

Outline

Q: For classification, what do you call a neural net with no hidden layers?

1 Back-Propagation

$x \rightarrow \text{linear} \rightarrow \text{sigmoid/softmax} \rightarrow y$

logistic regression

neural net

x_1

↓

linear activation

2 Convolutional Networks

new problem

- try linear / logistic regression approach

- training error high
→ try more complicated model

Announcements

midterm

project

lecture during tutorial

Wednesday 7pm (virtual)

Goal today

$f(x; \Theta)$
neural network \leftarrow
 w_1, b_1
 w_2, b_2

and so on

1 Back-Propagation

2 Convolutional Networks

compute

$\nabla w_1, \nabla b_1$

and so on

Learning Weights in a Neural Network

- Goal is to learn weights in a multi-layer neural network using gradient descent.
- Weight space for a multi-layer neural net: one set of weights for each unit in every layer of the network
- Define a loss \mathcal{L} and compute the gradient of the cost $d\mathcal{J}/d\mathbf{w}$, the average loss over all the training examples.
- Let's look at how we can calculate $d\mathcal{L}/d\mathbf{w}$.

Example: Two-Layer Neural Network




Figure: Two-Layer Neural Network

Computations for Two-Layer Neural Network

A neural network computes a composition of functions.

$$z_1^{(1)} = w_{01}^{(1)} \cdot 1 + w_{11}^{(1)} \cdot x_1 + w_{21}^{(1)} \cdot x_2$$

$$h_1 = \sigma(z_1)$$

$$z_1^{(2)} = w_{01}^{(2)} \cdot 1 + w_{11}^{(2)} \cdot h_1 + w_{21}^{(2)} \cdot h_2$$

$$y_1 = z_1$$

$$z_2^{(1)} =$$

$$h_2 =$$

$$z_2^{(2)} =$$

$$y_2 =$$

$$L = \frac{1}{2} \left((y_1 - t_1)^2 + (y_2 - t_2)^2 \right)$$

Simplified Example: Logistic Least Squares


$$z = wx + b$$

$$y = \sigma(z)$$

sigmoid


$$\mathcal{L} = \frac{1}{2}(y - t)^2$$

loss




Computation Graph

- The nodes represent the inputs and computed quantities.
- The edges represent which nodes are computed directly as a function of which other nodes.



Uni-variate Chain Rule

Let $z = f(y)$ and $y = g(x)$ be uni-variate functions.
Then $z = f(g(x))$.

$$\frac{dz}{dx} = \frac{dz}{dy} \frac{dy}{dx}$$


Univariate Chain Rule

How you would have done it in calculus class

$$\begin{aligned}\mathcal{L} &= \frac{1}{2}(\sigma(wx + b) - t)^2 \\ \frac{\partial \mathcal{L}}{\partial w} &= \frac{\partial}{\partial w} \left[\frac{1}{2}(\sigma(wx + b) - t)^2 \right] \\ &= \frac{1}{2} \frac{\partial}{\partial w} (\sigma(wx + b) - t)^2 \\ &= (\sigma(wx + b) - t) \frac{\partial}{\partial w} (\sigma(wx + b) - t) \\ &= (\sigma(wx + b) - t) \sigma'(wx + b) \frac{\partial}{\partial w} (wx + b) \\ &= (\sigma(wx + b) - t) \sigma'(wx + b) x\end{aligned}$$

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial b} &= \frac{\partial}{\partial b} \left[\frac{1}{2}(\sigma(wx + b) - t)^2 \right] \\ &= \frac{1}{2} \frac{\partial}{\partial b} (\sigma(wx + b) - t)^2 \\ &= (\sigma(wx + b) - t) \frac{\partial}{\partial b} (\sigma(wx + b) - t) \\ &= (\sigma(wx + b) - t) \sigma'(wx + b) \frac{\partial}{\partial b} (wx + b) \\ &= (\sigma(wx + b) - t) \sigma'(wx + b)\end{aligned}$$

What are the disadvantages of this approach?

computationally expensive (repeats terms)

hard to understand / easy to make mistakes

Logistic Least Squares: Gradient for w

Computing the gradient for w :

$$\frac{\partial \mathcal{L}}{\partial w} = \frac{\partial \mathcal{L}}{\partial y} \frac{\partial y}{\partial w}$$

$$= \frac{\partial \mathcal{L}}{\partial y} \frac{\partial y}{\partial z} \frac{\partial z}{\partial w}$$

$$= (y - t) \sigma'(z) x$$

$$= (\sigma(wx + b) - t) \sigma'(wx + b) x$$


$$\frac{\partial \mathcal{L}}{\partial y} \frac{\partial y}{\partial z} \frac{\partial z}{\partial b} \quad \Delta$$

Computing the loss:

$$z = wx + b$$

$$y = \sigma(z)$$

$$\mathcal{L} = \frac{1}{2} (y - t)^2$$



$$\frac{\partial \mathcal{L}}{\partial y} = y - t$$

$$\frac{\partial y}{\partial z} = \sigma'(z)$$

$$= \sigma(z) (1 - \sigma(z))$$

$$\frac{\partial z}{\partial w} = x$$

Logistic Least Squares: Gradient for b

Computing the gradient for b :

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial b} &= \\ &= \\ &= \\ &= \end{aligned}$$

Computing the loss:

$$z = wx + b$$

$$y = \sigma(z)$$

$$\mathcal{L} = \frac{1}{2}(y - t)^2$$

Logistic Least Squares: Gradient for b

Computing the gradient for b :

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial b} &= \frac{\partial \mathcal{L}}{\partial y} \frac{\partial y}{\partial b} \\ &= \frac{\partial \mathcal{L}}{\partial y} \frac{\partial y}{\partial z} \frac{\partial z}{\partial b} \\ &= (y - t) \sigma'(z) 1 \\ &= (\sigma(wx + b) - t) \sigma'(wx + b) 1\end{aligned}$$

Computing the loss:

$$z = wx + b$$

$$y = \sigma(z)$$

$$\mathcal{L} = \frac{1}{2} (y - t)^2$$

Comparing Gradient Computations for w and b

Computing the gradient for w : Computing the gradient for b :

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial w} &= \frac{\partial \mathcal{L}}{\partial y} \frac{\partial y}{\partial z} \frac{\partial z}{\partial w} \\ &= (y - t) \sigma'(z) x\end{aligned}$$

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial b} &= \frac{\partial \mathcal{L}}{\partial y} \frac{\partial y}{\partial z} \frac{\partial z}{\partial b} \\ &= (y - t) \sigma'(z) 1\end{aligned}$$

Computing the loss:

$$z = wx + b$$

$$y = \sigma(z)$$

$$\mathcal{L} = \frac{1}{2}(y - t)^2$$

Structured Way of Computing Gradients

Computing the gradients:

$$\frac{\partial \mathcal{L}}{\partial y} = (y - t)$$

$$\frac{\partial \mathcal{L}}{\partial z} = \frac{\partial \mathcal{L}}{\partial y} \sigma'(z)$$

$$\frac{\partial \mathcal{L}}{\partial w} = \frac{d\mathcal{L}}{dz} \frac{dz}{dw} = \frac{d\mathcal{L}}{dz} x$$

$$\frac{\partial \mathcal{L}}{\partial b} = \frac{d\mathcal{L}}{dz} \frac{dz}{db} = \frac{d\mathcal{L}}{dz} 1$$

Computing the loss:

$$z = wx + b$$

$$y = \sigma(z)$$

$$\mathcal{L} = \frac{1}{2}(y - t)^2$$

Error Signal Notation

$\nearrow \frac{\partial \mathcal{L}}{\partial y}$ (numeric values) at the current ^{value} of the weights and given input data

- Let \bar{y} denote the derivative $d\mathcal{L}/dy$, called the **error signal**.
- Error signals are just values our program is computing (rather than a mathematical operation).

Computing the loss:

$$z = wx + b$$

$$y = \sigma(z)$$

$$\mathcal{L} = \frac{1}{2}(y - t)^2$$

one single data point

$$w_{t+1} = w_t - \alpha \bar{w}_t$$

update using GD


Computing the derivatives:

$$\bar{L} = 1$$

$$\bar{y} = (y - t) \leftarrow \frac{\partial \mathcal{L}}{\partial y} = \frac{\partial \mathcal{L}}{\partial L} \frac{\partial L}{\partial y} = (y - t)$$

$$\bar{z} = \bar{y} \sigma'(z)$$

$$\bar{w} = \bar{z} x \quad \bar{b} = \bar{z}$$




$$\frac{\partial \mathcal{L}}{\partial z} = \bar{y} \cdot \frac{\partial y}{\partial z} = \bar{y} \sigma'(z)$$

Computation Graph has a Fan-Out > 1

$$\bar{w} = \bar{z} \cdot \frac{\partial z}{\partial w}$$
$$= \bar{z} \cdot x$$


L_2 -Regularized Regression



$$z = wx + b$$
$$y = \sigma(z)$$
$$\mathcal{L} = \frac{1}{2}(y - t)^2$$
$$\mathcal{R} = \frac{1}{2}w^2$$
$$\mathcal{L}_{\text{reg}} = \mathcal{L} + \lambda\mathcal{R}$$

Computation Graph has a Fan-Out > 1

Softmax Regression



$$z_\ell = \sum_j w_{\ell j} x_j + b_\ell$$


$$y_k = \frac{e^{z_k}}{\sum_\ell e^{z_\ell}}$$

$$\mathcal{L} = - \sum_k t_k \log y_k$$

Multi-variate Chain Rule

Suppose we have functions $f(x, y)$, $x(t)$, and $y(t)$.

$$\frac{d}{dt} f(x(t), y(t)) = \frac{\partial f}{\partial x} \frac{dx}{dt} + \frac{\partial f}{\partial y} \frac{dy}{dt}$$



$$\frac{d}{dx} e^{3x} = 3e^{3x}$$

Example:

$$\frac{\partial f}{\partial t} = \underbrace{\frac{\partial f}{\partial x}}_x + \underbrace{\frac{\partial f}{\partial y}}_y$$

contribution through x

$$f(x, y) = y + e^{xy}$$

$$x(t) = \cos t$$

$$y(t) = t^2$$

$$\frac{df}{dt} = \frac{\partial f}{\partial x} \frac{dx}{dt} + \frac{\partial f}{\partial y} \frac{dy}{dt}$$

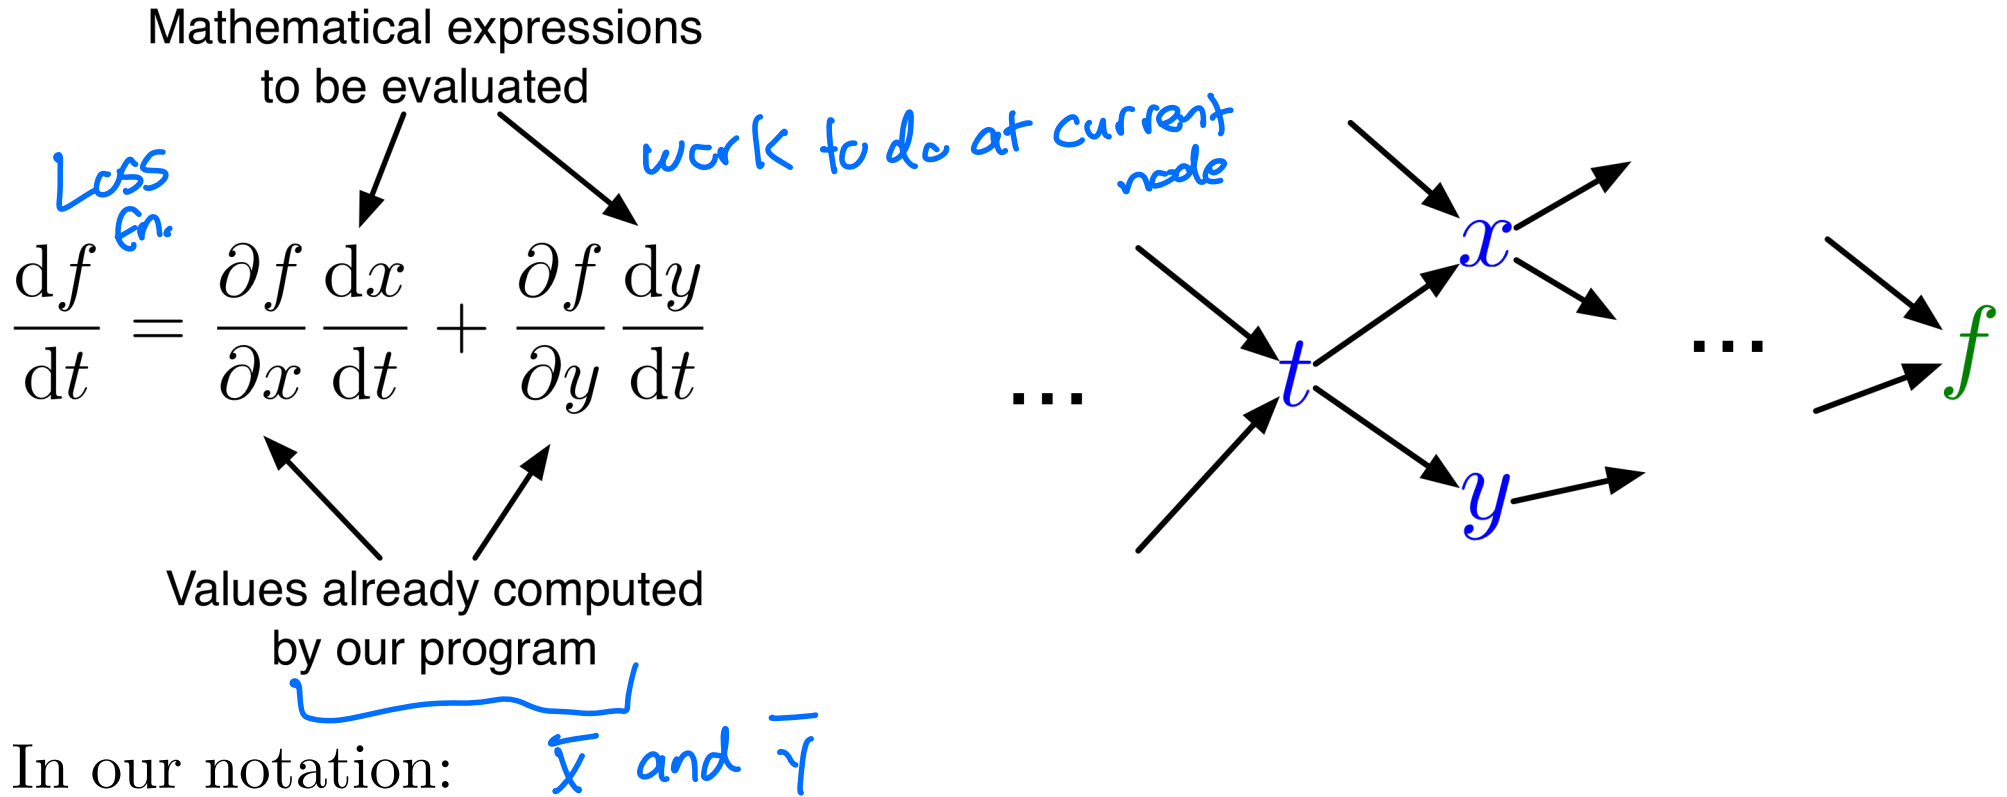
$$= (ye^{xy}) \cdot (-\sin t) + (1 + xe^{xy}) \cdot 2t$$

$$(-\sin t)$$

$$ye^{xy} \cdot (-\sin t)$$

Multi-variate Chain Rule

In the context of back-propagation:



$$\bar{t} = \bar{x} \frac{dx}{dt} + \bar{y} \frac{dy}{dt}$$

Full Backpropagation Algorithm:

Let v_1, \dots, v_N be a **topological ordering** of the computation graph (i.e. parents come before children.)

v_N denotes the variable for which we're trying to compute gradients.

- forward pass:

Compute the composition of functions

For $i = 1, \dots, N$,


Compute v_i as a function of Parents(v_i).

- backward pass:

For $i = N - 1, \dots, 1$,

$$\bar{v}_i = \sum_{j \in \text{Children}(v_i)} \bar{v}_j \frac{\partial v_j}{\partial v_i}$$

$\frac{\partial L}{\partial v_i}$



Forward pass:

$$z = wx + b$$

$$y = \sigma(z)$$

$$\mathcal{L} = \frac{1}{2}(y - t)^2$$

$$\mathcal{R} = \frac{1}{2}w^2$$

$$\mathcal{L}_{\text{reg}} = \mathcal{L} + \lambda \mathcal{R}$$



SGD

sample $(x, y) \sim D$

compute $L = f(x, y, w_t)$

Get $\bar{w}_t = \frac{\partial f}{\partial w_t}$ using backprop
for each layer

$$w_{t+1} = w_t - \alpha \bar{w}_t$$

$$b_{t+1} = b_t - \alpha \bar{b}_t$$

initialize $w_0^{(1)}, b_0^{(1)}$

$w_0^{(2)}, b_0^{(2)}$

small random numbers

$$\bar{\mathcal{L}}_{\text{reg}} = 1$$

$$\bar{\mathcal{L}} = \bar{\mathcal{L}}_{\text{reg}} \frac{\partial \mathcal{L}_{\text{reg}}}{\partial \mathcal{L}} = \bar{\mathcal{L}}_{\text{reg}}$$

$$\bar{\mathcal{R}} = \bar{\mathcal{L}}_{\text{reg}} \frac{\partial \mathcal{L}_{\text{reg}}}{\partial \mathcal{R}} = \bar{\mathcal{L}}_{\text{reg}} \lambda$$

\bar{t} is unnecessary

$$\bar{y} = \bar{\mathcal{L}} \frac{\partial \mathcal{L}}{\partial y} = \bar{\mathcal{L}} (y - t)$$

compute \bar{b} \bar{w}

$$\bar{z} = \bar{y} \frac{\partial y}{\partial z} = \bar{y} \sigma'(z)$$


\downarrow
 $\sigma(z)(1 - \sigma(z))$

$$= \bar{\mathcal{R}} \frac{\partial \mathcal{R}}{\partial w} + \bar{z} \frac{\partial z}{\partial w}$$

$$= \bar{\mathcal{R}} w + \bar{z} x$$

$$\bar{b} = \bar{z} \frac{\partial z}{\partial b} = \bar{z} 1 = \bar{z}$$

Backpropagation for Regularized Logistic Least Squares



Forward pass:

$$z = wx + b$$

$$y = \sigma(z)$$

$$\mathcal{L} = \frac{1}{2}(y - t)^2$$

$$\mathcal{R} = \frac{1}{2}w^2$$

$$\mathcal{L}_{\text{reg}} = \mathcal{L} + \lambda\mathcal{R}$$

Backward pass:

$$\overline{\mathcal{L}_{\text{reg}}} = 1$$

$$\begin{aligned} \overline{\mathcal{R}} &= \overline{\mathcal{L}_{\text{reg}}} \frac{d\mathcal{L}_{\text{reg}}}{d\mathcal{R}} \\ &= \overline{\mathcal{L}_{\text{reg}}} \lambda \end{aligned}$$

$$\begin{aligned} \overline{\mathcal{L}} &= \overline{\mathcal{L}_{\text{reg}}} \frac{d\mathcal{L}_{\text{reg}}}{d\mathcal{L}} \\ &= \overline{\mathcal{L}_{\text{reg}}} \end{aligned}$$


$$\begin{aligned} \overline{y} &= \overline{\mathcal{L}} \frac{d\mathcal{L}}{dy} \\ &= \overline{\mathcal{L}} (y - t) \end{aligned}$$

$$\begin{aligned} \overline{z} &= \overline{y} \frac{dy}{dz} \\ &= \overline{y} \sigma'(z) \end{aligned}$$

$$\begin{aligned} \overline{w} &= \overline{z} \frac{\partial z}{\partial w} + \overline{\mathcal{R}} \frac{d\mathcal{R}}{dw} \\ &= \overline{z} x + \overline{\mathcal{R}} w \end{aligned}$$

$$\begin{aligned} \overline{b} &= \overline{z} \frac{\partial z}{\partial b} \\ &= \overline{z} \end{aligned}$$

Backpropagation for Two-Layer Neural Network



Forward pass:


$$z_i = \sum_j w_{ij}^{(1)} x_j + b_i^{(1)}$$

$$h_i = \sigma(z_i)$$

$$y_k = \sum_i w_{ki}^{(2)} h_i + b_k^{(2)}$$

$$\mathcal{L} = \frac{1}{2} \sum_k (y_k - t_k)^2$$

Backward pass:



$$\bar{\mathcal{L}} = 1$$

$$\bar{y}_k = \bar{\mathcal{L}} (y_k - t_k)$$

$$\bar{w}_{ki}^{(2)} = \bar{y}_k h_i$$

$$\bar{b}_k^{(2)} = \bar{y}_k$$


k depends on number of hidden units you choose

$$\bar{h}_i = \sum_k \bar{y}_k w_{ki}^{(2)}$$

$$\bar{z}_i = \bar{h}_i \sigma'(z_i)$$

$$\bar{w}_{ij}^{(1)} = \bar{z}_i x_j$$

$$\bar{b}_i^{(1)} = \bar{z}_i$$



$x \rightarrow z \rightarrow h \rightarrow y \rightarrow \mathcal{L}$
 linear activation linear loss
 fn

$$\bar{\mathcal{L}} = 1$$

$$\bar{y}_i = \bar{\mathcal{L}} \frac{\partial \mathcal{L}}{\partial y_i} = \bar{\mathcal{L}} (y_i - t_i)$$

$$\star \bar{w}_{11}^{(2)} = \bar{y}_1 \frac{\partial y_1}{\partial w_{11}^{(2)}} = \bar{y}_1 h_1$$

$$\star \bar{b}_i^{(2)} = \bar{y}_i \frac{\partial y_i}{\partial b_i^{(2)}} = \bar{y}_i \cdot 1$$

$$z_i = \sum_j w_{ij}^{(1)} x_j + b_i^{(1)}$$

$$h_i = \sigma(z_i)$$

$$\star y_k = \sum_i w_{ki}^{(2)} h_i + b_k^{(2)} =$$


$$\mathcal{L} = \frac{1}{2} \sum_{k=1}^2 (y_k - t_k)^2 = \frac{(y_1 - t_1)^2 + (y_2 - t_2)^2}{2}$$

$$\bar{h}_1 = \sum_k \bar{y}_k \frac{\partial y_k}{\partial h_1}$$

$$\star = \sum_k \bar{y}_k w_{k1}^{(2)}$$

Backpropagation for Two-Layer Neural Network

In vectorized form:



Forward pass:

$$\mathbf{z} = \mathbf{W}^{(1)}\mathbf{x} + \mathbf{b}^{(1)}$$

$$\mathbf{h} = \sigma(\mathbf{z})$$

$$\mathbf{y} = \mathbf{W}^{(2)}\mathbf{h} + \mathbf{b}^{(2)}$$

$$\mathcal{L} = \frac{1}{2} \|\mathbf{t} - \mathbf{y}\|^2$$

Backward pass:

$$\bar{\mathcal{L}} = 1$$

$$\bar{\mathbf{y}} = \bar{\mathcal{L}}(\mathbf{y} - \mathbf{t})$$

$$\bar{\mathbf{W}}^{(2)} = \bar{\mathbf{y}}\mathbf{h}^\top$$

$$\bar{\mathbf{b}}^{(2)} = \bar{\mathbf{y}}$$

$$\bar{\mathbf{h}} = \mathbf{W}^{(2)\top}\bar{\mathbf{y}}$$

$$\bar{\mathbf{z}} = \bar{\mathbf{h}} \circ \sigma'(\mathbf{z})$$

$$\bar{\mathbf{W}}^{(1)} = \bar{\mathbf{z}}\mathbf{x}^\top$$

$$\bar{\mathbf{b}}^{(1)} = \bar{\mathbf{z}}$$

*update
with
gradient
descent*

resume

12:20

Computational Cost

- Computational cost of forward pass:
one add-multiply operation per weight

$$z_i = \sum_j w_{ij}^{(1)} x_j + b_i^{(1)}$$

$$f(x_1, x_2, \dots, x_n)$$

x_ϵ

$$f(x_1, x_2, \dots, x_{\epsilon+\epsilon})$$

$$- f(x_1, \dots, x_{\epsilon-\epsilon})$$

- Computational cost of backward pass:
two add-multiply operations per weight

$$\overline{w}_{ki}^{(2)} = \overline{y}_k h_i \text{ (hidden unit)}$$

$$\overline{h}_i = \sum_k \overline{y}_k w_{ki}^{(2)}$$

(how you affect the next layer) $z \in$

- One backward pass is as expensive as two forward passes.
- For a multilayer perceptron, this means the cost is linear in the number of layers, quadratic in the number of units per layer.

1000 \rightarrow 1000 units

Backpropagation

$$10^3 \cdot 10^3 = 10^6 \text{ weights}$$

- The algorithm for efficiently computing gradients in neural nets.
- Gradient descent with gradients computed via backprop is used to train the overwhelming majority of neural nets today.
- Even optimization algorithms much fancier than gradient descent (e.g. second-order methods) use backprop to compute the gradients.
- Despite its practical success, backprop is believed to be neurally implausible.

Auto-Differentiation


`loss.backward()`

- Suppose we construct our networks out of a series of “primitive” operations (e.g., add, multiply) with specified routines for computing derivatives.
- **Autodifferentiation** performs backprop in a completely mechanical and automatic way.
- Many autodiff libraries: PyTorch, Tensorflow, Jax, etc.
- Although autodiff automates the backward pass for you, it’s still important to know how things work under the hood.
- In CSC413, learn more about how autodiff works and use an autodiff framework to build complex neural networks.

1 Back-Propagation

2 Convolutional Networks

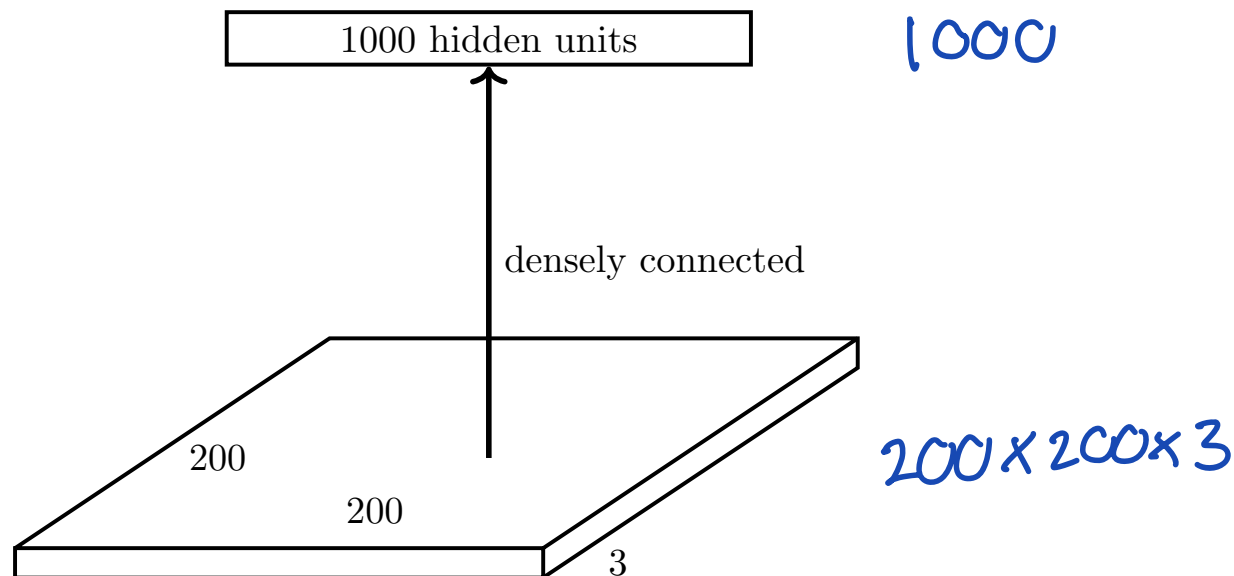
Robust to Transformations



- Must be robust to transformations or distortions:
 - ▶ change in pose/viewpoint
 - ▶ change in illumination
 - ▶ deformation
 - ▶ occlusion (some objects are hidden behind others)
- We would like the network to be **invariant**:
if the image is transformed slightly,
the classification shouldn't change.

Too Many Parameters

Want to train a network that takes a 200×200 RGB image as input.



What is the problem with having this as the first layer?

Too many parameters! Input size = $200 \times 200 \times 3 = 120\text{K}$.

Parameters = $120\text{K} \times 1000 = 120$ million.



120M in 1 FC layer

Shared Structures in the Network

- Some features, e.g. edges, corners, contours, object parts, may be useful in multiple locations in the image.
- We want feature detectors that are applicable in multiple locations in the image.

Convolution Layers



Fully connected layers:



Each hidden unit looks at the entire image.

Convolution Layers



Locally connected layers:



Each set of hidden units looks at a small region of the image.

Convolution Layers


Convolution layers:



Each set of hidden units looks at a small region of the image, and the weights are shared between all image locations.

Going Deeply Convolutional

Convolution layers can be stacked:




1-D Convolution


We have two signals/arrays x and w .

- x is an input signal (e.g. a waveform or an image).
- w is a set of k weights (also referred to as a kernel or filter).
- Often zero pad x to an infinite array

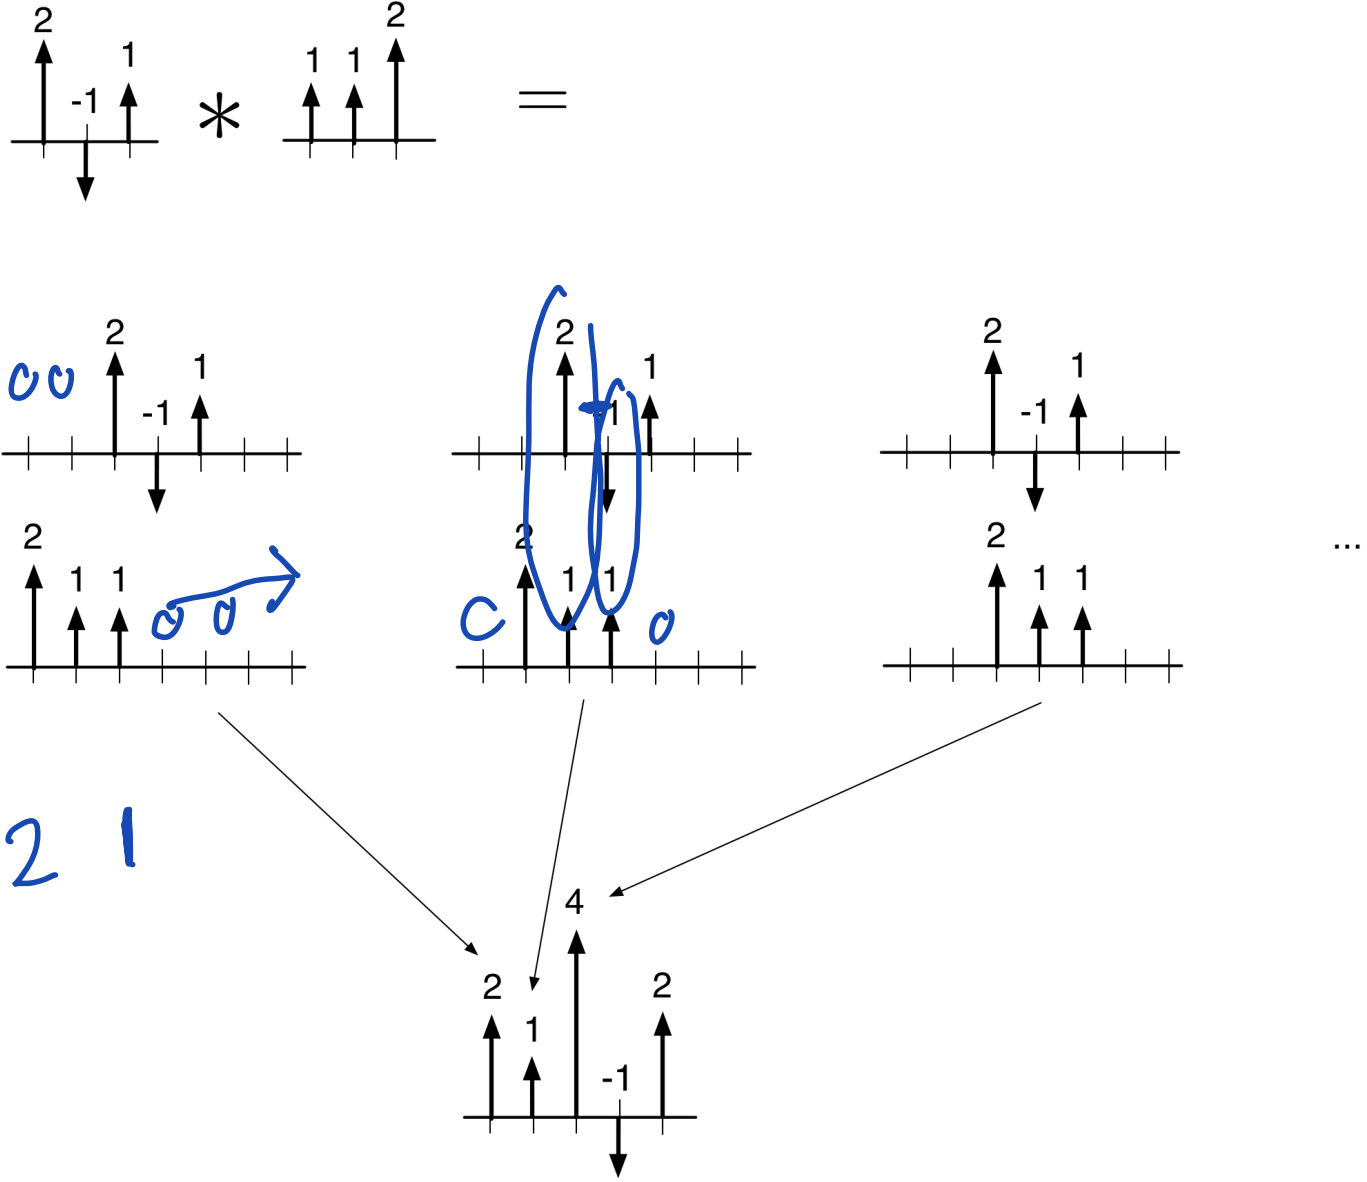
The t -th value in the convolution is defined below.

$$(x * w)[t] = \sum_{\tau=0}^{k-1} x[t - \tau]w[\tau].$$


Convolution Method 1: Translate-And-Scale



Convolution Method 2: Flip-And-Filter



Properties of Convolution

- Commutativity

$$a * b = b * a$$

- Linearity

$$a * (\lambda_1 b + \lambda_2 c) = \lambda_1 a * b + \lambda_2 a * c$$

2-D Convolution

2-D convolution is defined analogously to 1-D convolution.

If x and w are two 2-D arrays, then:

$$(x * w)[i, j] = \sum_s \sum_t x[i - s, j - t] * w[s, t].$$


2-D Convolution: Translate-and-Scale

$$\begin{array}{|c|c|c|} \hline 1 & 3 & 1 \\ \hline 0 & -1 & 1 \\ \hline 2 & 2 & -1 \\ \hline \end{array} * \begin{array}{|c|c|} \hline 1 & 2 \\ \hline 0 & -1 \\ \hline \end{array} = 1 \times \begin{array}{|c|c|c|c|} \hline 1 & 3 & 1 & \\ \hline 0 & -1 & 1 & \\ \hline 2 & 2 & -1 & \\ \hline & & & \\ \hline \end{array} + 2 \times \begin{array}{|c|c|c|c|} \hline & 1 & 3 & 1 \\ \hline & 0 & -1 & 1 \\ \hline & 2 & 2 & -1 \\ \hline & & & \\ \hline \end{array} = \begin{array}{|c|c|c|c|} \hline 1 & 5 & 7 & 2 \\ \hline 0 & -2 & -4 & 1 \\ \hline 2 & 6 & 4 & -3 \\ \hline 0 & -2 & -2 & 1 \\ \hline \end{array} + -1 \times \begin{array}{|c|c|c|c|} \hline & & & \\ \hline & 1 & 3 & 1 \\ \hline & 0 & -1 & 1 \\ \hline & 2 & 2 & -1 \\ \hline \end{array}$$

2-D Convolution: Flip-and-Filter

typical perspective
imagine dragging filter
across image

$$\begin{bmatrix} 1 & 3 & 1 \\ 0 & -1 & 1 \\ 2 & 2 & -1 \end{bmatrix} * \begin{bmatrix} 1 & 2 \\ 0 & -1 \end{bmatrix}$$



Example 1: What does this convolution kernel do?

blurring effect



*

0	1	0
1	4	1
0	1	0



Example 2: What does this convolution kernel do?

Sharpening



*

0	-1	0
-1	8	-1
0	-1	0



Example 3: What does this convolution kernel do?

edge detection




*

1	0	-1
2	0	-2
1	0	-1




Convolution Layer in Convolutional Networks

- Two types of layers: **convolution layers** (or detection layer), and **pooling layers**. *average pooling*
max pooling
- The convolution layer has a set of filters and produces a set of feature maps.
- Each feature map is a result of convolving the image with a filter.




$$h^{(2)} = \max(h_1^{(1)}, h_2^{(1)}, h_3^{(1)}, h_4^{(1)})$$

Example first-layer filters



convolution

$\overline{h}_i^{(1)}$ look like?
 $= \overline{h}^{(2)}$.
 $\begin{cases} 1 & \text{if } h_i^{(1)} \text{ is max} \\ 0 & \text{otherwise} \end{cases}$



(Zeiler and Fergus, 2013, Visualizing and understanding convolutional networks)

Non-linearity in Convolutional Networks

Common to apply a linear rectification nonlinearity:

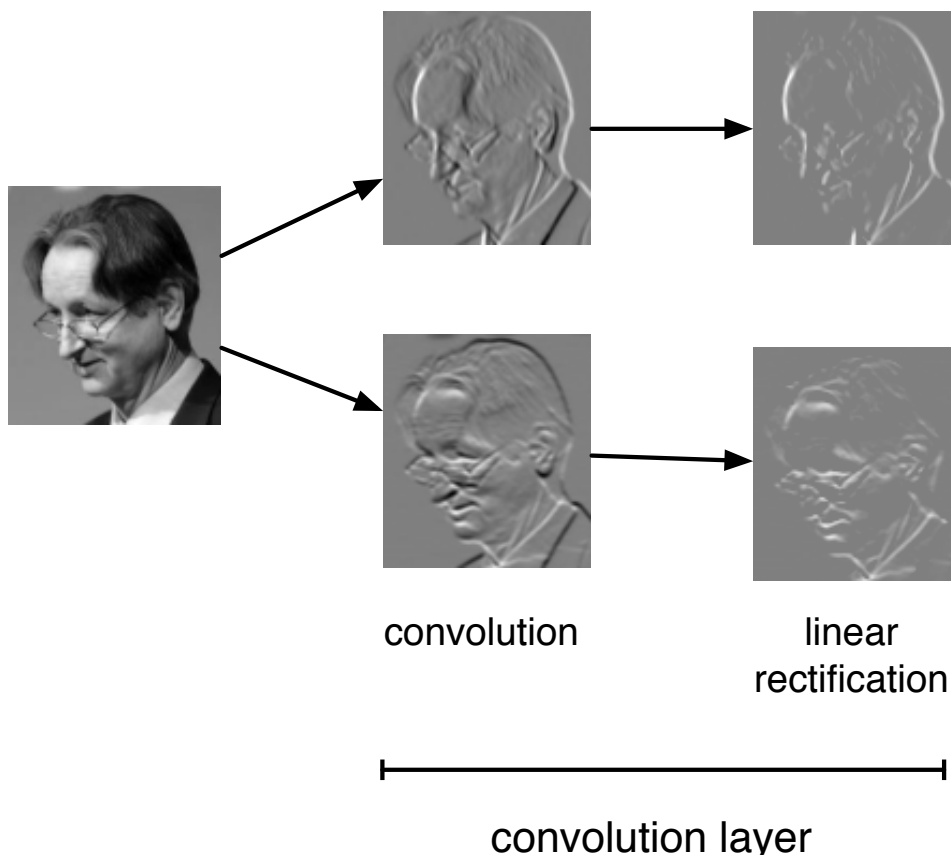
$$y_i = \max(z_i, 0).$$

ReLU

activation fn. for CNNs


Why might we do this?

Convolution is a linear operation. Therefore, we need a nonlinearity, otherwise 2 convolution layers would be no more powerful than 1.



Pooling Layers

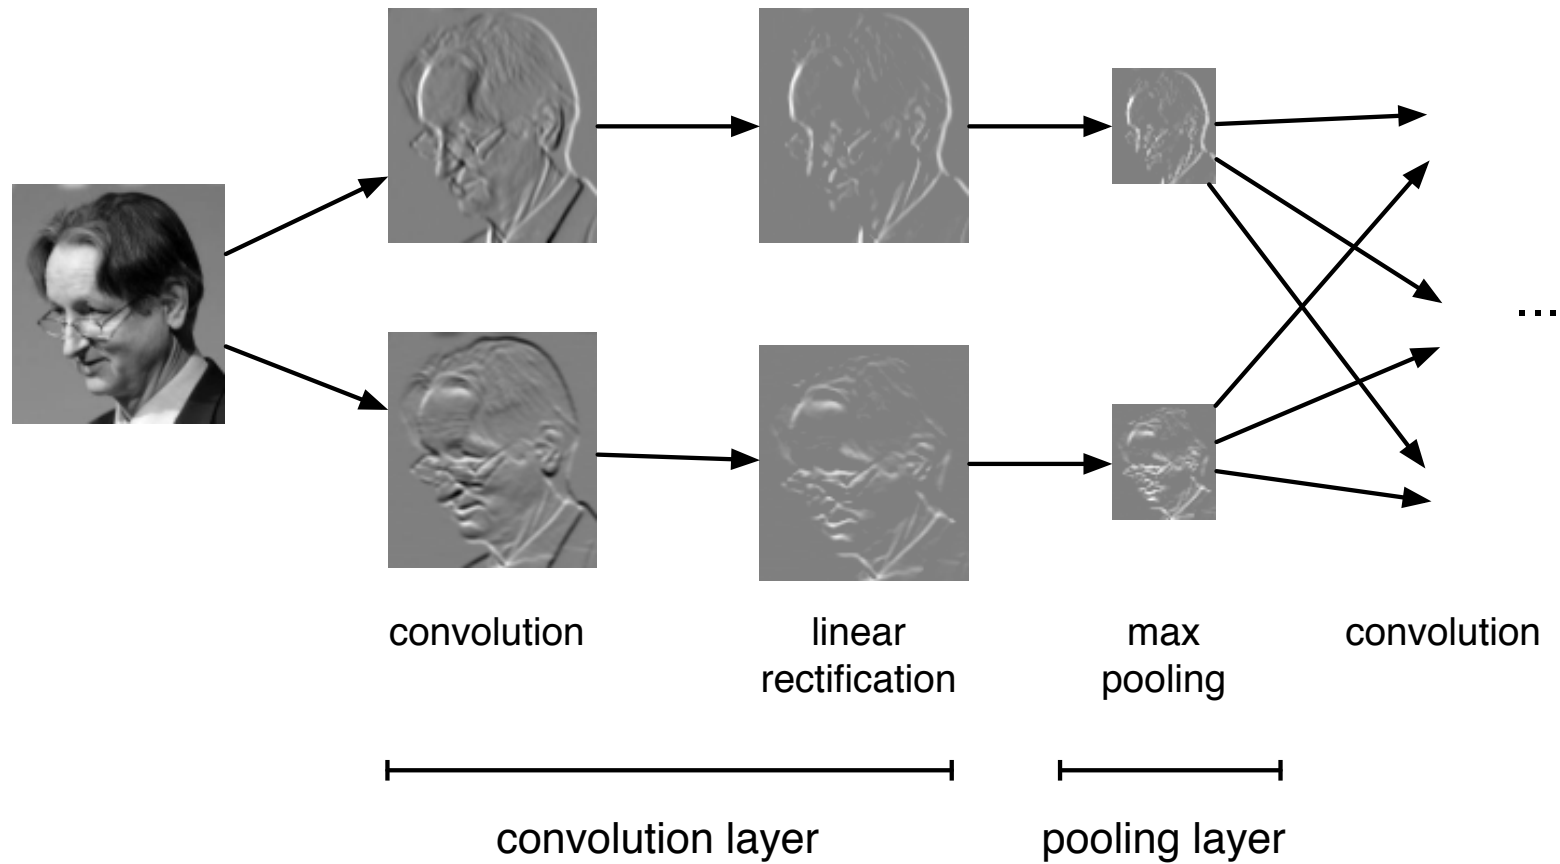
These layers reduce the size of the representation and build in in-variance to small transformations.



Most commonly, we use [max-pooling](#), which computes the maximum value of the units in a [pooling group](#):

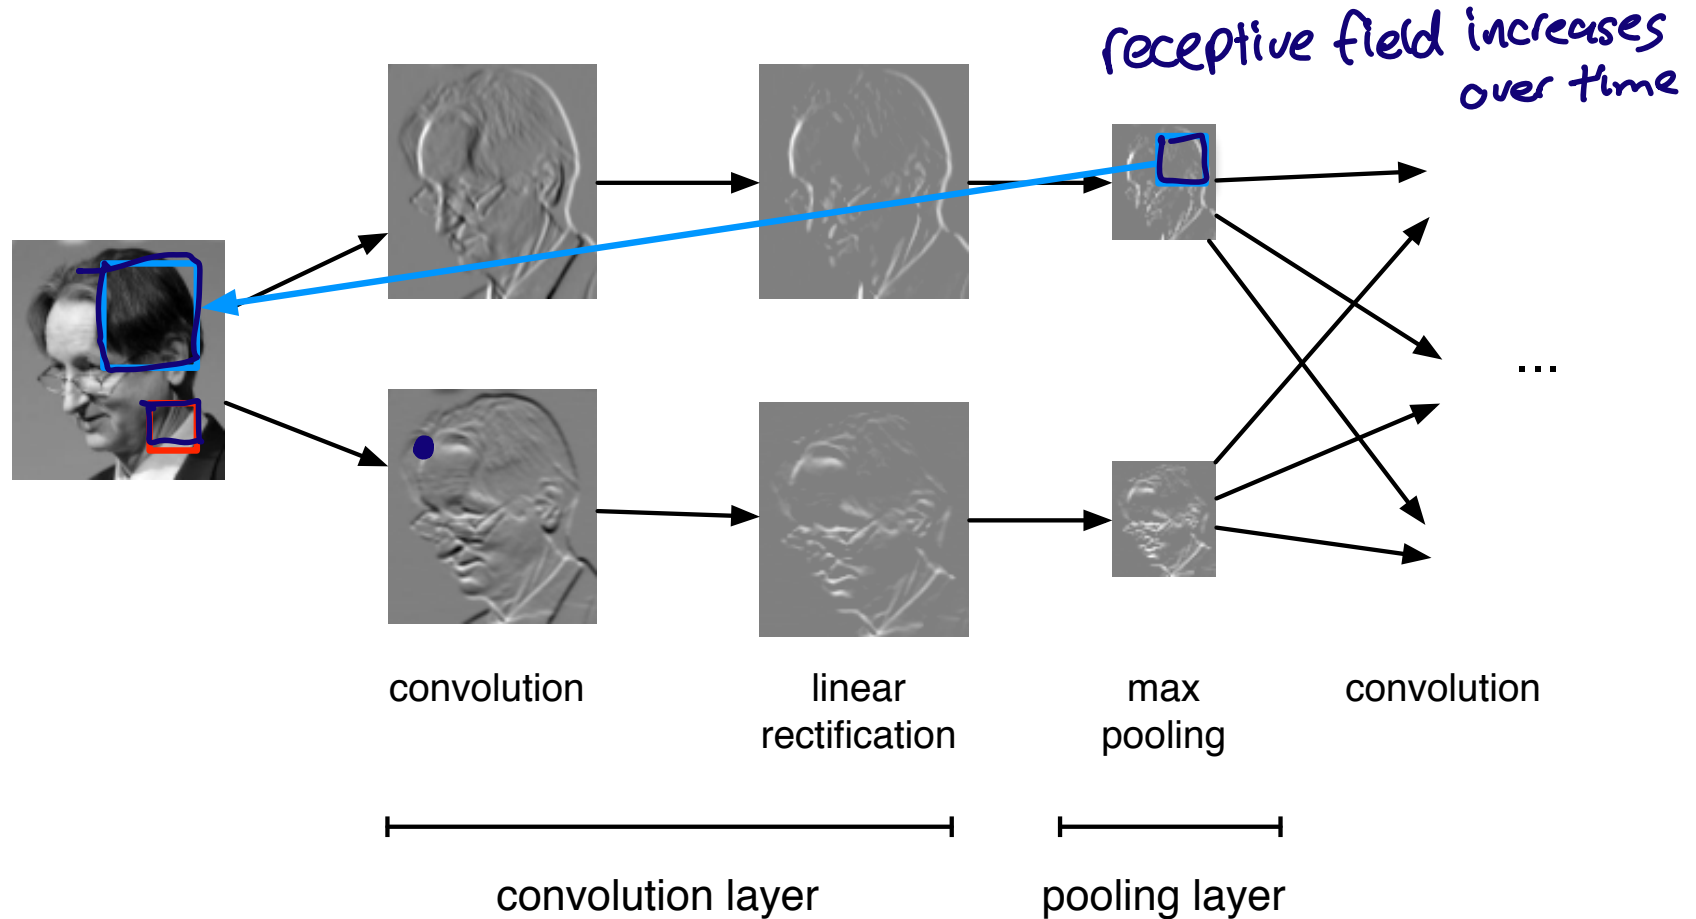
$$y_i = \max_{j \text{ in pooling group}} z_j$$

Convolutional networks



Convolutional Network Structure


Because of pooling, higher-layer filters can cover a larger region of the input than equal-sized filters in the lower layers.



Equivariance and Invariance

The network's responses should be robust to translations of the input. But this can mean two different things.

- Convolution layers are **equivariant**: if you translate the inputs, the outputs are translated by the same amount.
- Want the network's predictions to be **invariant**: if you translate the inputs, the prediction should not change. Pooling layers provide invariance to small translations.



Convolution Layers

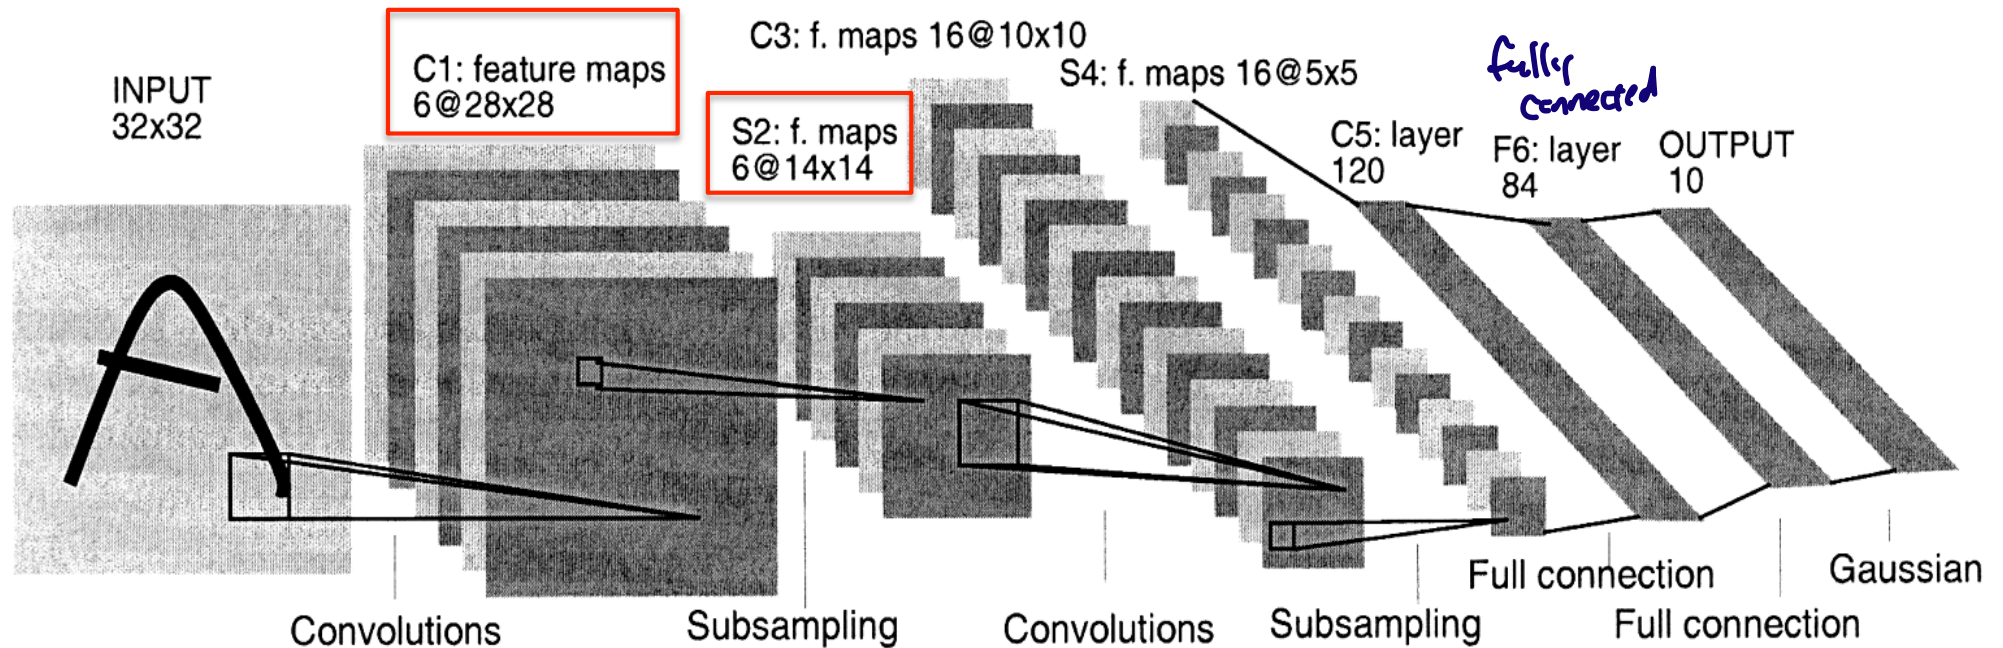
Each layer consists of several **feature maps**, or **channels** each of which is an array.

- If the input layer represents a grayscale image, it consists of one channel. If it represents a color image, it consists of three channels.

Each unit is connected to each unit within its receptive field in the previous layer. This includes *all* of the previous layer's feature maps.


LeNet

The **LeNet** architecture applied to handwritten digit recognition on MNIST in 1998:



AlexNet

AlexNet, like LeNet but scaled up in every way (more layers, more units, more connections, etc.):



(Krizhevsky et al., 2012)

AlexNet's stunning performance on the ImageNet competition is what got everyone excited about deep learning in 2012.


ImageNet Results Over the Years

There are 1000 classes. Top-5 errors mean that the network can make 5 guesses for each image. So chance is 0.5%.

Year	Model	Top-5 error
2010	Hand-designed descriptors + SVM	28.2%
2011	Compressed Fisher Vectors + SVM	25.8%
2012	AlexNet	16.4%
2013	a variant of AlexNet	11.7%
2014	GoogLeNet	6.6%
2015	deep residual nets	4.5%

neural network
↓

ResNet



$$h_2 = \sigma(W h_1)$$

ResNet

$$h_2 = h_1 + \sigma(W h_1)$$

Human-level performance is around 5.1%.

No longer running the object recognition competition because the performance is already so good.