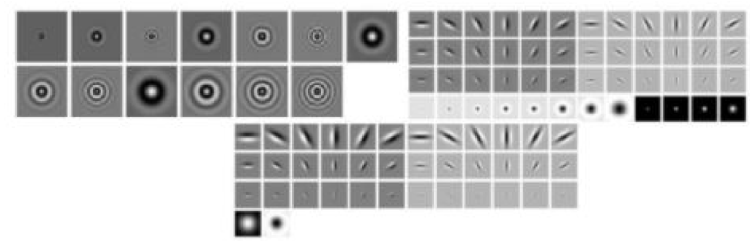

Textons

and many others:

SURF, MSER, LBP, Color-SIFT, Color histogram, GLOH, …..

Computer vision features

SIFT

Computer vision features

SIFT

HOG

Computer vision features

SIFT

LBP

HOG

Computer vision features

SIFT

LBP

HOG

SURF

Computer vision features

SIFT

LBP

HOG

MSER

SURF

Computer vision features

SIFT

LBP

Textons

HOG

MSER

SURF

Computer vision features

SIFT

LBP

GLOH

Textons

HOG

MSER

SURF

Computer vision features

SIFT

LBP

GLOH

Textons

HOG

Color-SIFT

MSER

SURF

Computer vision features

SIFT

SPIN

LBP

GLOH

Textons

HOG

Color-SIFT

MSER

SURF

Computer vision features

SIFT

SPIN

LBP

GLOH

Textons

HOG

PHOG

Color-SIFT

MSER

SURF

Computer vision features

SIFT

SPIN

LBP

GLOH

Textons

HOG

PHOG

Color-SIFT

MSER

SURF

...

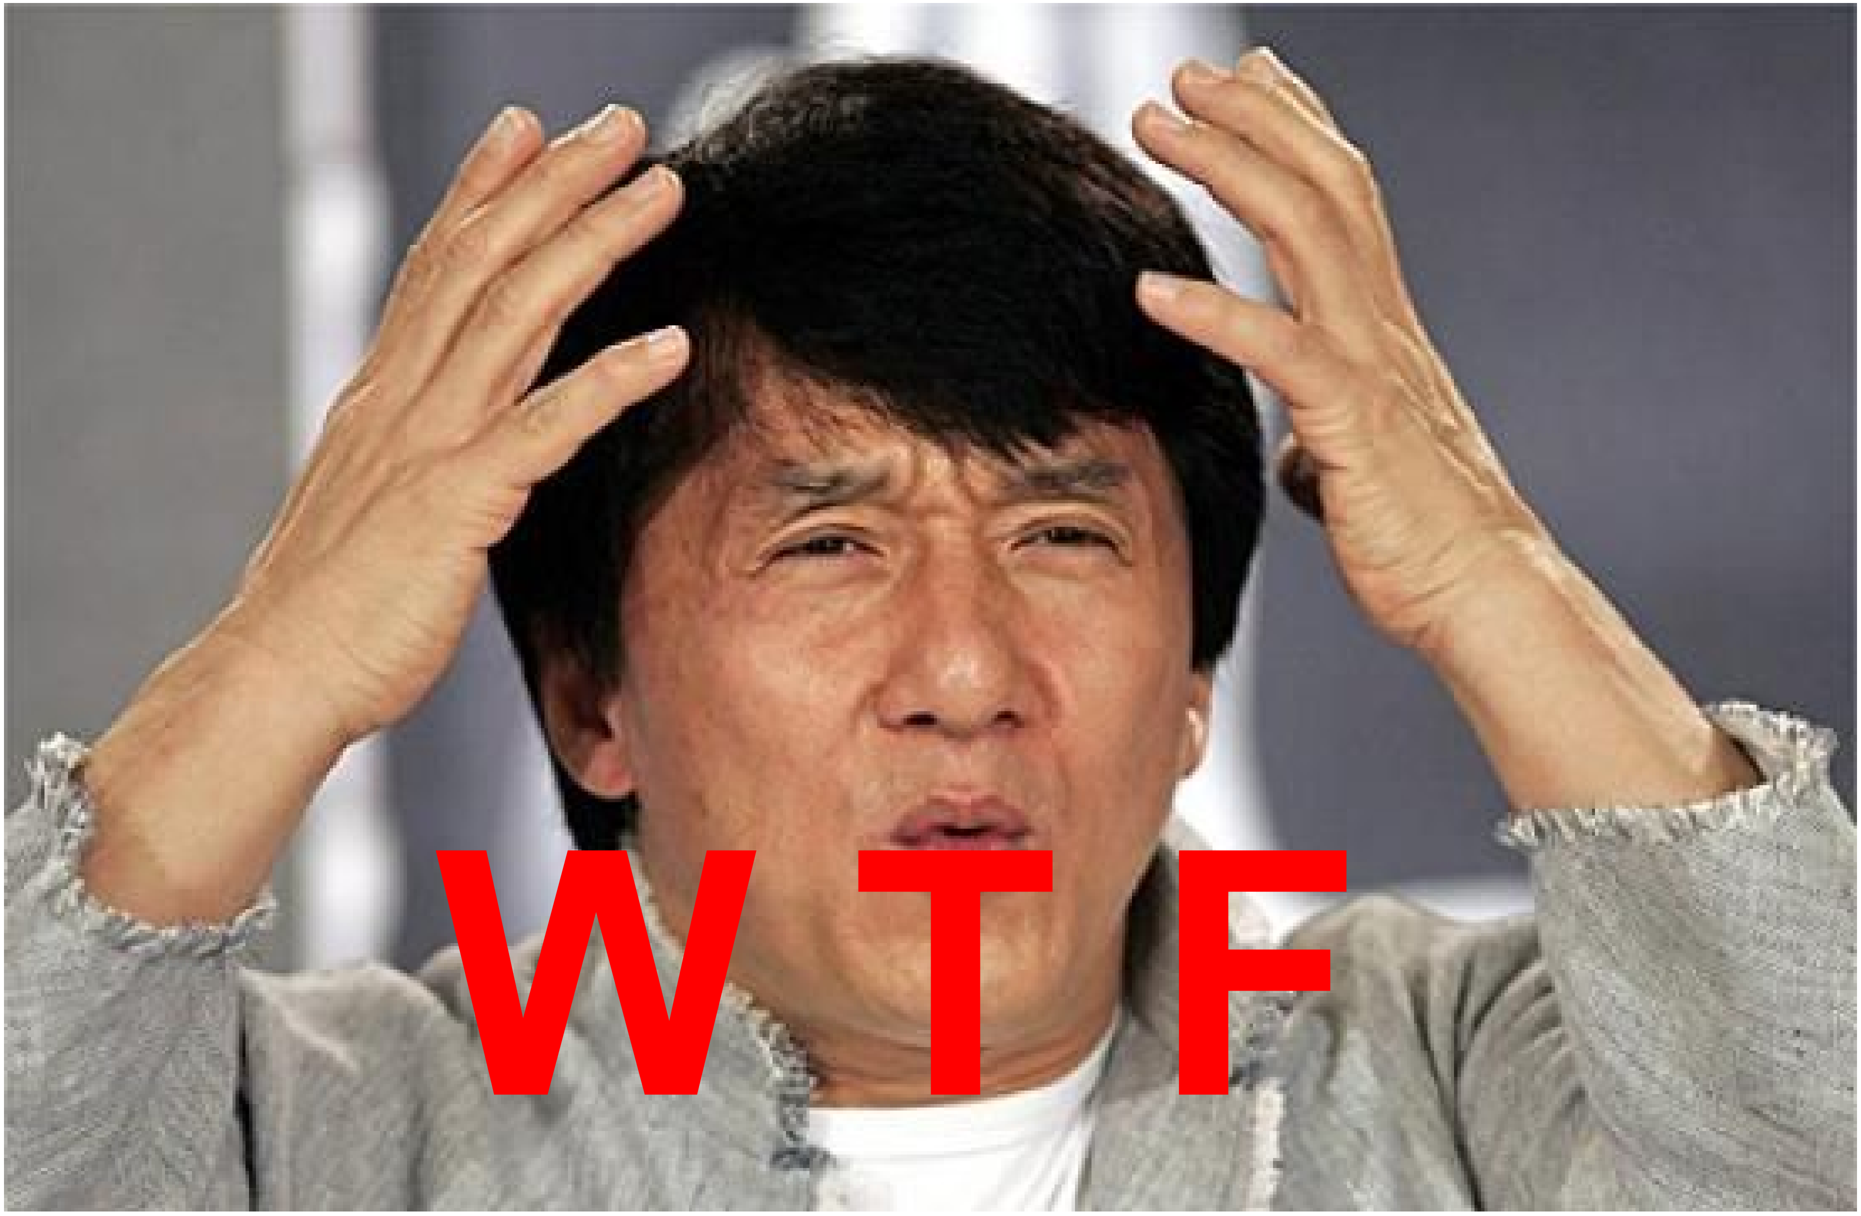
W T F


Which
Training
Features to use?

Can we learn the features (internal representations)?

Recall: The Neuron Metaphor

- Neurons
  - accept information from multiple inputs,
  - transmit information to other neurons.
- Multiply inputs by weights along edges
- Apply some function to the set of inputs at each node

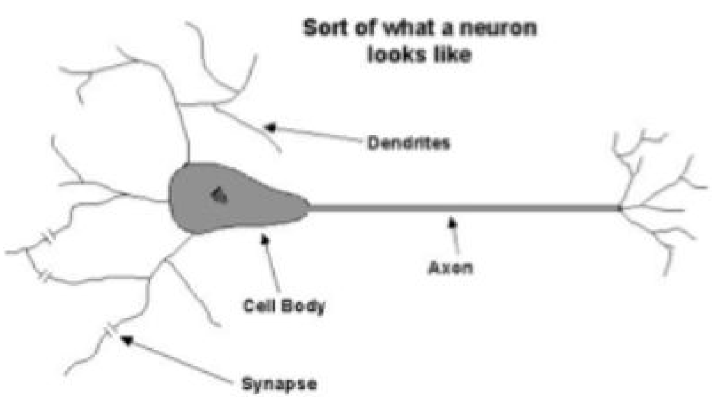

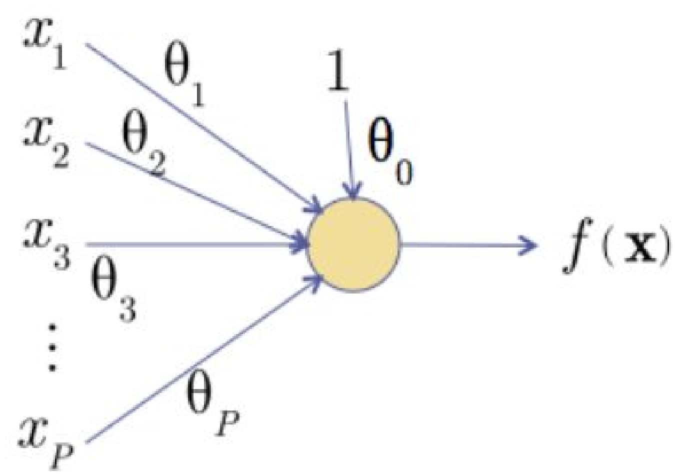

Types of Neurons

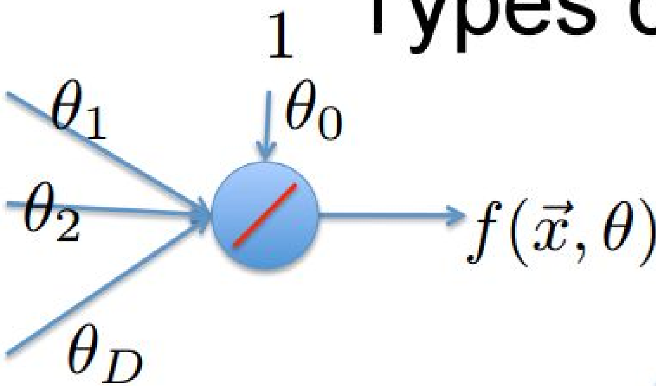

Linear Neuron

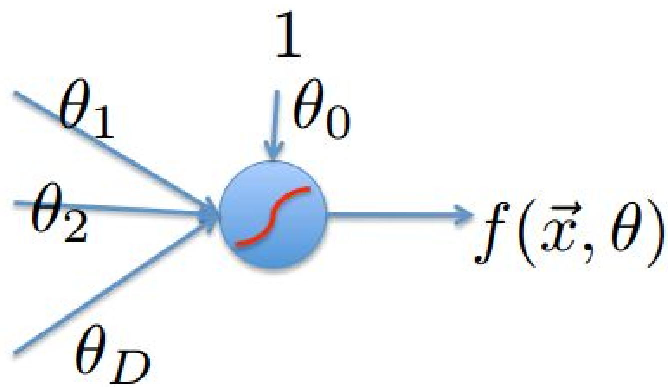

Logistic Neuron

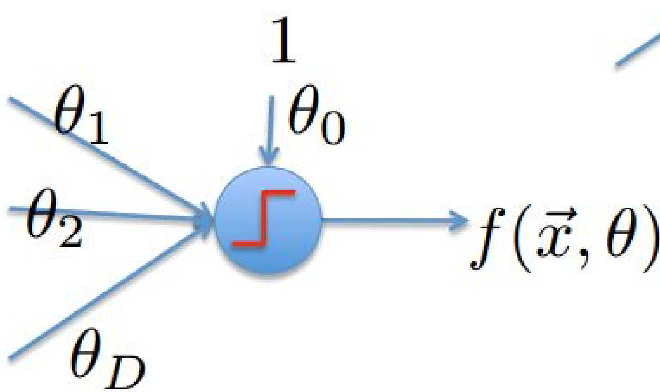

Perceptron

Potentially more. Require a convex loss function for gradient descent training.

Limitation

- A single “neuron” is still a linear decision boundary
- What to do?
- Idea: Stack a bunch of them together!

Multilayer Networks

- Cascade Neurons together
- The output from one layer is the input to the next
- Each Layer has its own sets of weights

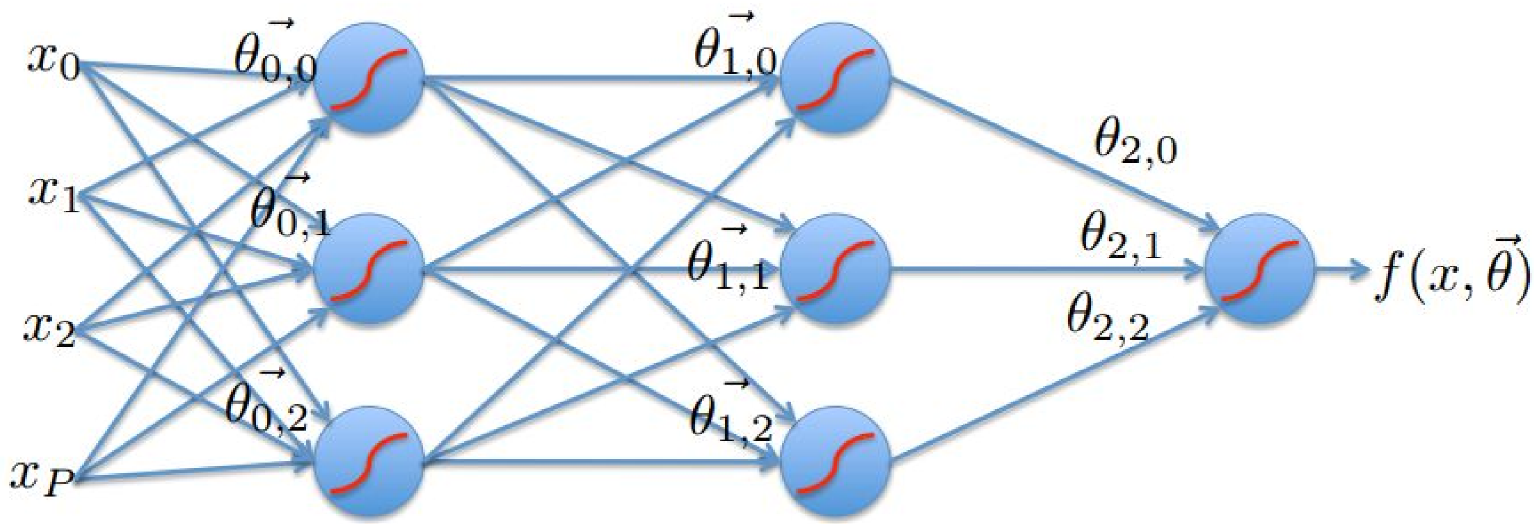

A quick note

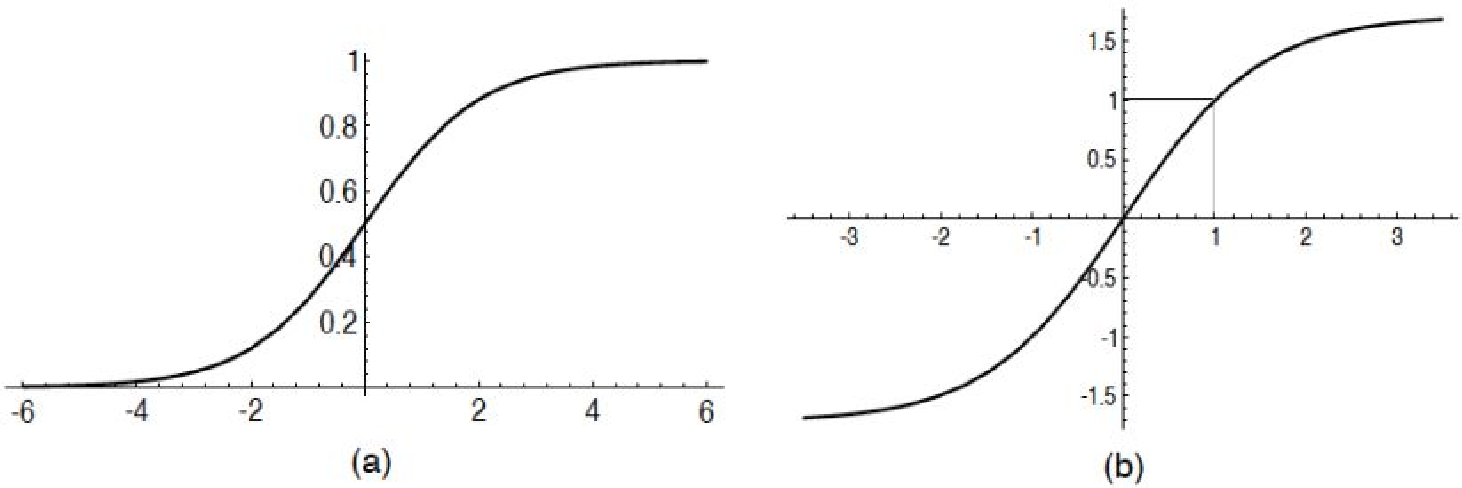

Fig. 4. (a) Not recommended: the standard logistic function, $f(x) = 1/(1+e^{-x})$. (b) Hyperbolic tangent, $f(x) = 1.7159 \tanh\left(\frac{2}{3}x\right)$.

Rectified Linear Units (ReLU)

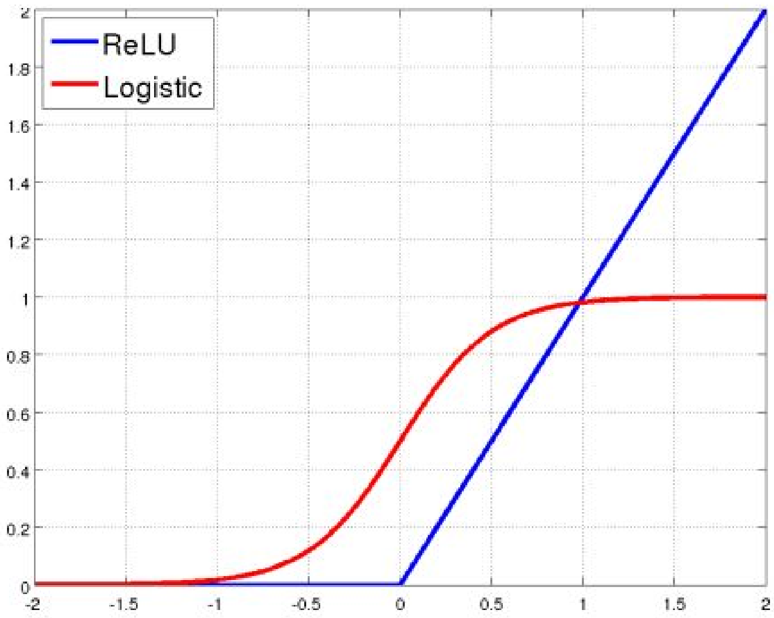

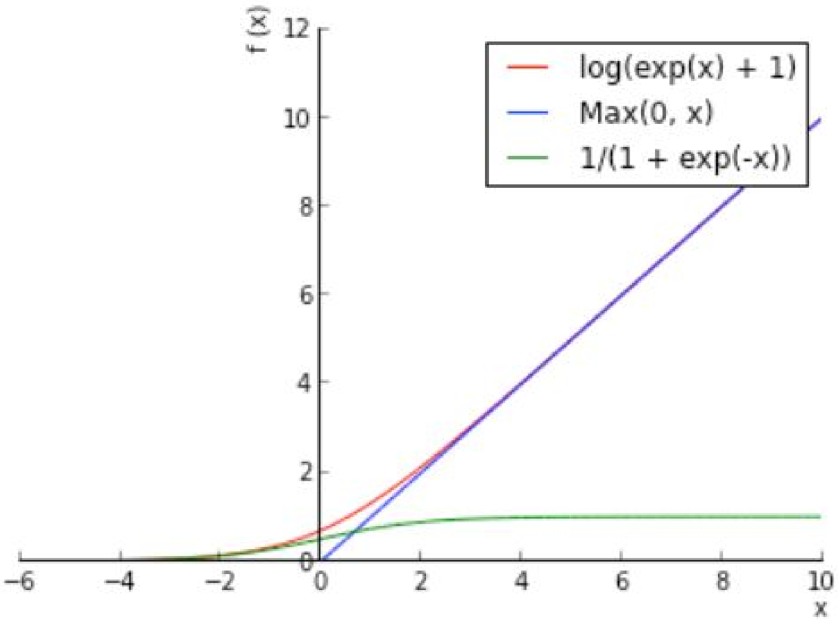

Supervised Learning

$\{(\boldsymbol{x}^i, y^i), i=1 \ldots P\}$ training dataset

$\boldsymbol{x}^i$ i-th input training example

y^i i-th target label

P number of training examples

Goal: predict the target label of unseen inputs.

Supervised Learning: Examples

Classification

"dog"

classification

Denoising

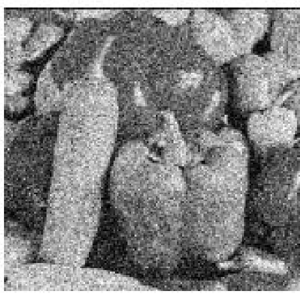

regression

OCR

"2 3 4 5"

structured prediction

Supervised Deep Learning

Classification

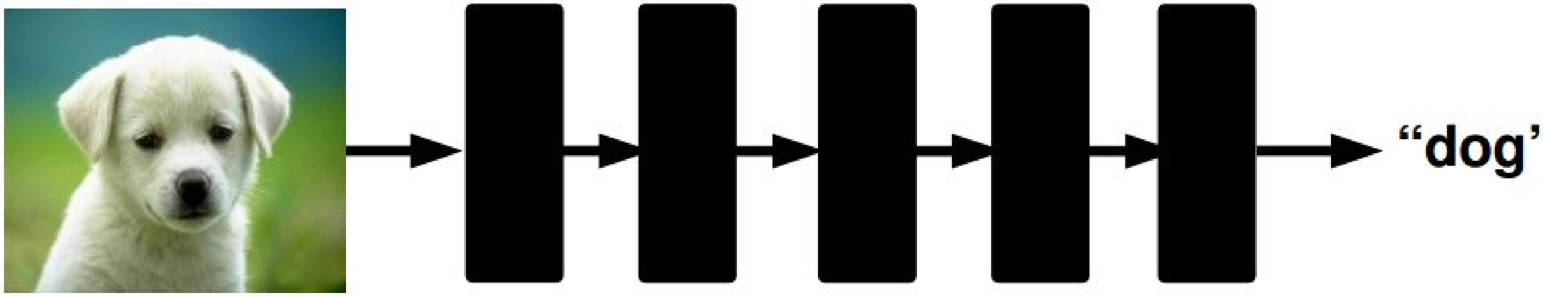

Denoising

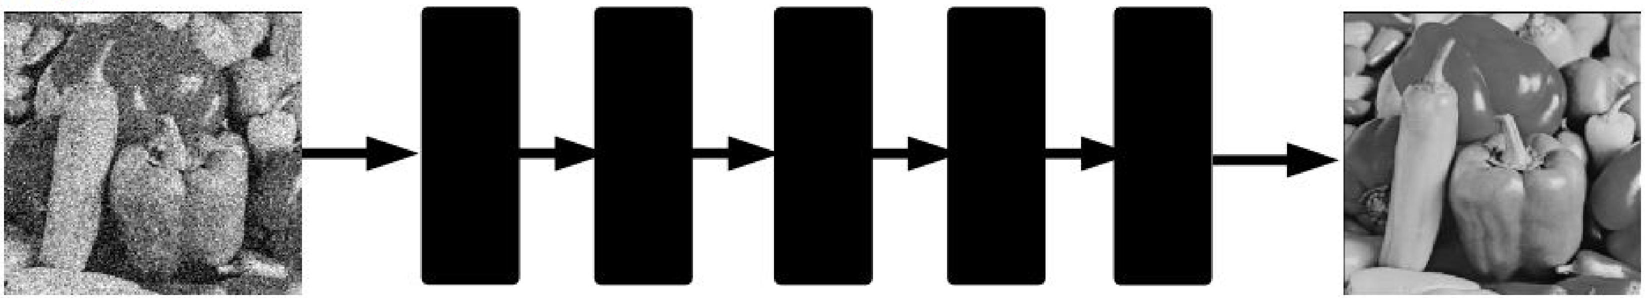

OCR

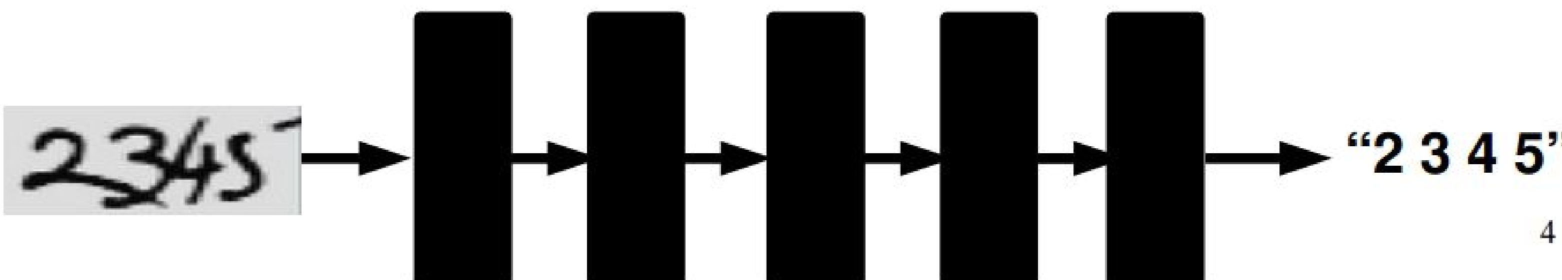

Forward Propagation

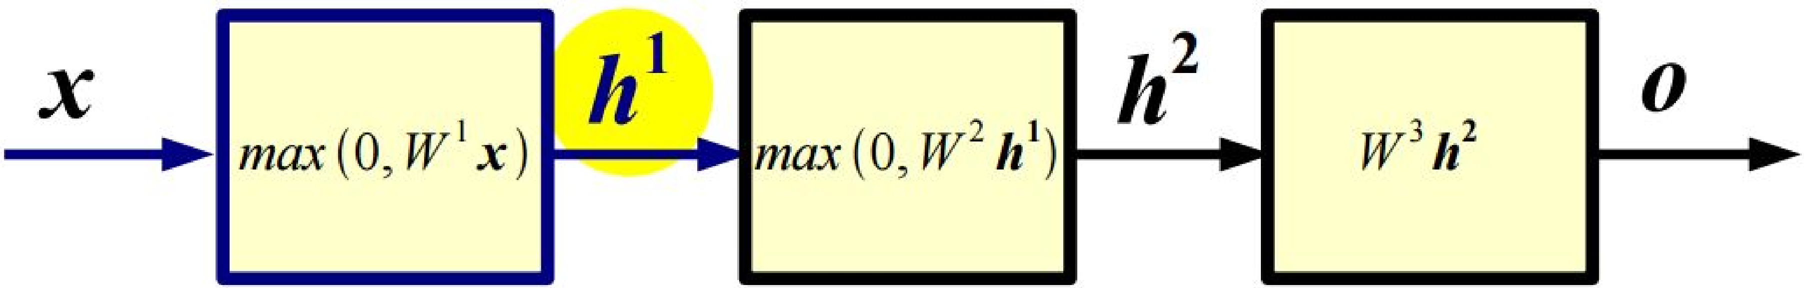

$\boldsymbol{x} \in R^D \quad W^1 \in R^{N_1 \times D} \quad \boldsymbol{b}^1 \in R^{N_1} \quad \boldsymbol{h}^1 \in R^{N_1}$

$$\boldsymbol{h}^1 = max(0, W^1 \boldsymbol{x} + \boldsymbol{b}^1)$$

W^1 1-st layer weight matrix or weights
$\boldsymbol{b}^1$ 1-st layer biases

The non-linearity $u = max(0, v)$ is called **ReLU** in the DL literature. Each output hidden unit takes as input all the units at the previous layer: each such layer is called "**fully connected**".

Forward Propagation

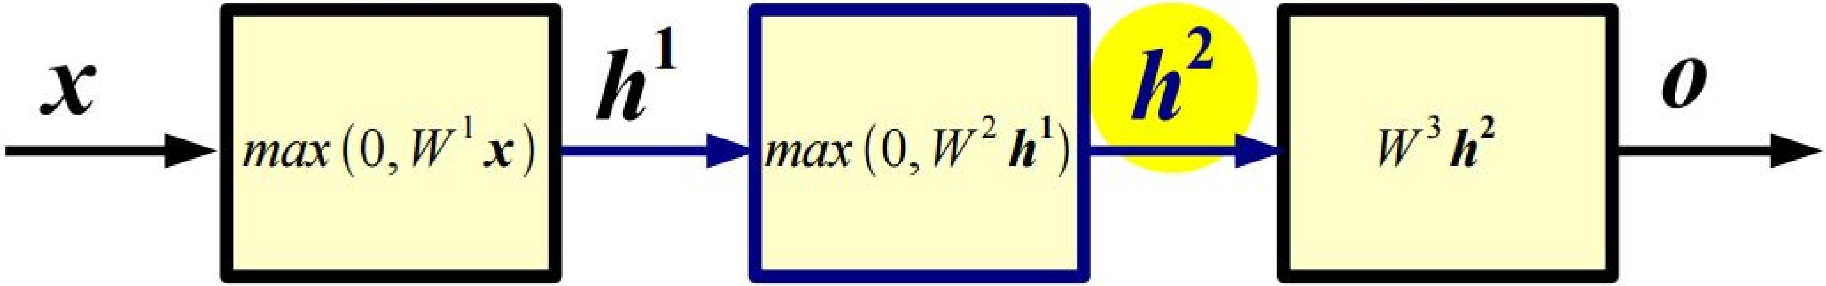

$h^1 \in R^{N_1} \quad W^2 \in R^{N_2 \times N_1} \quad b^2 \in R^{N_2} \qquad h^2 \in R^{N_2}$

$$h^2 = max(0, W^2 h^1 + b^2)$$

W^2 2-nd layer weight matrix or weights
b^2 2-nd layer biases

Ranzato

Forward Propagation

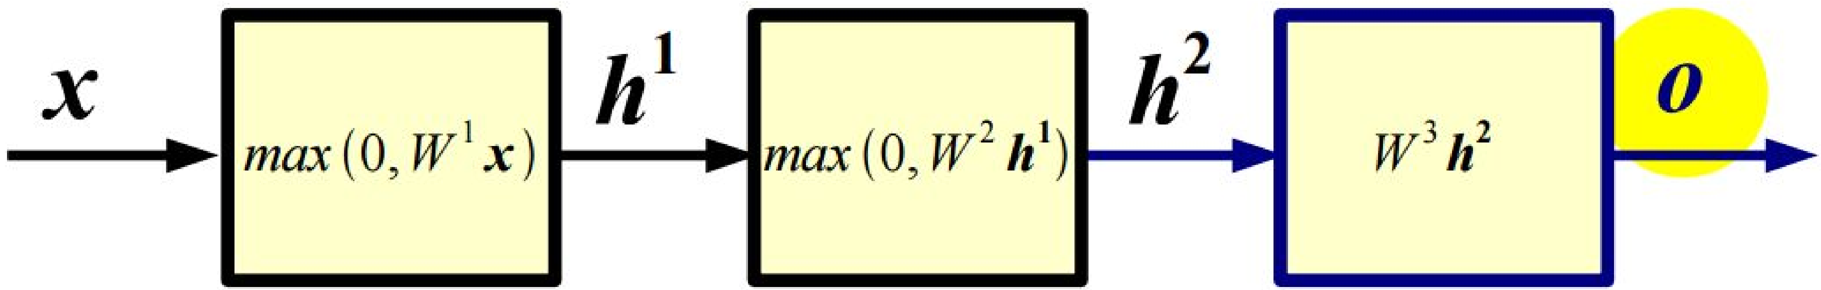

$\boldsymbol{h}^2 \in R^{N_2} \quad W^3 \in R^{N_3 \times N_2} \quad \boldsymbol{b}^3 \in R^{N_3} \qquad \boldsymbol{o} \in R^{N_3}$

$$\boldsymbol{o} = max(0, W^3 \boldsymbol{h}^2 + \boldsymbol{b}^3)$$

W^3 3-rd layer weight matrix or weights
$\boldsymbol{b}^3$ 3-rd layer biases

Training

Learning consists of minimizing the loss (plus some regularization term) w.r.t. parameters over the whole training set.

$$\boldsymbol{\theta}^* = arg\ min_{\boldsymbol{\theta}} \sum_{n=1}^{P} L(\boldsymbol{x}^n, y^n; \boldsymbol{\theta})$$

Question: How to minimize a complicated function of the parameters?

Answer: Chain rule, a.k.a. **Backpropagation**! That is the procedure to compute gradients of the loss w.r.t. parameters in a multi-layer neural network.

Rumelhart et al. "Learning internal representations by back-propagating.." Nature 1986

Backward Propagation

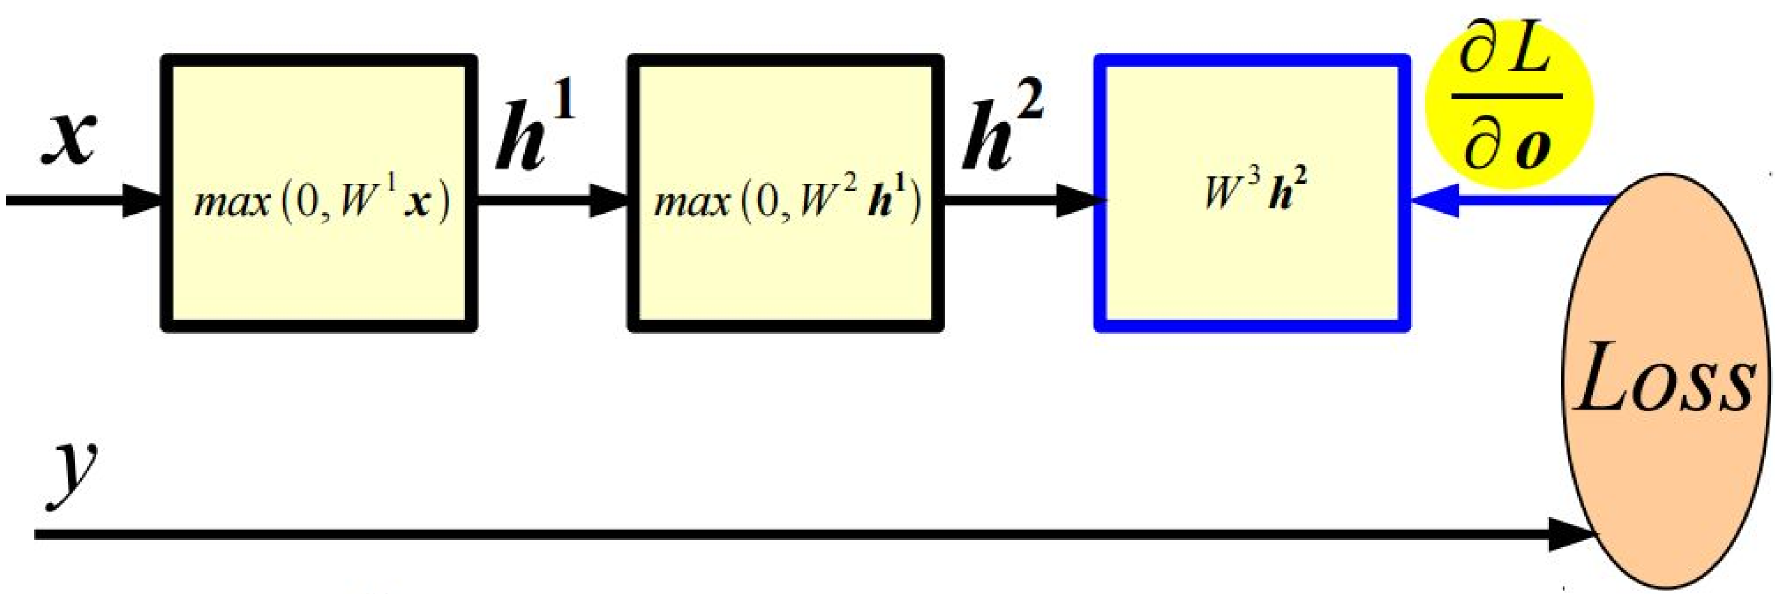

Given $\partial L / \partial \boldsymbol{o}$ and assuming we can easily compute the Jacobian of each module, we have:

$$\frac{\partial L}{\partial W^3} = \frac{\partial L}{\partial \boldsymbol{o}} \frac{\partial \boldsymbol{o}}{\partial W^3} \qquad \frac{\partial L}{\partial \boldsymbol{h}^2} = \frac{\partial L}{\partial \boldsymbol{o}} \frac{\partial \boldsymbol{o}}{\partial \boldsymbol{h}^2}$$

Backward Propagation

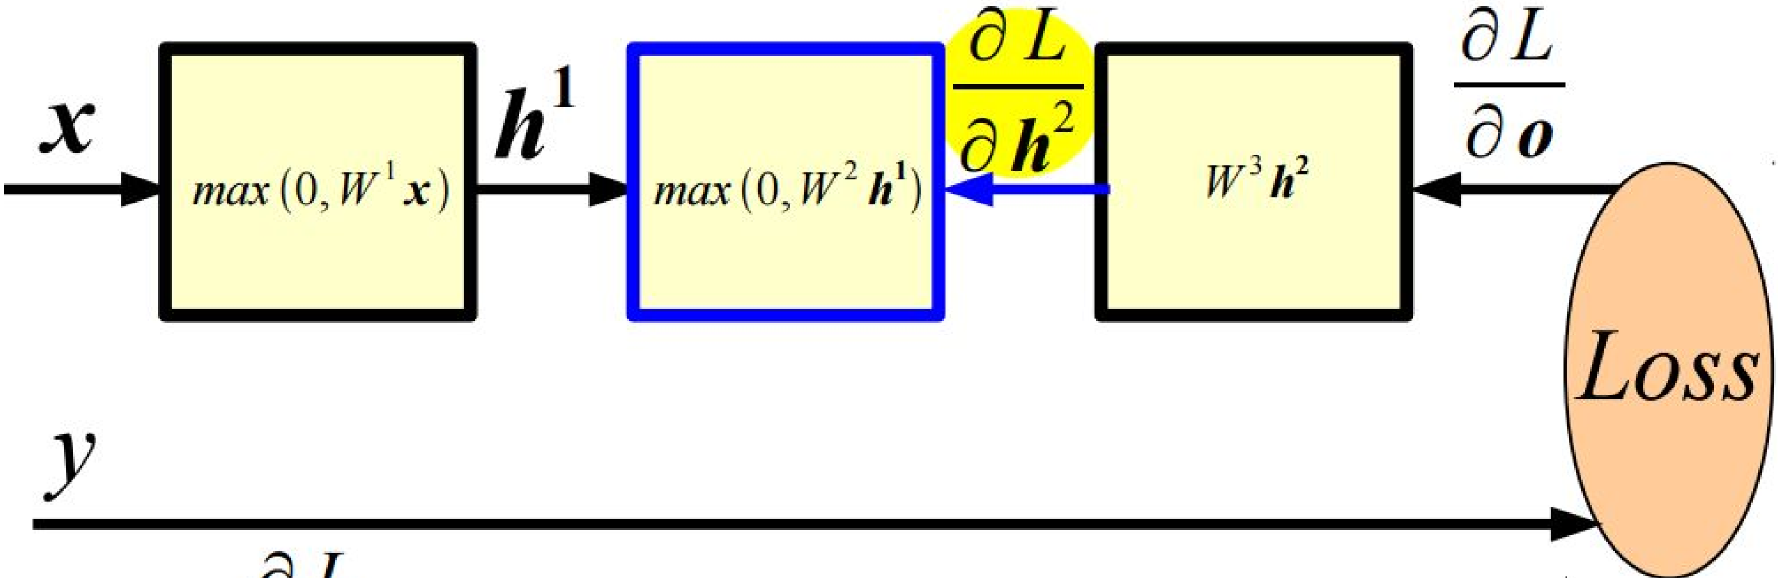

Given $\frac{\partial L}{\partial \boldsymbol{h}^2}$ we can compute now:

$$\frac{\partial L}{\partial W^2} = \frac{\partial L}{\partial \boldsymbol{h}^2} \frac{\partial \boldsymbol{h}^2}{\partial W^2} \qquad \frac{\partial L}{\partial \boldsymbol{h}^1} = \frac{\partial L}{\partial \boldsymbol{h}^2} \frac{\partial \boldsymbol{h}^2}{\partial \boldsymbol{h}^1}$$

Ranzato

Backward Propagation

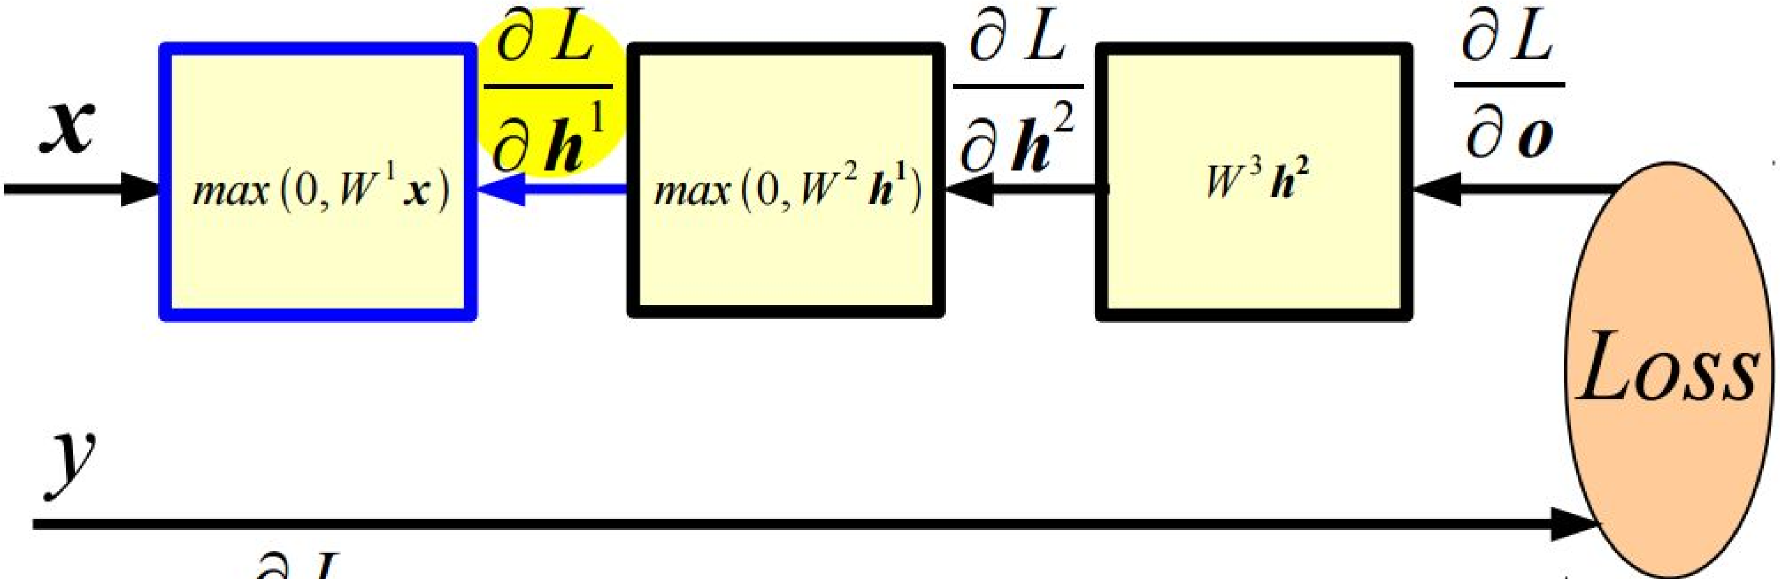

Given $\frac{\partial L}{\partial \boldsymbol{h}^1}$ we can compute now:

$$\frac{\partial L}{\partial W^1} = \frac{\partial L}{\partial \boldsymbol{h}^1} \frac{\partial \boldsymbol{h}^1}{\partial W^1}$$

Convolutional Neural Networks

Fully Connected Layer

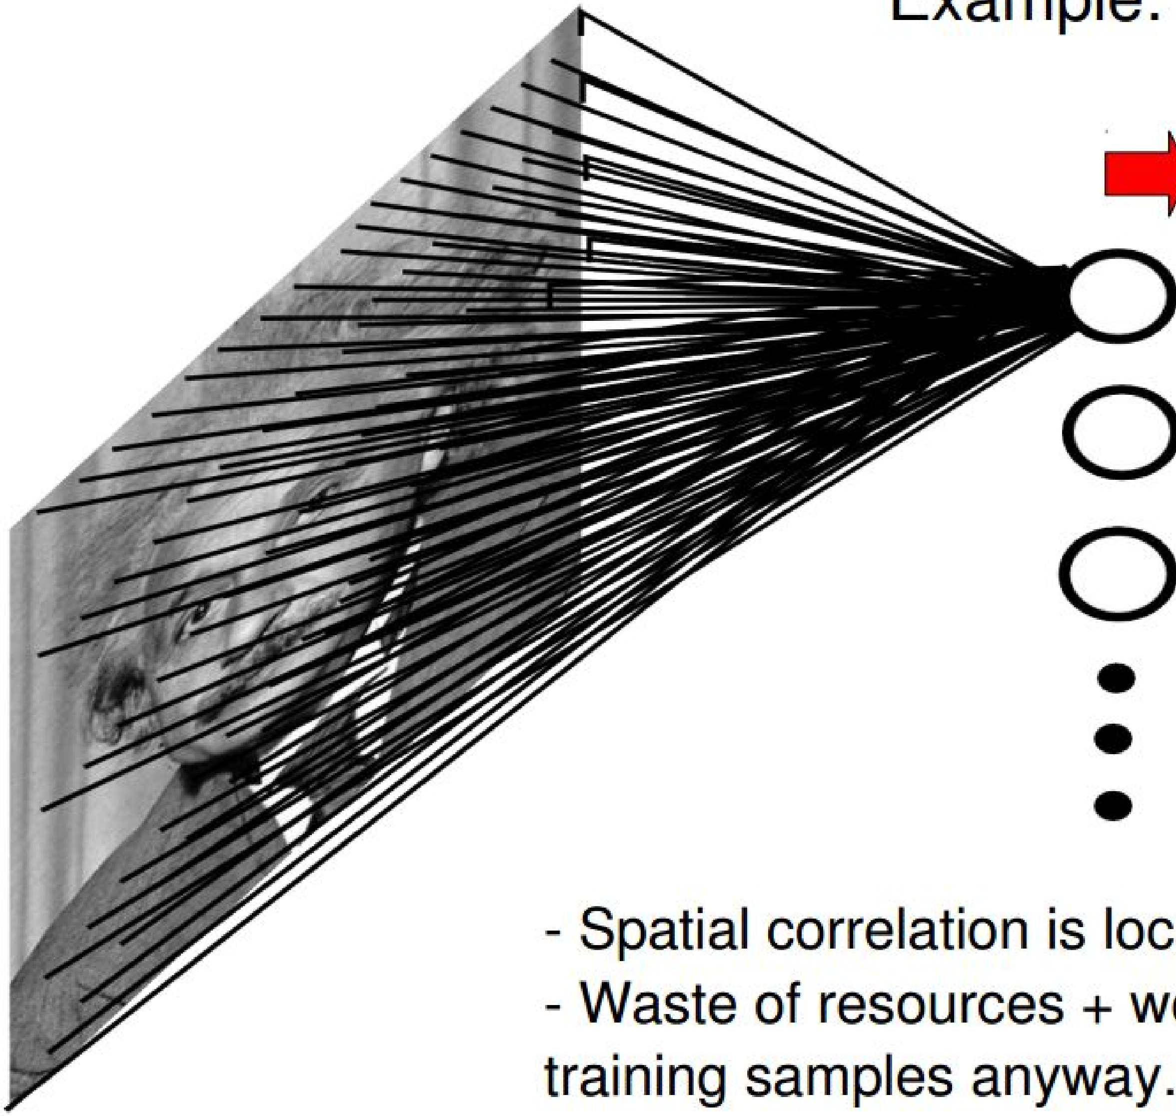

Example: 200x200 image
40K hidden units
~2B parameters!!!

- Spatial correlation is local
- Waste of resources + we have not enough training samples anyway..

Ranzato

Locally Connected Layer

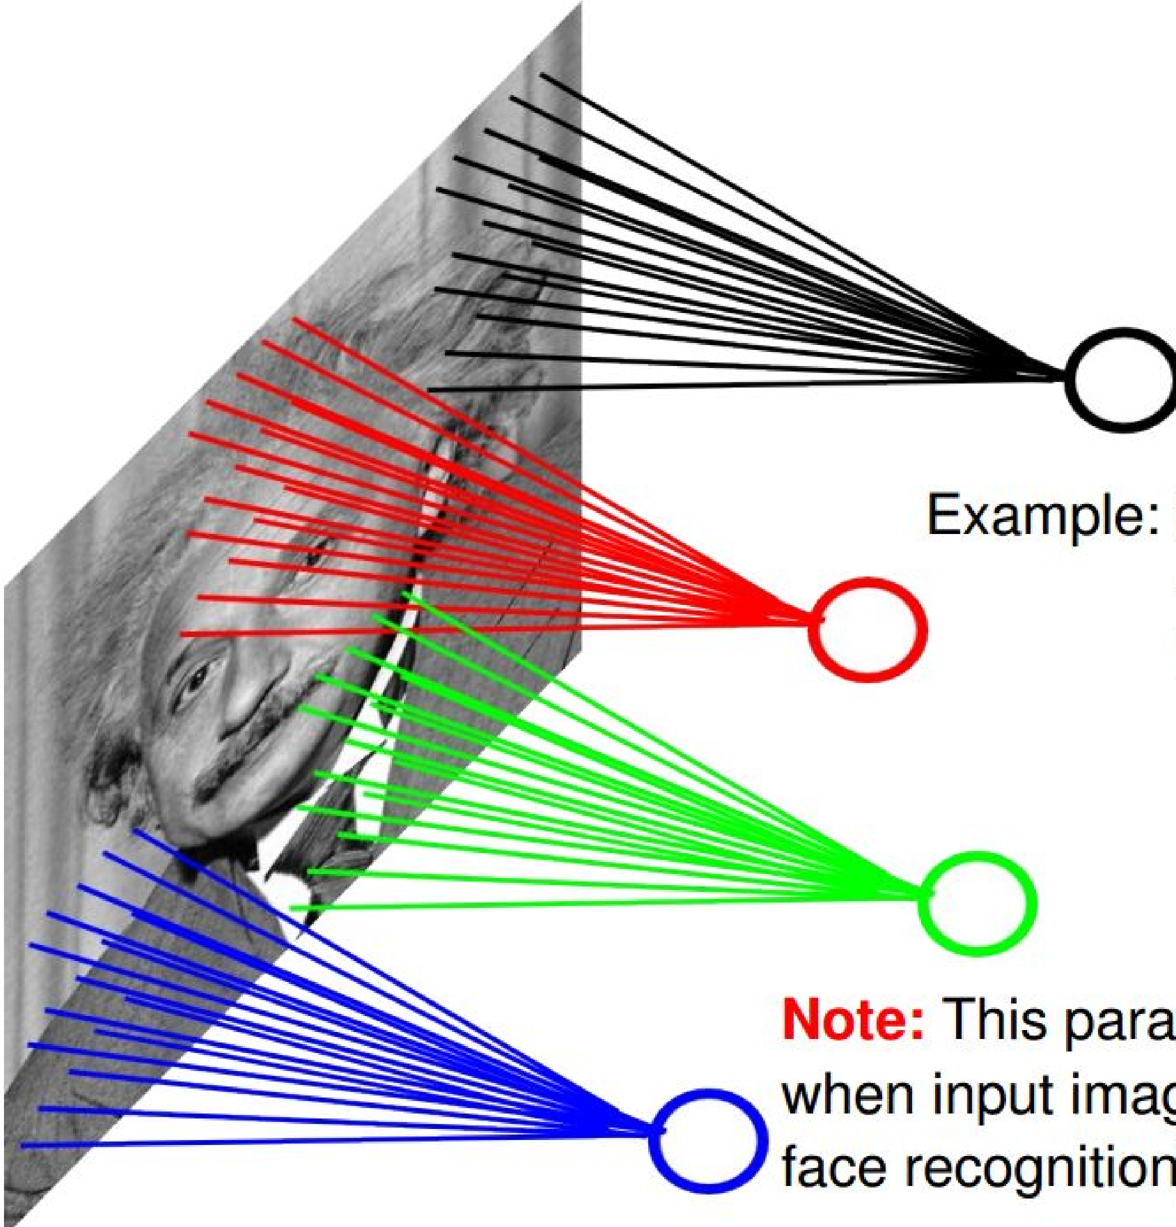

Example: 200x200 image
40K hidden units
Filter size: 10x10
4M parameters

Note: This parameterization is good when input image is registered (e.g., face recognition).

Ranzato

Locally Connected Layer

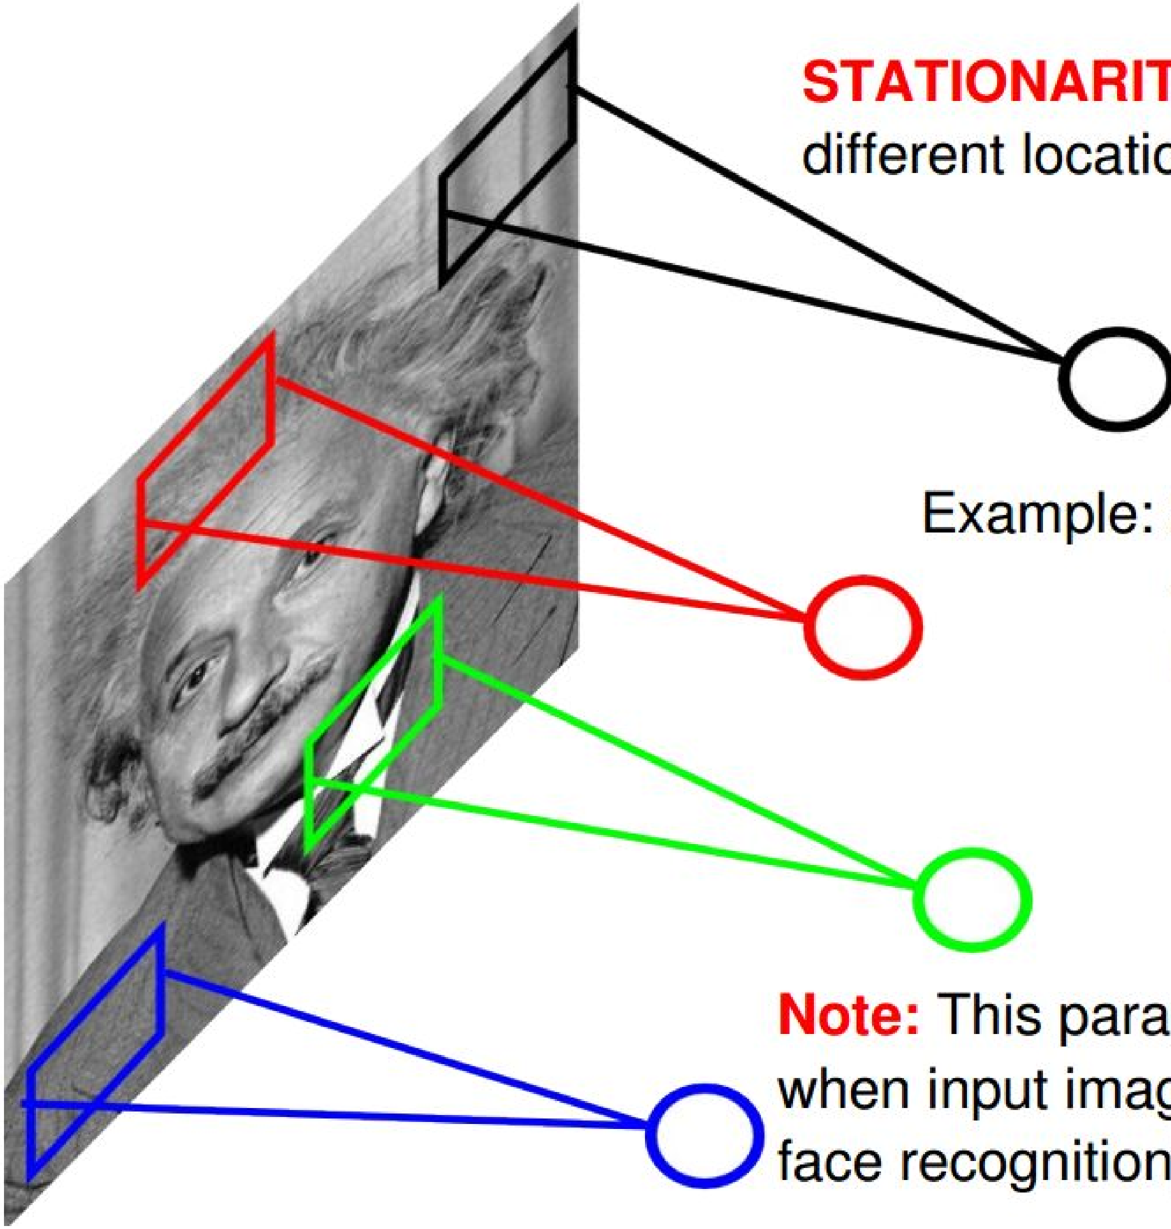

STATIONARITY? Statistics is similar at different locations

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).

Ranzato

Convolutional Layer

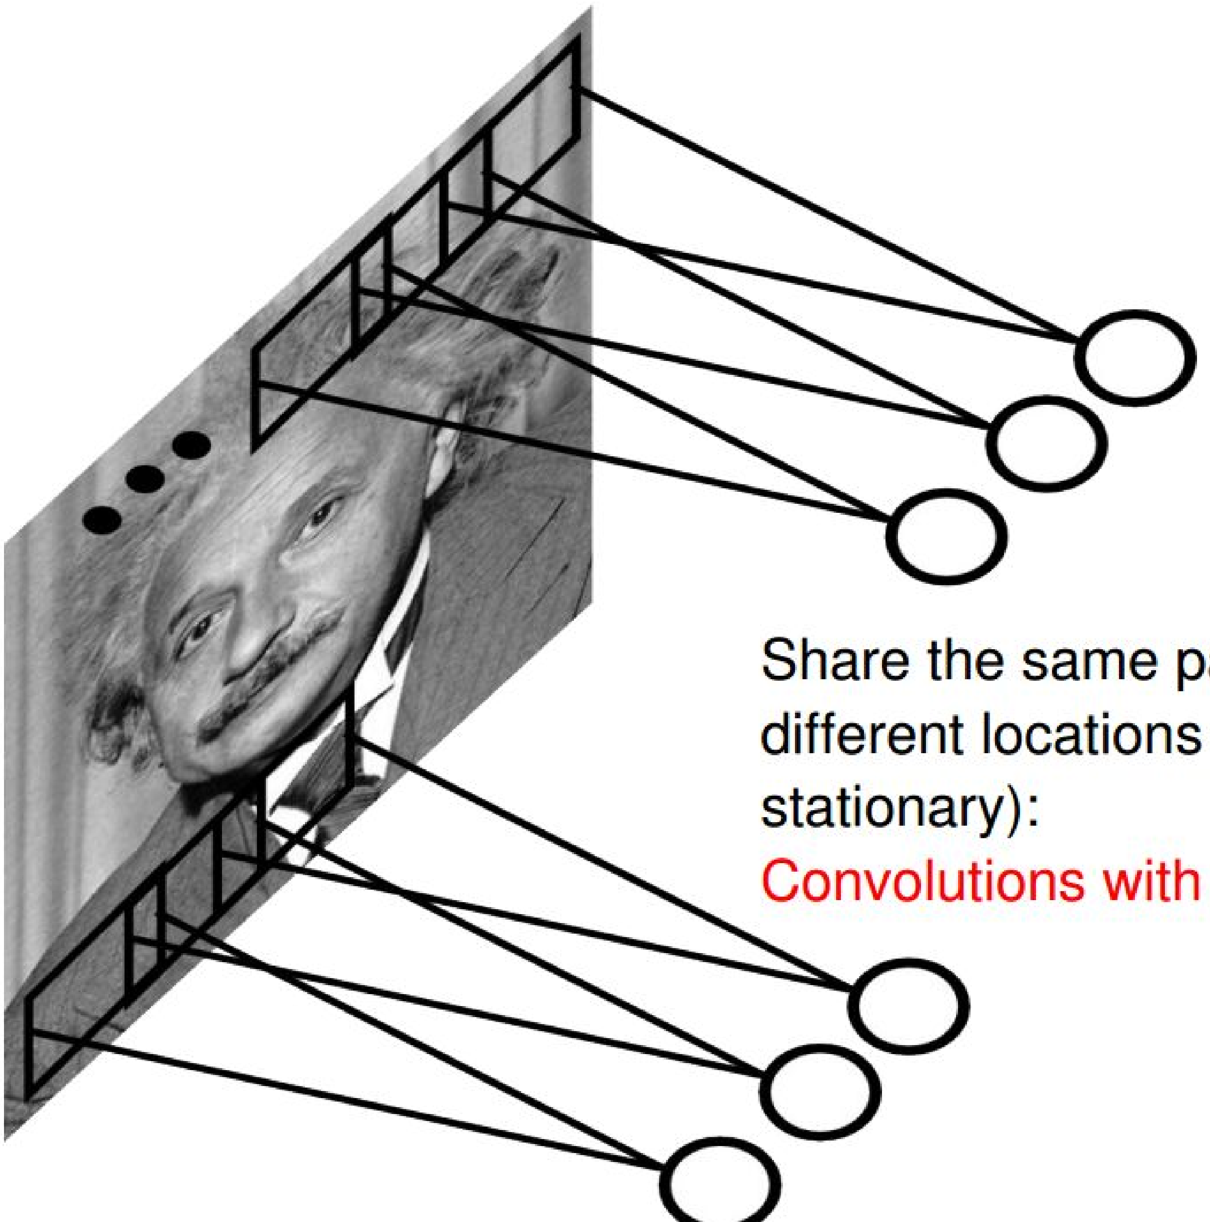

Share the same parameters across different locations (assuming input is stationary):
Convolutions with learned kernels

Convolutional Layer

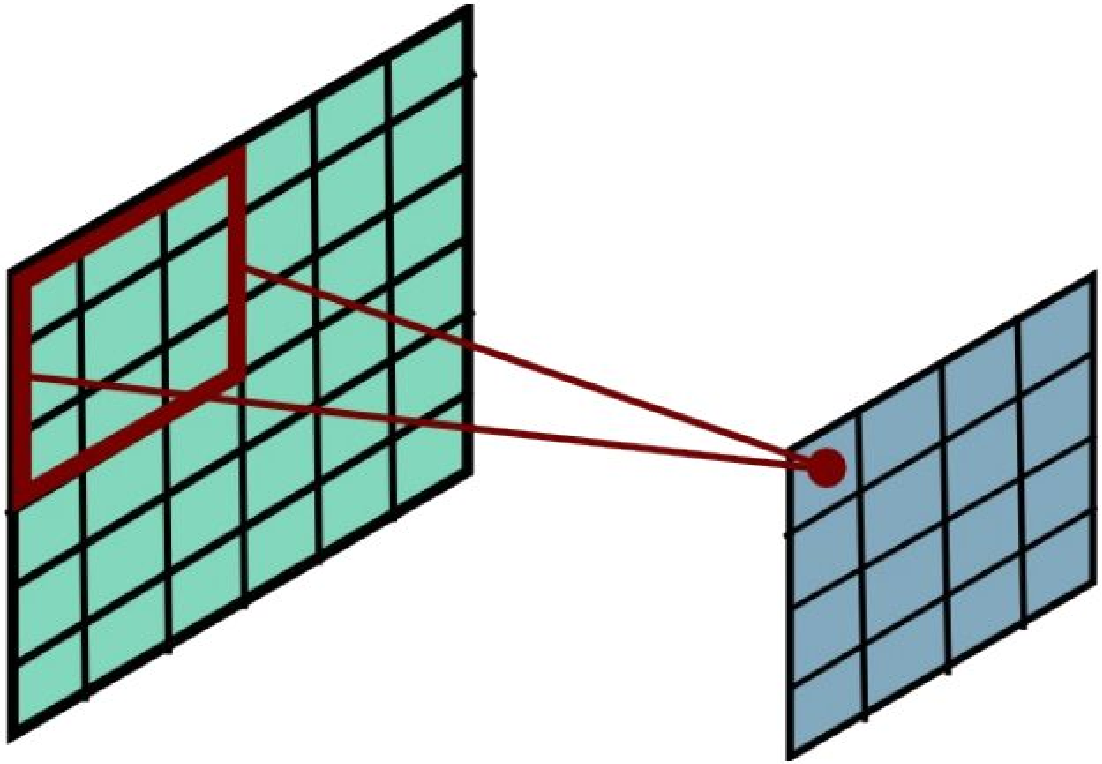

Ranzato

Convolutional Layer

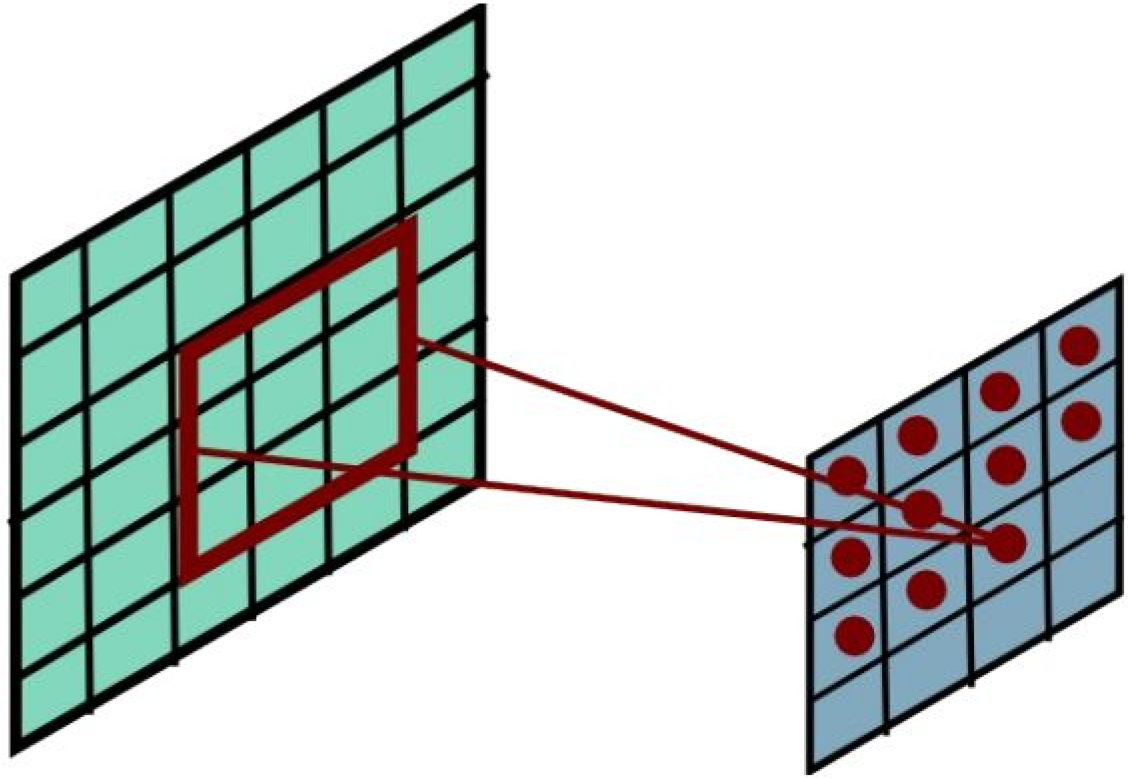

Ranzato

Convolutional Layer

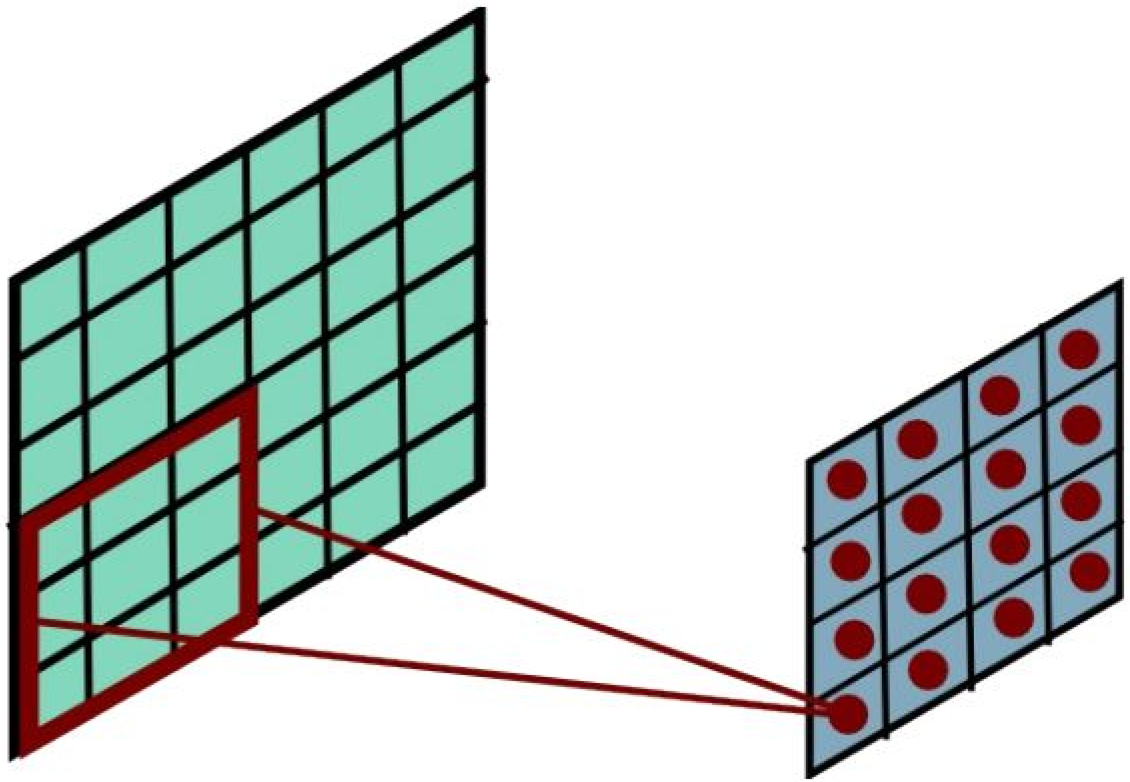

Ranzato

Convolutional Layer

$$* \begin{bmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{bmatrix} =$$

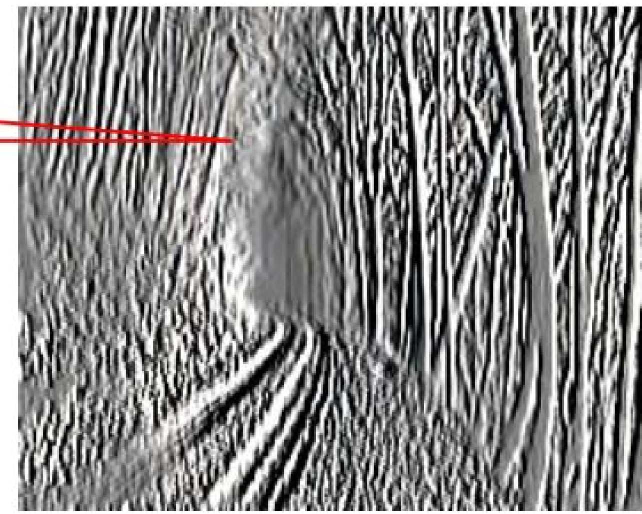

Convolutional Layer

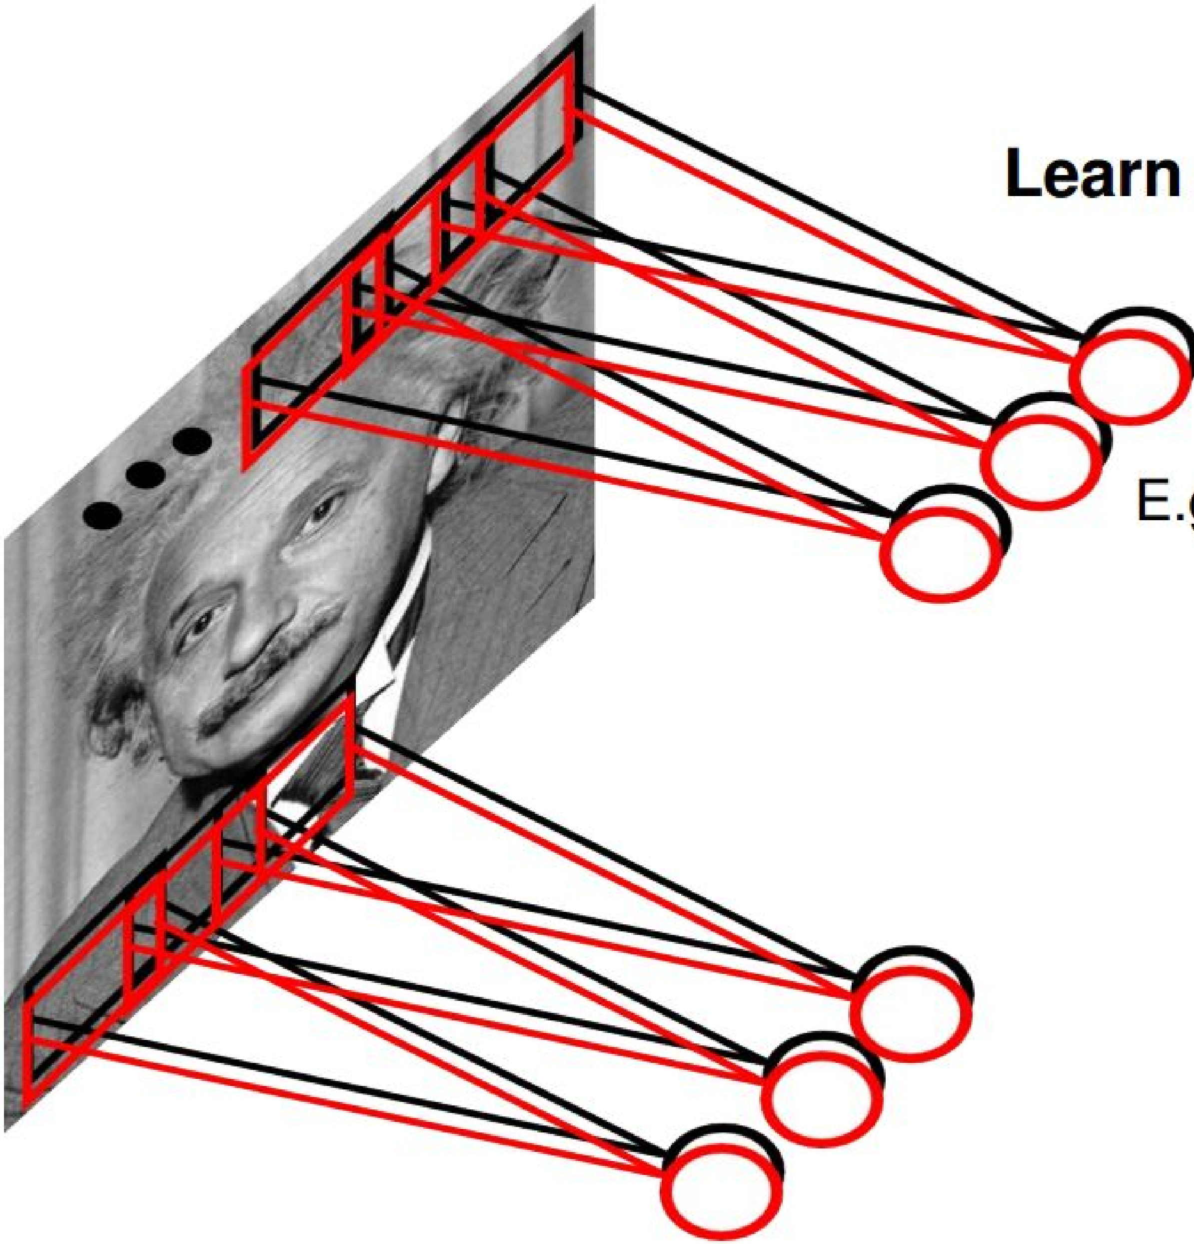

Learn multiple filters.

E.g.: 200x200 image
100 Filters
Filter size: 10x10
10K parameters

Ranzato

Convolutional Layer

$$h_j^n = max\left(0, \sum_{k=1}^{K} h_k^{n-1} * w_{kj}^n\right)$$

output feature map

input feature map

kernel

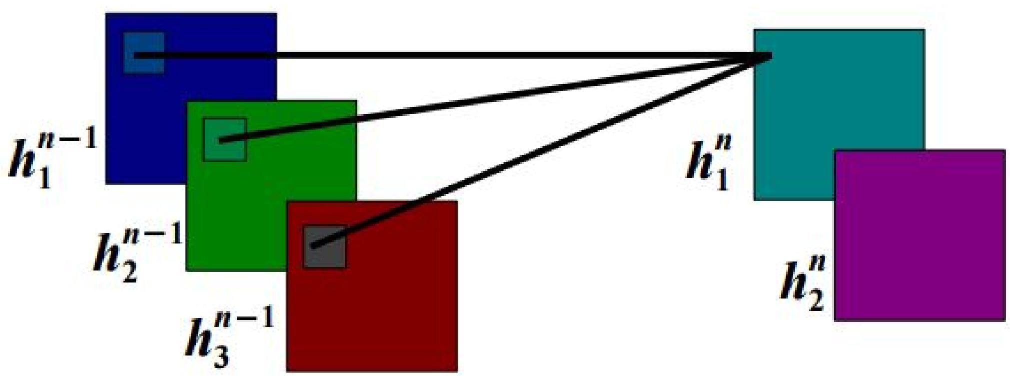

Pooling Layer

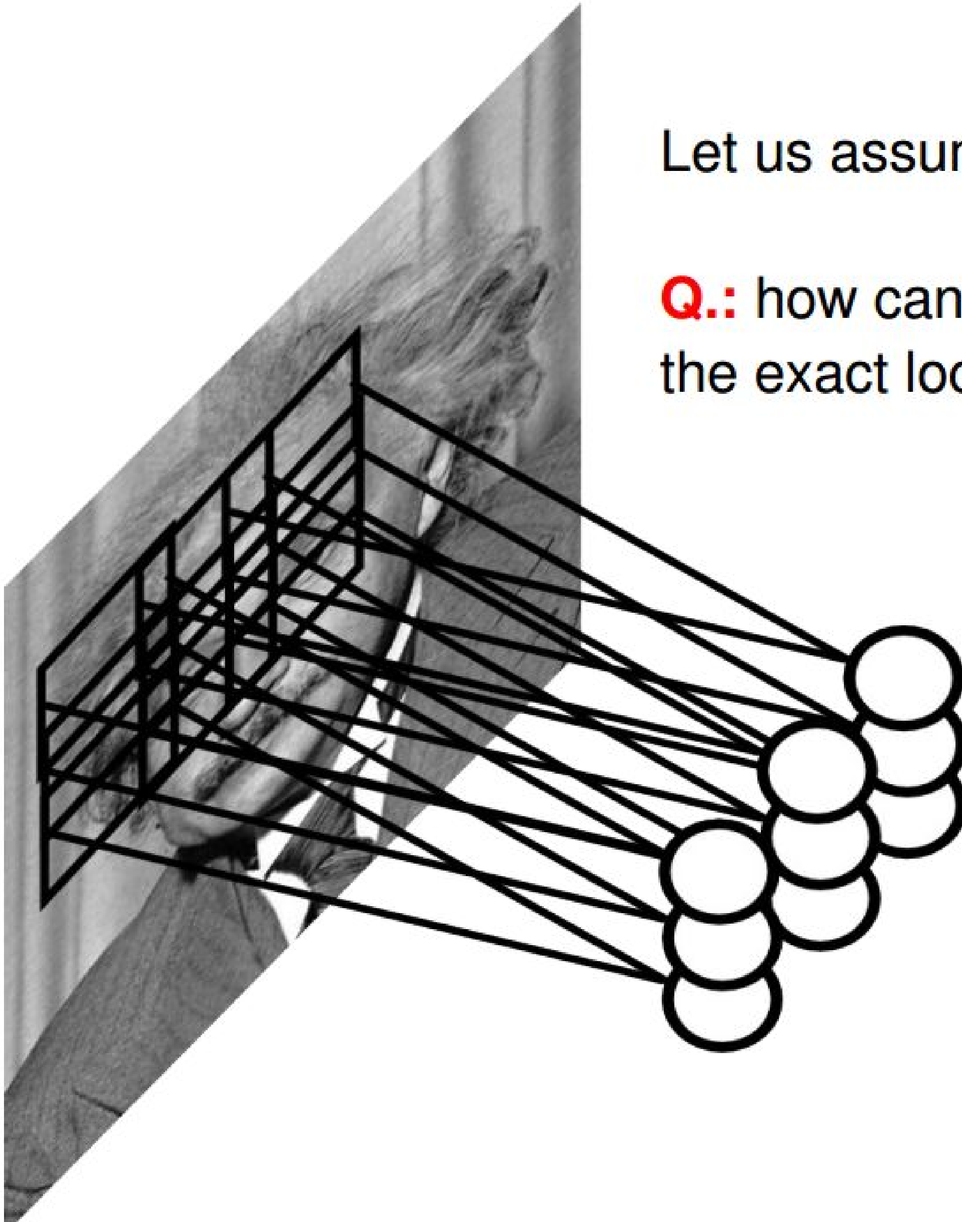

Let us assume filter is an "eye" detector.

Q.: how can we make the detection robust to the exact location of the eye?

Ranzato

Pooling Layer

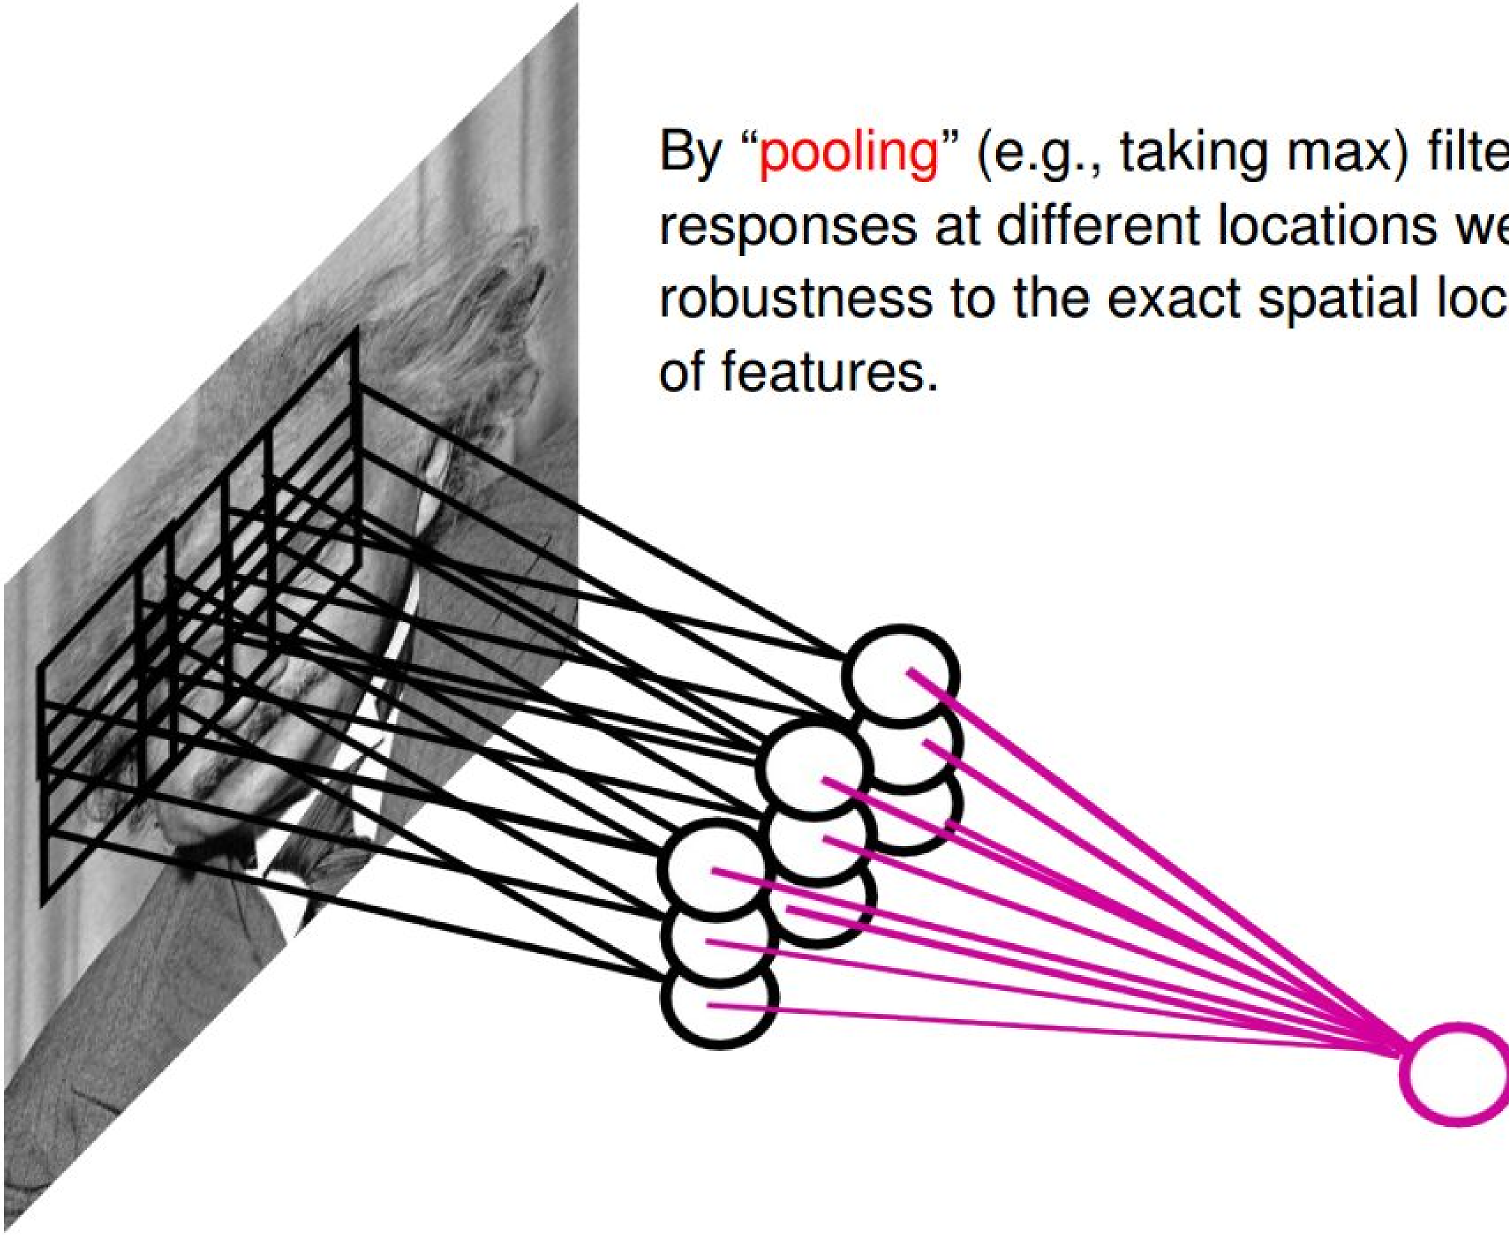

By "pooling" (e.g., taking max) filter responses at different locations we gain robustness to the exact spatial location of features.

Ranzato

Pooling Layer: Examples

Max-pooling:

$$h_j^n(x,y) = max_{\bar{x} \in N(x),\, \bar{y} \in N(y)} h_j^{n-1}(\bar{x},\bar{y})$$

Average-pooling:

$$h_j^n(x,y) = 1/K \sum_{\bar{x} \in N(x),\, \bar{y} \in N(y)} h_j^{n-1}(\bar{x},\bar{y})$$

L2-pooling:

$$h_j^n(x,y) = \sqrt{\sum_{\bar{x} \in N(x),\, \bar{y} \in N(y)} h_j^{n-1}(\bar{x},\bar{y})^2}$$

L2-pooling over features:

$$h_j^n(x,y) = \sqrt{\sum_{k \in N(j)} h_k^{n-1}(x,y)^2}$$

Convolutional Nets

- Example:
  - http://yann.lecun.com/exdb/lenet/index.html

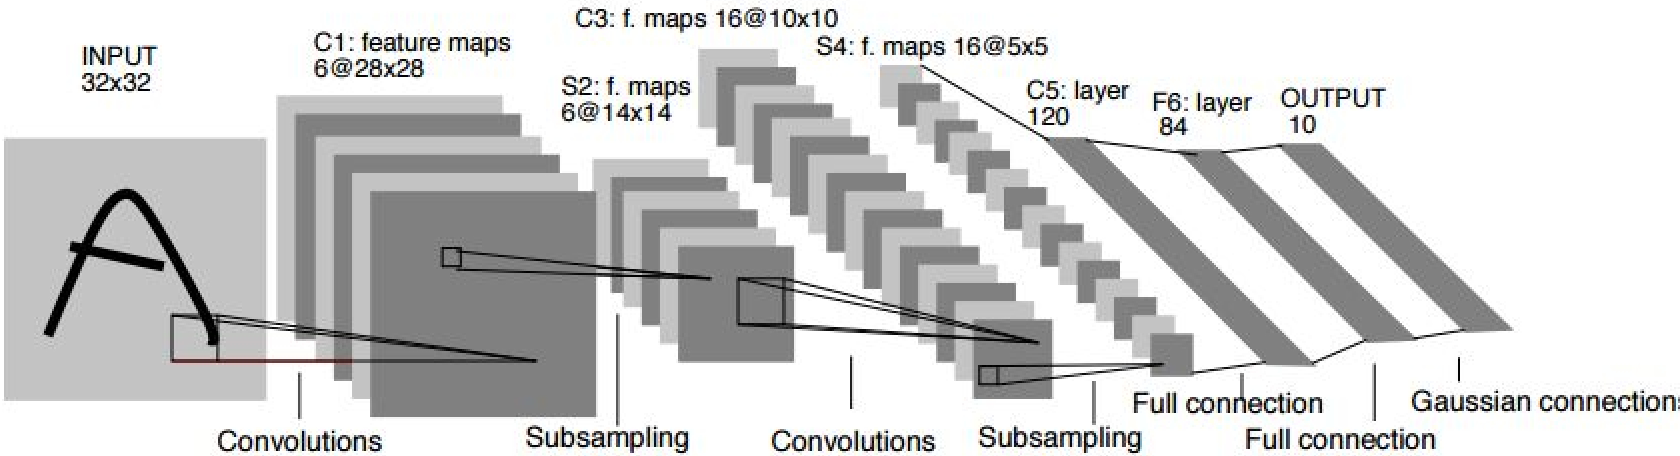

Image Credit: Yann LeCun, Kevin Murphy

Alex Net

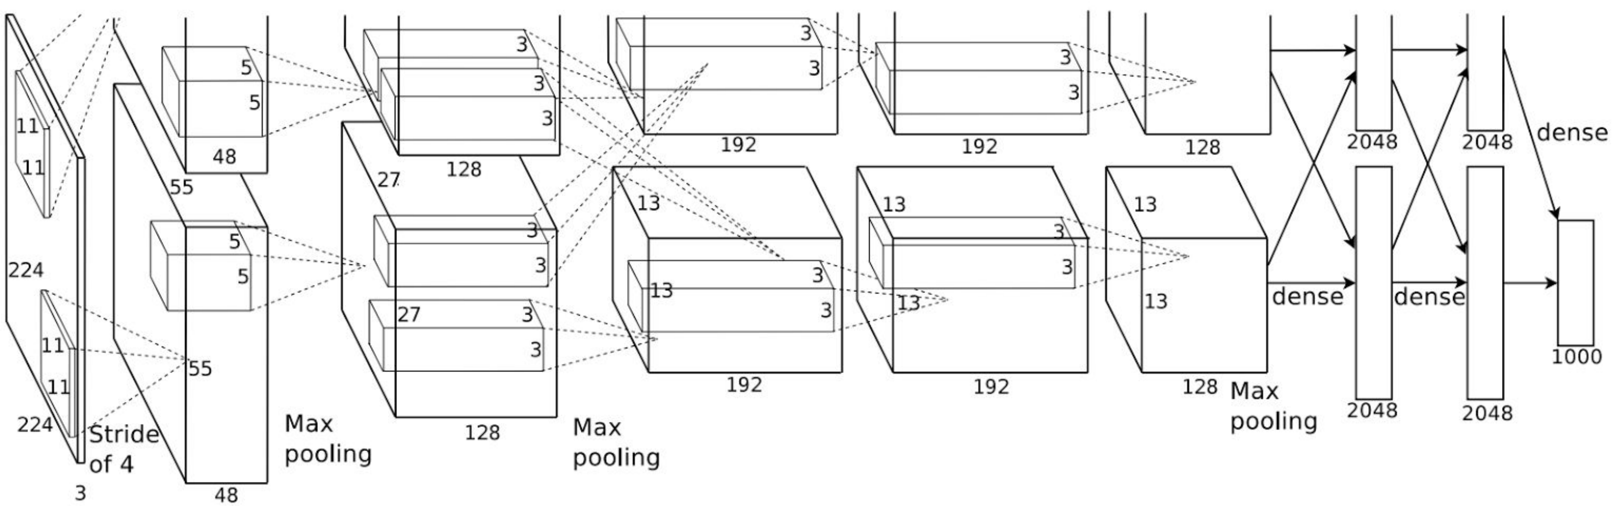

Krizhevsky et al. NIPS 2012

Visualizing Learned Filters

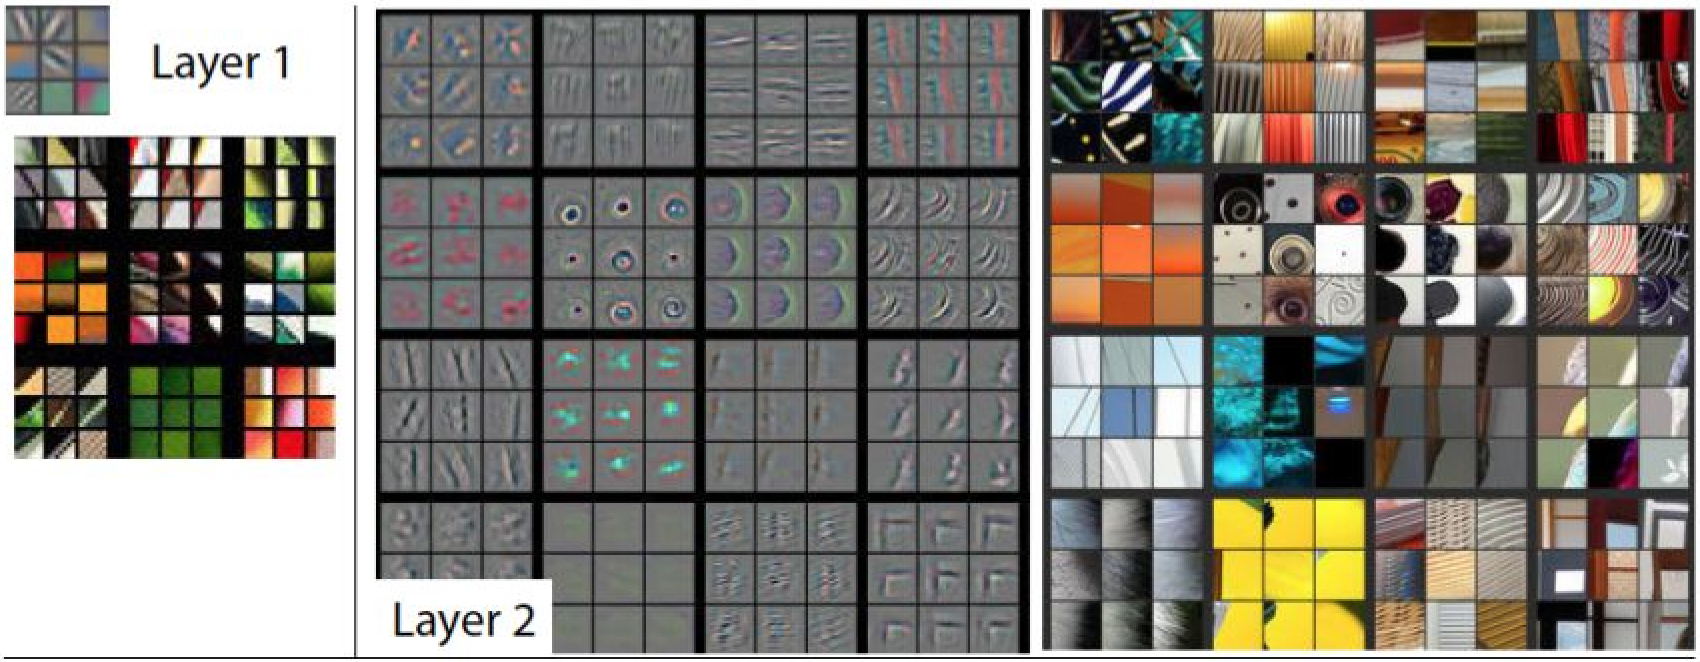

Visualizing Learned Filters

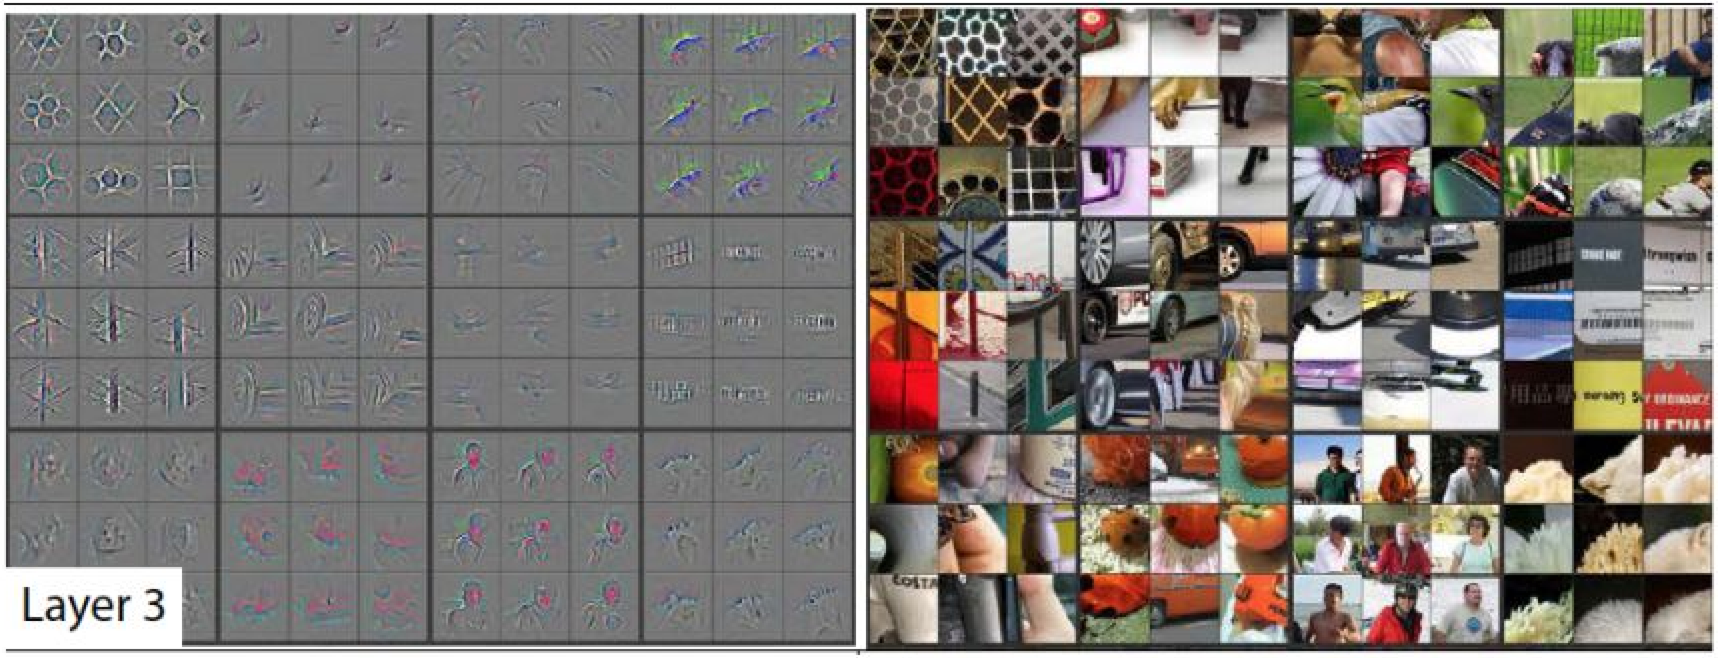

Figure Credit: [Zeiler & Fergus ECCV14]

Visualizing Learned Filters

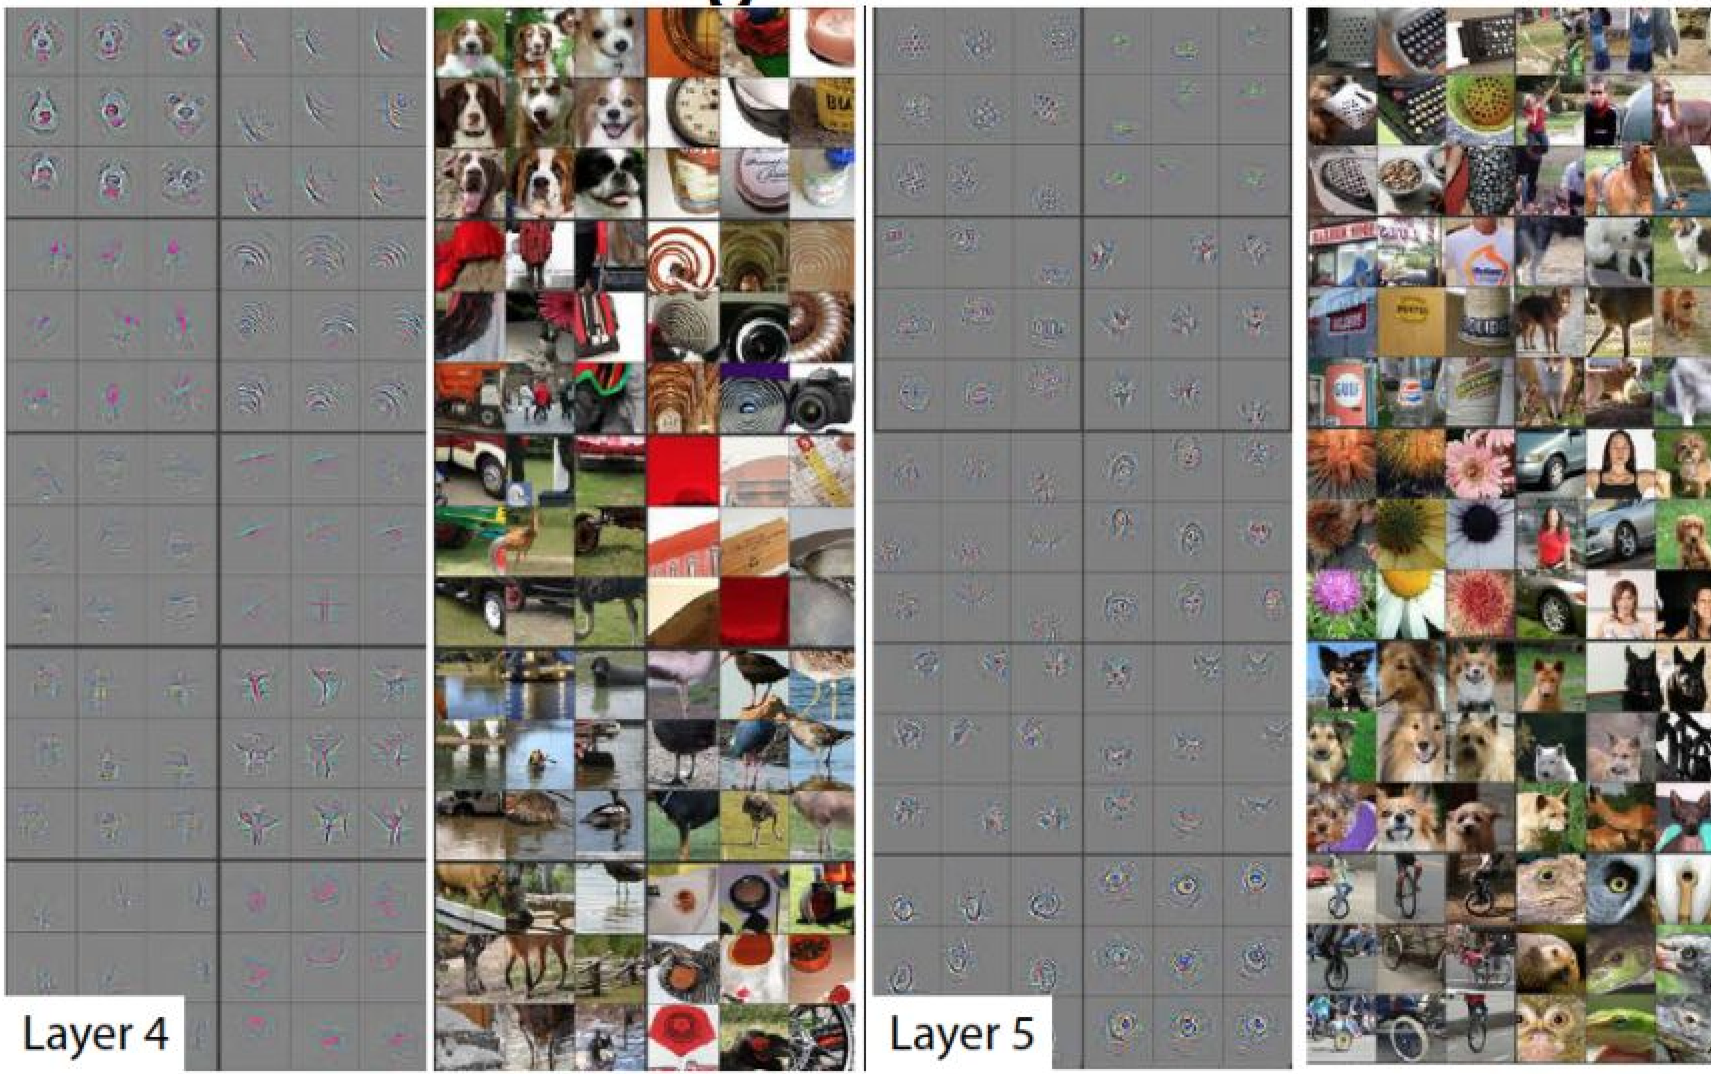

Figure Credit: [Zeiler & Fergus ECCV14]

Industry Deployment

- Used in Facebook, Google, Microsoft
- Image Recognition, Speech Recognition, ….
- Fast at test time

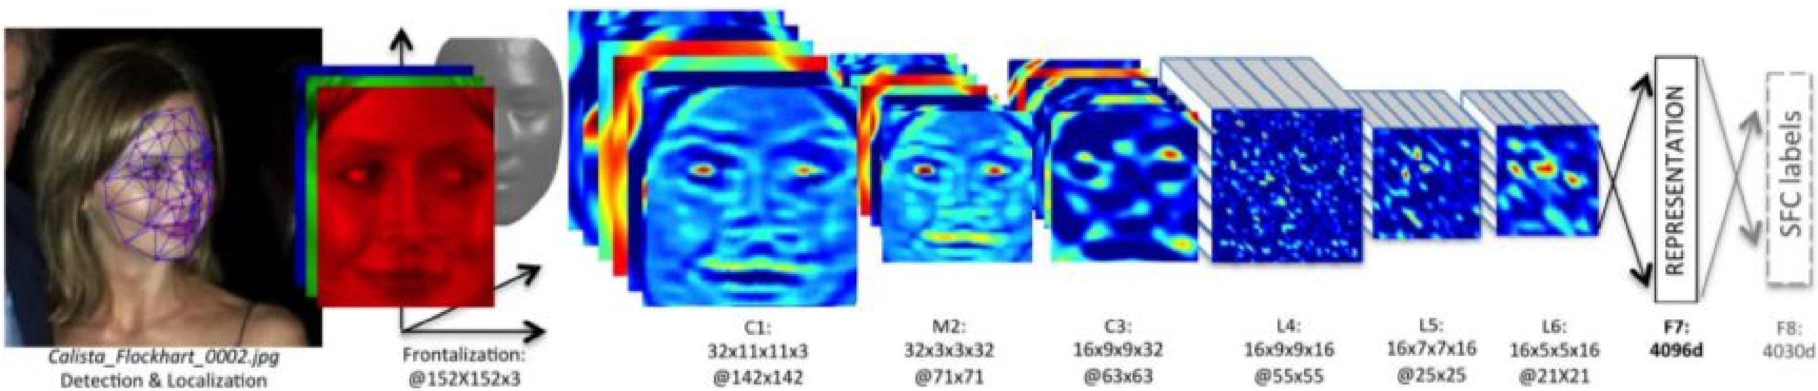

Taigman et al. DeepFace: Closing the Gap to Human-Level Performance in Face Verification, CVPR'14

Slide: R. Fergus