



**ARAMIS  
LAB**  
BRAIN DATA SCIENCE

[www.aramislab.fr](http://www.aramislab.fr)

[ninonburgos.com](http://ninonburgos.com)

10<sup>th</sup> March 2020



**Inserm**



**SORBONNE  
UNIVERSITÉ**



*inria*  
informatics mathematics

**PR[AI]RIE**  
PaRis Artificial Intelligence Research InstitutE

# Introduction to deep learning

**Ninon Burgos, CNRS Researcher**

Aramis Lab, Paris, France

“When you’re fundraising, it’s AI

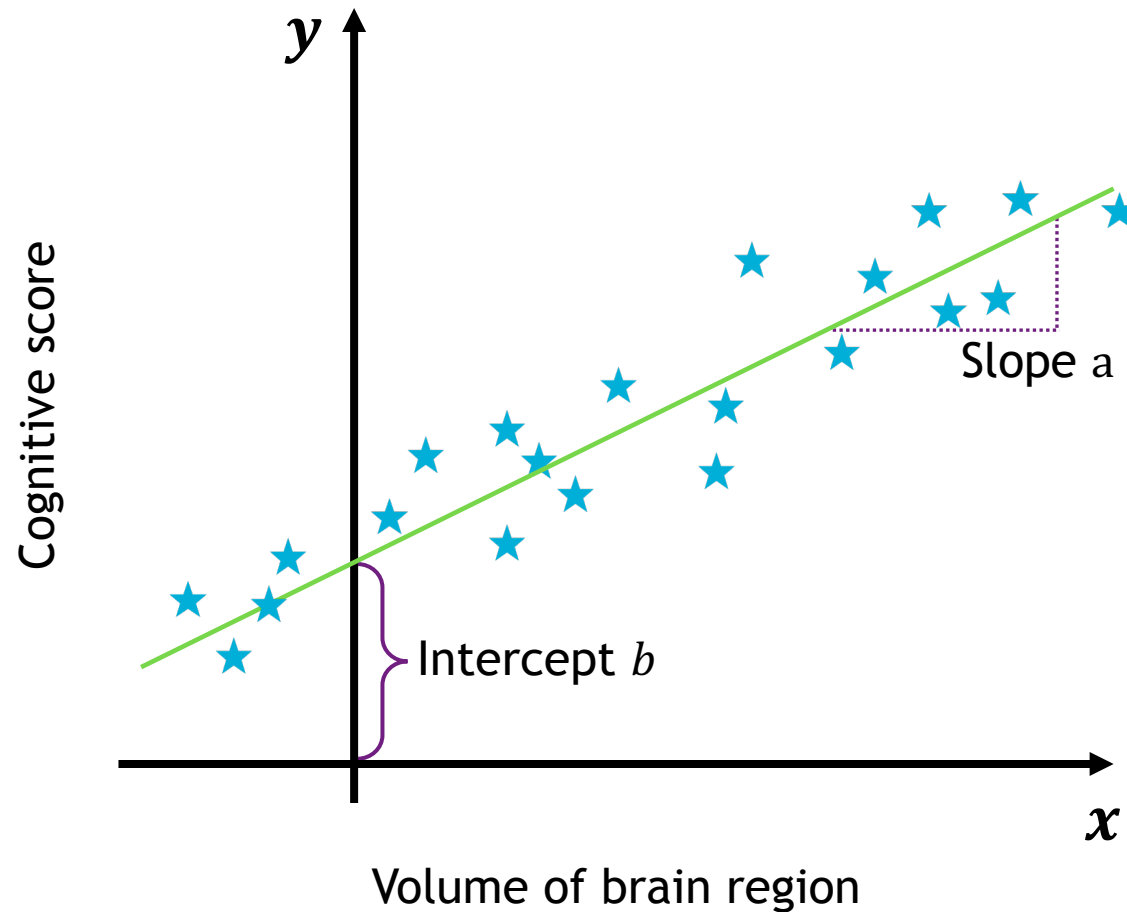
When you’re hiring, it’s ML

When you’re implementing, it’s linear regression”

– Baron Schwartz


# The basics

## What is linear regression?



$$y = f(x) = ax + b$$


## Cost function



Mean squared error:

$$J = \frac{1}{n} \sum_{i=1}^n (\hat{y}_i - y_i)^2$$

## Learning




Find the model with the minimal error:


$$\hat{f} = \arg \min_{f \in \mathcal{F}} J(f)$$

## Gradient descent

Small learning rate



Large learning rate



## Gradient descent

**Cost function:**

$$J = \frac{1}{n} \sum_{i=1}^n (\hat{y}_i - y_i)^2 = \frac{1}{n} \sum_{i=1}^n (ax_i + b - y_i)^2$$

**Partial derivatives:**

$$\frac{\partial J}{\partial a} = \frac{2}{n} \sum_{i=1}^n (ax_i + b - y_i) \cdot x_i = \frac{2}{n} \sum_{i=1}^n (\hat{y}_i - y_i) \cdot x_i$$

$$\frac{\partial J}{\partial b} = \frac{2}{n} \sum_{i=1}^n (ax_i + b - y_i) = \frac{2}{n} \sum_{i=1}^n (\hat{y}_i - y_i)$$

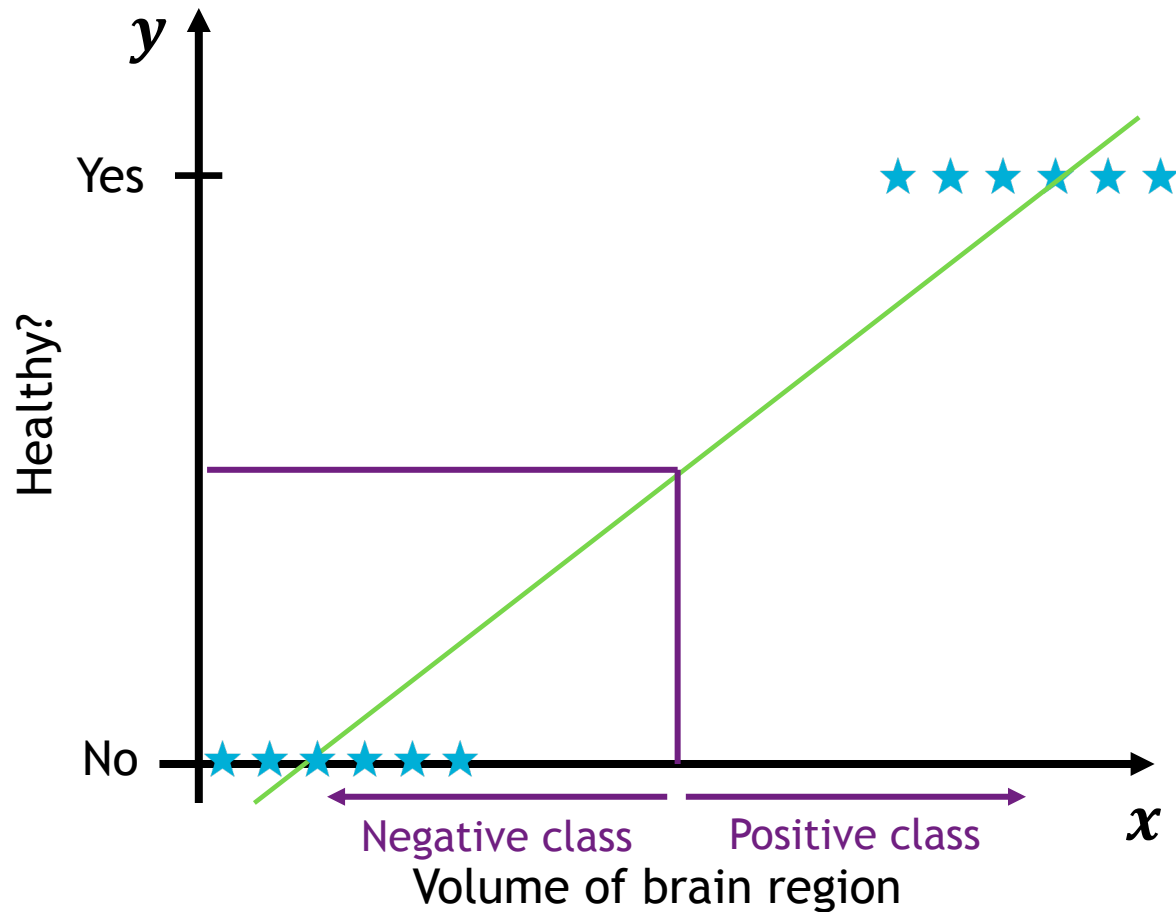
**Update of  $a$  and  $b$ :**

$$a \leftarrow a - \eta \frac{\partial J}{\partial a} = a - \eta \frac{2}{n} \sum_{i=1}^n (\hat{y}_i - y_i)$$

$$b \leftarrow b - \eta \frac{\partial J}{\partial b} = b - \eta \frac{2}{n} \sum_{i=1}^n (\hat{y}_i - y_i) \cdot x_i$$

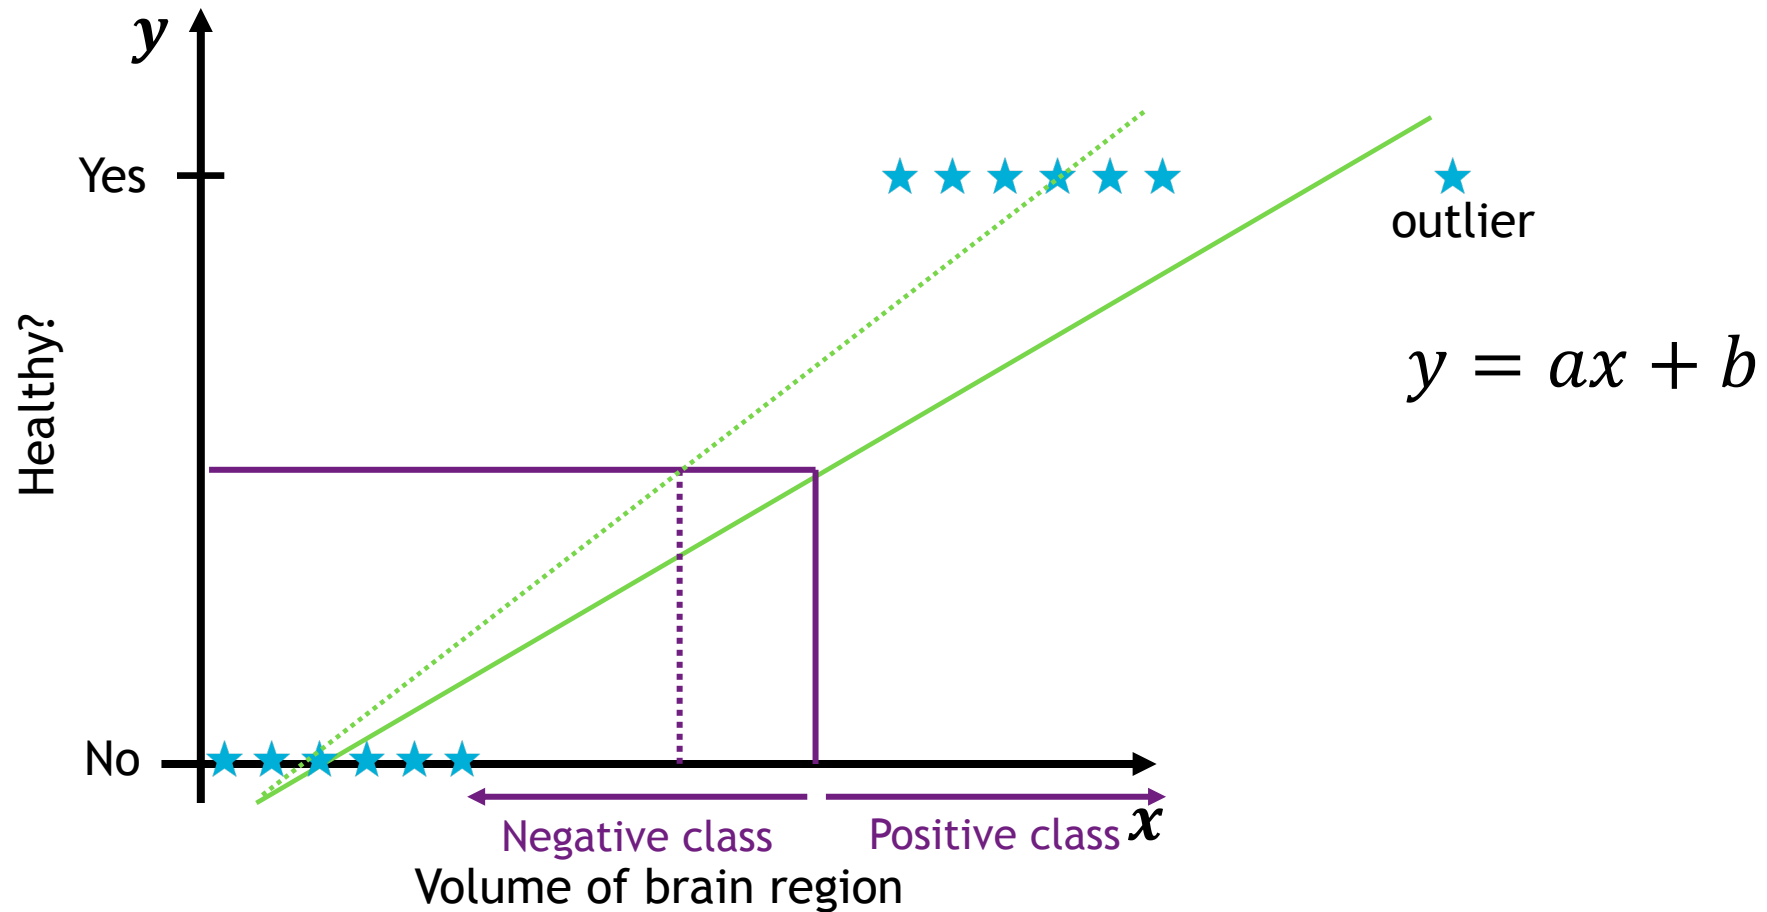


## Linear function

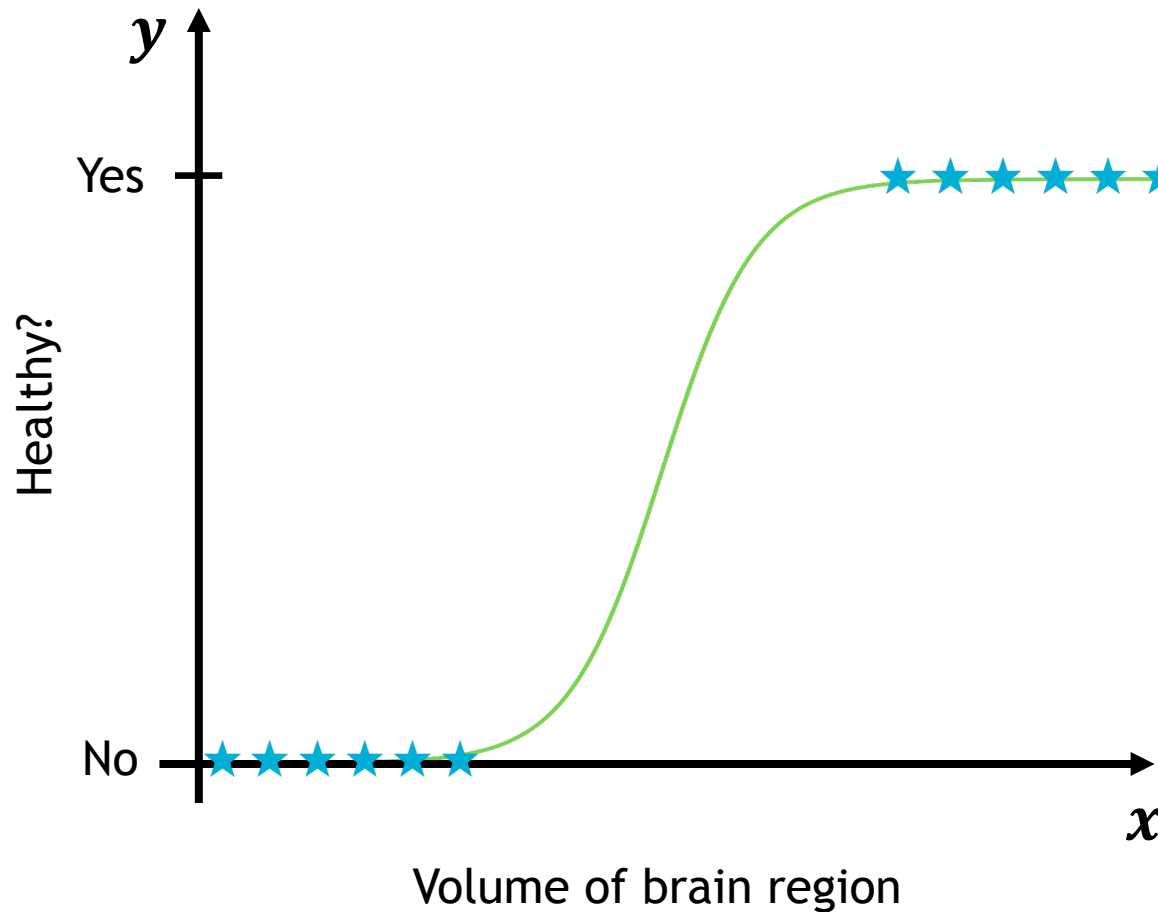


$$y = ax + b$$

## Linear function



## Sigmoid function



$$y = ax + b$$

$$f = \frac{1}{1 + e^{-y}} = \frac{1}{1 + e^{-(ax+b)}}$$

## Multiple input variables

$$z = w_0 + w_1x_1 + w_2x_2 + \dots + w_mx_m = w_0 + \sum_{j=1}^m x_jw_j$$

$$f = \frac{1}{1 + e^{-z}} = \frac{1}{1 + e^{-(w_0 + \sum_{j=1}^m x_jw_j)}}$$


## Loss function

- Low if prediction is correct
- High if prediction is wrong

$\left\{ \begin{array}{l} \text{if } y = 1, l(f(\mathbf{x}), y) \text{ should be low when } f(\mathbf{x}) \text{ is high} \\ \text{if } y = 0, l(f(\mathbf{x}), y) \text{ should be low when } f(\mathbf{x}) \text{ is low} \end{array} \right.$

$$l(f(\mathbf{x}), y) = \begin{cases} -\log(f(\mathbf{x})), & \text{if } y = 1 \\ -\log(1 - f(\mathbf{x})), & \text{if } y = 0 \end{cases}$$

$$l(f(\mathbf{x}), y) = -y \log(f(\mathbf{x})) - (1 - y) \log(1 - f(\mathbf{x}))$$



## Cost function

$$J(f) = \frac{1}{n} \sum_{i=1}^n l(f(\mathbf{x}_i), y_i)$$

$$J(f) = -\frac{1}{n} \sum_{i=1}^n (y_i \log(f(\mathbf{x}_i)) + (1 - y_i) \log(1 - f(\mathbf{x}_i)))$$

## Learning

Find the model with the minimal error:

$$\hat{f} = \arg \min_{f \in \mathcal{F}} J(f)$$

## Gradient descent

**Cost function:**

$$J(f) = -\frac{1}{n} \sum_{i=1}^n (y_i \log(f(\mathbf{x}_i)) + (1 - y_i) \log(1 - f(\mathbf{x}_i)))$$

**Partial derivatives:**

$$\frac{\partial J}{\partial w_j} = -\frac{1}{n} \sum_{i=1}^n (y_{i,j} - f(\mathbf{x}_i)) \cdot x_{i,j}$$

**Update of  $a$  and  $b$ :**

$$w_j \leftarrow w_j - \eta \frac{\partial J}{\partial w_j}$$

Input variables:  $\mathbf{x} = \begin{bmatrix} x_1 \\ \vdots \\ x_m \end{bmatrix}$

Input features

Output:  $y$

Model:  $f, y = f(\mathbf{x})$

The "artificial intelligence"

Loss:  $l(f(\mathbf{x}), y)$

Quantifies how much the prediction is far from the true output

Cost function:  $J(f) = \frac{1}{n} \sum_{i=1}^n l(f(\mathbf{x}_i), y_i)$

Quantifies how much the prediction is far from the true output across all training examples

Learning:  $\hat{f} = \arg \min_{f \in \mathcal{F}} J(f)$

Find the model with the minimal error

Gradient descent:  $w \leftarrow w - \eta \frac{\partial J}{\partial w}$

Method to find the model with the minimal error






# ML vs DL : learning features

ML

Feature extraction



Classification



Glioblastoma

Herpes simplex encephalitis

DL




Glioblastoma

Herpes simplex encephalitis

# Neural networks

## Artificial neuron



Output

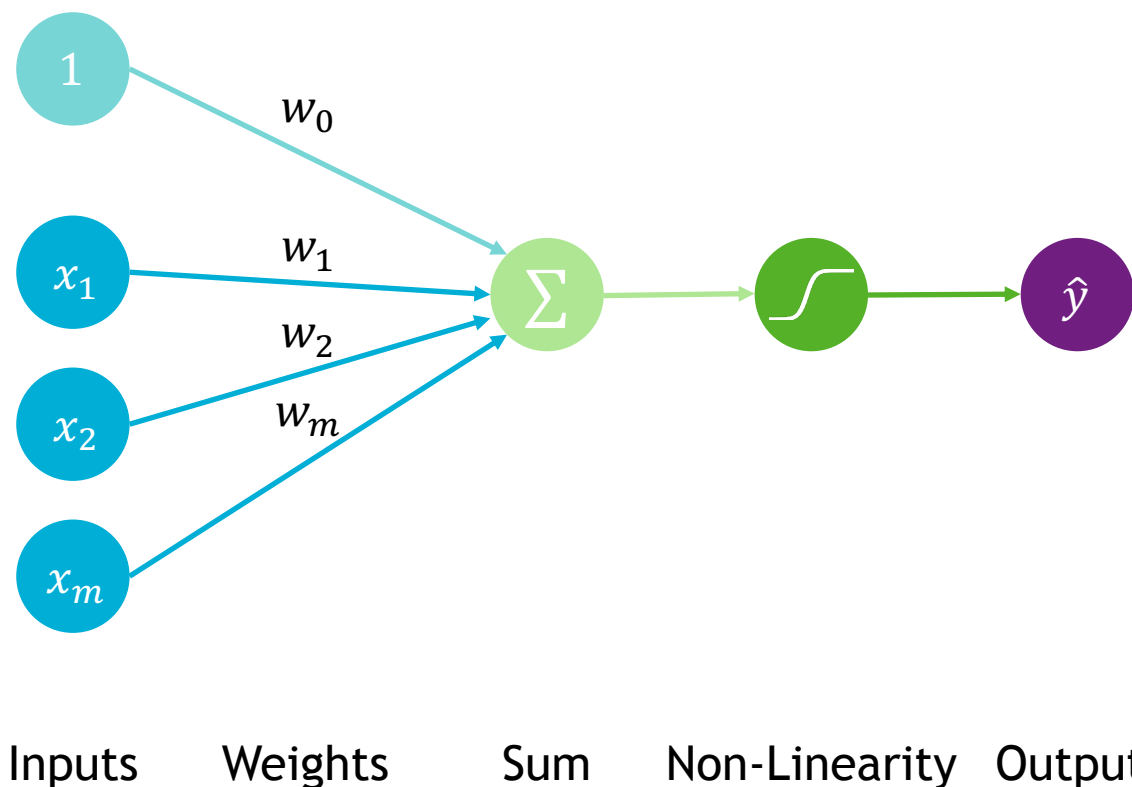
Linear combination of inputs

$$\hat{y} = h \left( \sum_{j=1}^m x_j w_j \right)$$

Non-linear activation function

Inputs    Weights    Sum    Non-Linearity    Output

## Artificial neuron



Output      Bias      Linear combination of inputs

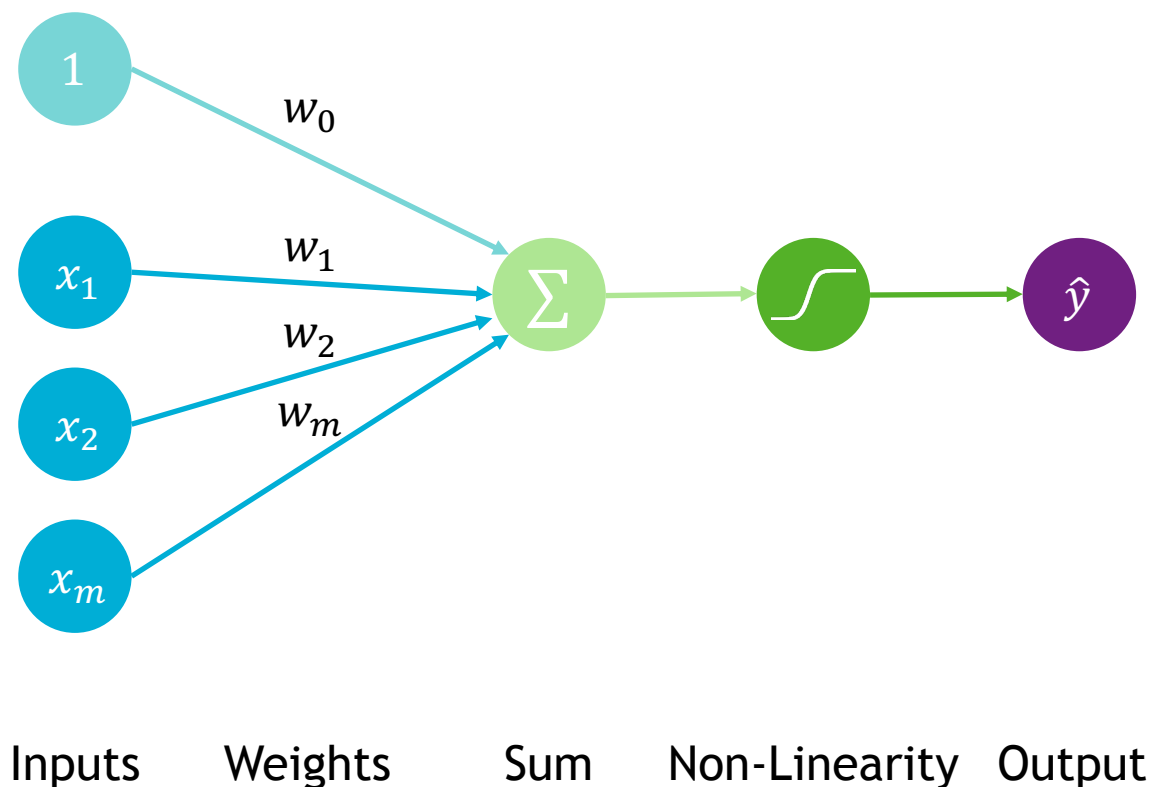
$$\hat{y} = h \left( w_0 + \sum_{j=1}^m x_j w_j \right)$$

Non-linear activation function

$$\hat{y} = h(w_0 + \mathbf{X}^T \mathbf{W})$$


where  $\mathbf{X} = \begin{bmatrix} x_1 \\ \vdots \\ x_m \end{bmatrix}$  and  $\mathbf{W} = \begin{bmatrix} w_1 \\ \vdots \\ w_m \end{bmatrix}$

## Activation function



$$\hat{y} = h(w_0 + \mathbf{X}^T \mathbf{W})$$


Sigmoid function:  $h(z) = \frac{1}{1 + e^{-z}}$



→ Logistic regression

## Examples of activation functions


### Sigmoid



$$h(z) = \frac{1}{1 + e^{-z}}$$

$$h'(z) = h(z)(1 - h(z))$$


### Hyperbolic tangent



$$h(z) = \frac{e^z - e^{-z}}{e^z + e^{-z}}$$

$$h'(z) = 1 - h(z)^2$$


### Rectified linear unit (ReLU)




$$h(z) = \max(0, z)$$

$$h'(z) = \begin{cases} 1, & z > 0 \\ 0, & \text{otherwise} \end{cases}$$

## Importance of non-linear activation functions

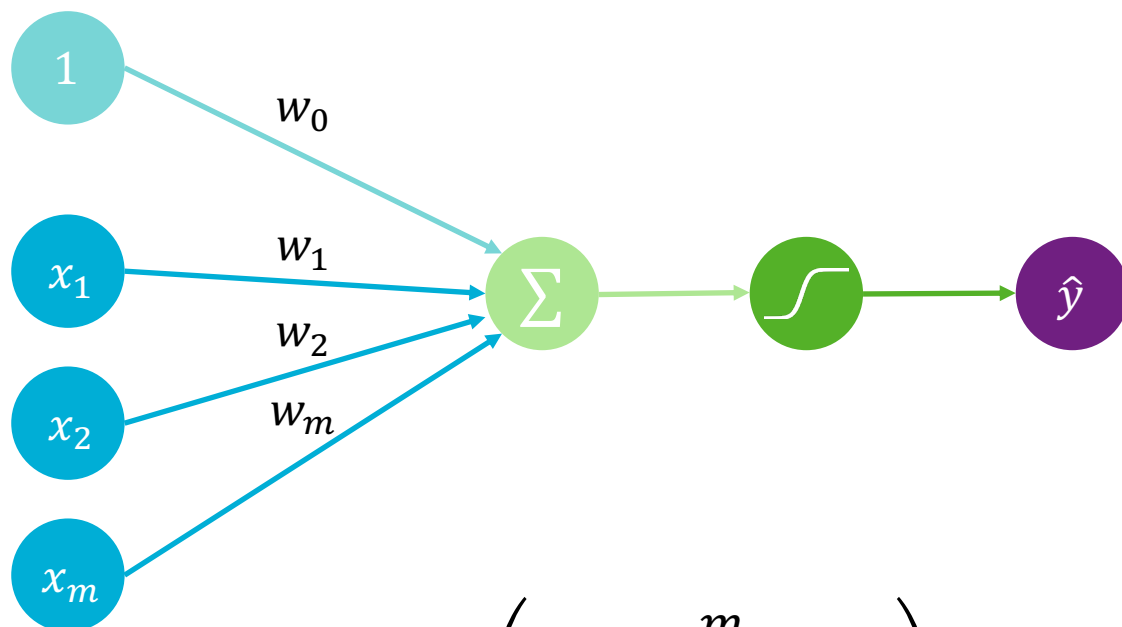


Linear activation functions  
→ linear decisions




Non-linear activation functions  
→ arbitrarily complex decisions

## Simplified notation




$$\hat{y} = h \left( w_0 + \sum_{j=1}^m x_j w_j \right)$$



$$z = w_0 + \sum_{j=1}^m x_j w_j$$




## Multi output neural network with dense layers



$$z_i = w_{0,i} + \sum_{j=1}^m x_j w_{j,i}$$


# Single layer neural network



$$z_i = w_{0,i}^{(1)} + \sum_{j=1}^m x_j w_{j,i}^{(1)}$$

$$\hat{y}_i = h \left( w_{0,i}^{(2)} + \sum_{j=1}^{d_1} z_j w_{j,i}^{(2)} \right)$$

# Single layer neural network



Inputs

Hidden


Output

$$z_2 = w_{0,2}^{(1)} + \sum_{j=1}^m x_j w_{j,2}^{(1)}$$

$$= w_{0,2}^{(1)} + x_1 w_{1,2}^{(1)} + x_2 w_{2,2}^{(1)} + x_m w_{m,2}^{(1)}$$

# Single layer neural network

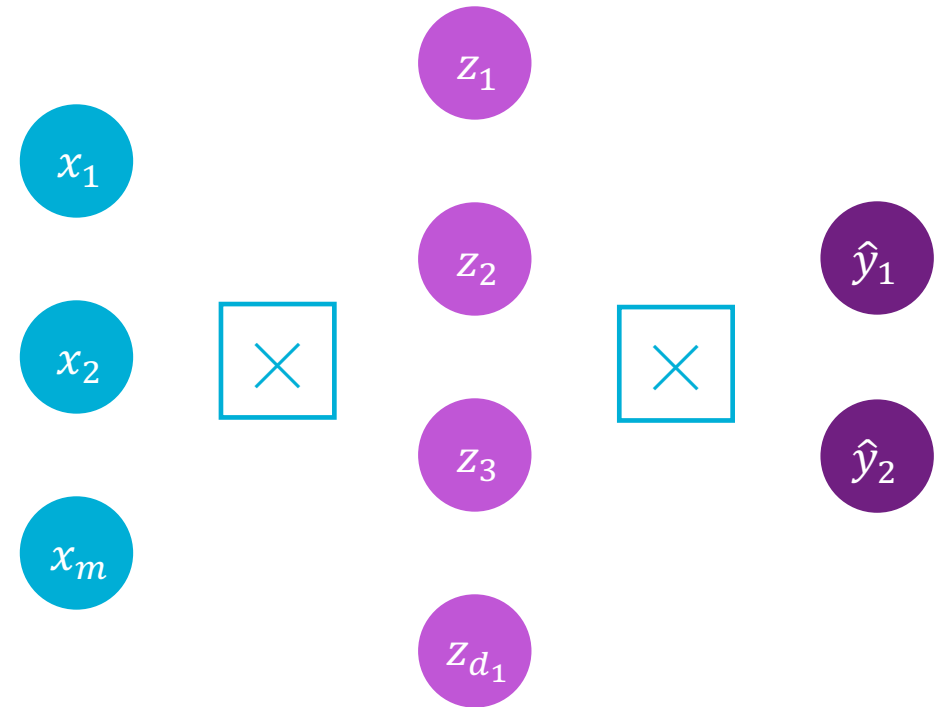
## Simplified notation



Inputs

Hidden

Output




Inputs

Hidden

Output

## Stacking layers




$$z_{k,i} = w_{0,i}^{(k)} + \sum_{j=1}^{d_{k-1}} h(z_{k-1,j}) w_{j,i}^{(k)}$$

## Diagnosis of dementia based on imaging biomarkers

Is this subject healthy?

$x_2$  = brain  
metabolism




★ Healthy

★ Dementia

$x_1$  = volume of fluid  
in the brain

## Diagnosis of dementia based on imaging biomarkers

Is this subject healthy?




Loss:  $l(\underbrace{f(x^{(i)}; \mathbf{W})}_{\text{Predicted}}, \underbrace{y^{(i)}}_{\text{Actual}})$

## Diagnosis of dementia based on imaging biomarkers

Is this subject healthy?

$$\mathbf{X} = \begin{bmatrix} 1 & 1.5 \\ 3.5 & 1.2 \\ 5 & 0.9 \\ \vdots & \vdots \end{bmatrix}$$



$$f(x) = \begin{bmatrix} 0.1 \\ 0.4 \\ 0.6 \\ \vdots \end{bmatrix} \quad y = \begin{bmatrix} 1 \\ 0 \\ 0 \\ \vdots \end{bmatrix}$$

Cost function:


$$J(\mathbf{W}) = \frac{1}{n} \sum_{i=1}^n l(\underbrace{f(x^{(i)}; \mathbf{W})}_{\text{Predicted}}, \underbrace{y^{(i)}}_{\text{Actual}})$$



## Diagnosis of dementia based on imaging biomarkers

Is this subject healthy?

$$\mathbf{X} = \begin{bmatrix} 1 & 1.5 \\ 3.5 & 1.2 \\ 5 & 0.9 \\ \vdots & \vdots \end{bmatrix}$$



$$\begin{array}{cc} f(x) & y \\ \begin{bmatrix} 0.1 \\ 0.4 \\ 0.6 \\ \vdots \end{bmatrix} & \begin{bmatrix} 1 \\ 0 \\ 0 \\ \vdots \end{bmatrix} \end{array}$$


Cost function with cross entropy loss:

$$J(\mathbf{W}) = \frac{1}{n} \sum_{i=1}^n \underbrace{-y^{(i)}}_{\text{Actual}} \log \left( \underbrace{f(x^{(i)}; \mathbf{W})}_{\text{Predicted}} \right) - \left( \underbrace{1 - y^{(i)}}_{\text{Actual}} \right) \log \left( \underbrace{1 - f(x^{(i)}; \mathbf{W})}_{\text{Predicted}} \right)$$

## Diagnosis of dementia based on imaging biomarkers

What is the dementia severity?

$$\mathbf{X} = \begin{bmatrix} 1 & 1.5 \\ 3.5 & 1.2 \\ 5 & 0.9 \\ \vdots & \vdots \end{bmatrix}$$



$$f(x) = \begin{bmatrix} 3 \\ 1 \\ 8 \\ \vdots \end{bmatrix} \quad y = \begin{bmatrix} 0 \\ 2 \\ 7 \\ \vdots \end{bmatrix}$$

Cost function with mean squared error loss:

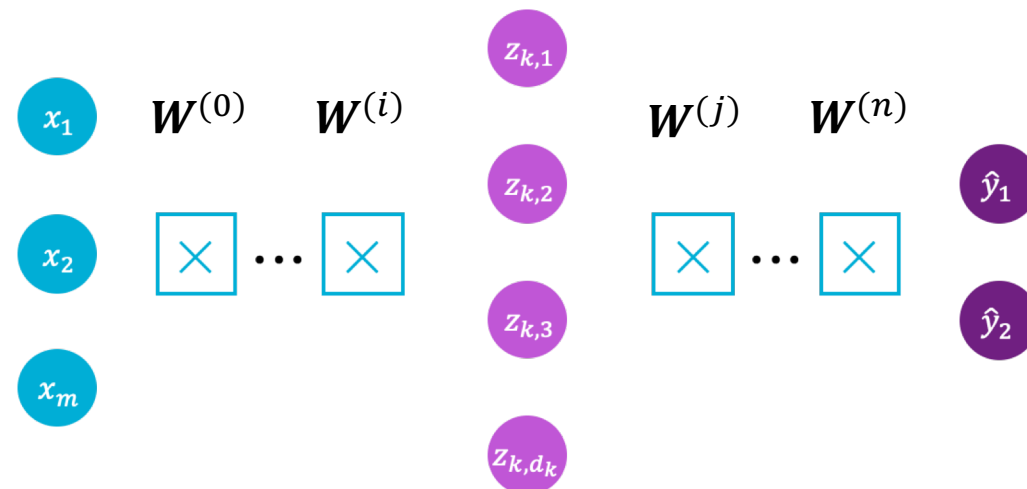
$$J(\mathbf{W}) = \frac{1}{n} \sum_{i=1}^n \left( \underbrace{y^{(i)}}_{\text{Actual}} - \underbrace{f(x^{(i)}; \mathbf{W})}_{\text{Predicted}} \right)^2$$

## Loss optimisation

- Find the network weights that achieve the lowest loss

$$\mathbf{W}^* = \underset{\mathbf{W}}{\operatorname{argmin}} J(\mathbf{W}) = \underset{\mathbf{W}}{\operatorname{argmin}} \frac{1}{n} \sum_{i=1}^n l(f(x^{(i)}; \mathbf{W}), y^{(i)})$$

$\uparrow$   
 $\mathbf{W} = \{\mathbf{W}^{(0)}, \mathbf{W}^{(1)}, \dots\}$



## Gradient descent


### Algorithm

1. Initialise weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence
  - a. Compute gradient  $\frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
  - b. Update weights  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
3. Return weights

## Computing gradients: backpropagation



## Computing gradients: backpropagation



$$\frac{\partial J(\mathbf{W})}{\partial w_2} = \frac{\partial J(\mathbf{W})}{\partial \hat{y}} \times \frac{\partial \hat{y}}{\partial w_2}$$

Below the equation, there are two colored bars: a purple bar under the term  $\frac{\partial J(\mathbf{W})}{\partial \hat{y}}$  and a green bar under the term  $\frac{\partial \hat{y}}{\partial w_2}$ .

## Computing gradients: backpropagation



$$\frac{\partial J(\mathbf{W})}{\partial w_2} = \frac{\partial J(\mathbf{W})}{\partial \hat{y}} \times \frac{\partial \hat{y}}{\partial w_2}$$

Below the equation, there are two colored bars: a purple bar under  $\frac{\partial J(\mathbf{W})}{\partial \hat{y}}$  and a green bar under  $\frac{\partial \hat{y}}{\partial w_2}$ .


$$\frac{\partial J(\mathbf{W})}{\partial w_1} = \frac{\partial J(\mathbf{W})}{\partial \hat{y}} \times \frac{\partial \hat{y}}{\partial w_1} = \frac{\partial J(\mathbf{W})}{\partial \hat{y}} \times \frac{\partial \hat{y}}{\partial z_1} \times \frac{\partial z_1}{\partial w_1}$$

Below the equation, there are three colored bars: a purple bar under  $\frac{\partial J(\mathbf{W})}{\partial \hat{y}}$ , a green bar under  $\frac{\partial \hat{y}}{\partial z_1}$ , and a blue bar under  $\frac{\partial z_1}{\partial w_1}$ .

## Gradient descent in practice

### Algorithm

1. Initialise weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence
  - a. Compute gradient  $\frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
  - b. Update weights  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
3. Return weights




Li et al., Visualizing the Loss Landscape of Neural Nets, NIPS, 2018



## Learning rate

$$W \leftarrow W - \eta \frac{\partial J(W)}{\partial W}$$


- If  $\eta$  is too small: slow to converge, may be trapped in local minima



## Learning rate

$$W \leftarrow W - \eta \frac{\partial J(W)}{\partial W}$$

- If  $\eta$  is too small: slow to converge, may be trapped in local minima
- If  $\eta$  is too large: may diverge




## Learning rate

$$W \leftarrow W - \eta \frac{\partial J(W)}{\partial W}$$

- If  $\eta$  is too small: slow to converge, may be trapped in local minima
- If  $\eta$  is too large: may diverge

→ Adaptive learning rate



## Adaptative learning rate

- Learning rates are no longer fixed
- Can be made larger or smaller depending on:
  - how large the gradient is
  - how fast learning is happening
  - the size of particular weights
  - etc.
- Algorithms:
  - Adam [Kingma et al., Adam: A Method for Stochastic Optimization, 2014]
  - Adadelta [Zeiler et al., ADADELTA: An Adaptive Learning Rate Method, 2012]
  - Adagrad [Duchi et al., Adaptive Subgradient Methods for Online Learning and Stochastic Optimization, 2011]
  - RMSProp [Hinton, Neural Networks for Machine Learning]

## Standard gradient descent

### Algorithm

1. Initialise weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence
  - a. Compute gradient  $\frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$  Can be expensive to compute
  - b. Update weights  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
3. Return weights

## Stochastic gradient descent with one sample

### Algorithm

1. Initialise weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence
  - a. Pick single data point  $i$
  - b. Compute gradient  $\frac{\partial J_i(\mathbf{W})}{\partial \mathbf{W}}$  Easy to compute but very noisy
  - c. Update weights  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$
3. Return weights

## Stochastic gradient descent with several samples (mini-batch)

### Algorithm

1. Initialise weights randomly  $\sim \mathcal{N}(0, \sigma^2)$
2. Loop until convergence

a. Pick batch of  $B$  data points


Fast to compute and good estimate of the true gradient

b. Compute gradient  $\frac{\partial J(\mathbf{W})}{\partial \mathbf{W}} = \frac{1}{B} \sum_{k=1}^B \frac{\partial J_k(\mathbf{W})}{\partial \mathbf{W}}$


c. Update weights  $\mathbf{W} \leftarrow \mathbf{W} - \eta \frac{\partial J(\mathbf{W})}{\partial \mathbf{W}}$

3. Return weights


## Overfitting



Underfitting




Ideal fit




Overfitting



## Regularisation: Dropout




## Regularisation: Early stopping




## The Perceptron

- Structural building blocks
- Nonlinear activation functions



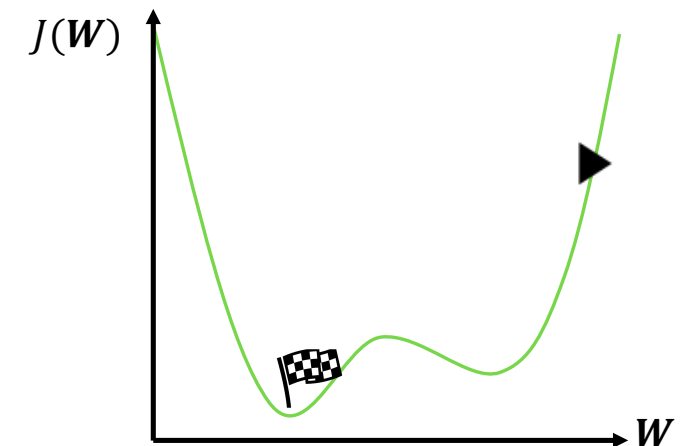
## Neural Networks

- Stacking Perceptrons to form neural networks
- Optimisation through backpropagation



## Training in Practice


- Adaptive learning
- Batching
- Regularisation



# Convolutional neural networks

## Diagnosis of dementia based on imaging biomarkers

Is this subject healthy?









## Fully connected neural network

0	0	0	0	0	0	1	2	4	4	6	4	5	3	5	4	3	5	5	4	5	4	3	1	0	0	0	0	0	0		
1	0	0	1	0	0	0	1	1	1	1	1	0	0	0	-1	-2	-1	-1	-1	-1	0	1	1	1	0	0	1	0	0		
1	1	1	1	1	1	1	1	1	-1	-2	0	7	16	20	24	29	30	32	33	25	14	3	-1	0	1	1	1	1	1		
1	1	1	1	1	1	1	-1	0	14	33	50	49	44	42	38	38	35	34	33	37	45	47	27	1	0	1	1	1	1		
1	1	1	0	1	1	-1	-11	45	51	39	37	33	33	49	61	57	60	70	61	49	46	33	46	51	9	-1	1	1	1		
1	1	1	1	1	-1	18	46	33	34	53	63	77	74	79	84	74	82	83	85	73	74	68	37	40	58	16	-1	1	1		
1	1	1	1	-1	18	40	37	61	70	87	89	88	87	83	83	87	86	86	88	73	81	78	72	35	33	54	6	0	1	1	
1	1	1	0	8	39	39	70	87	86	89	90	89	89	79	79	89	84	87	89	85	76	79	66	55	36	44	30	-2	1	1	1
1	1	1	0	38	38	71	84	87	84	81	80	82	84	83	92	92	85	83	84	79	75	84	70	81	56	31	45	4	0	1	1
1	1	0	12	43	54	82	83	84	76	80	89	91	90	93	94	94	91	83	80	82	90	90	86	76	78	46	39	22	-1	1	1
1	1	-2	28	42	64	83	81	75	82	93	91	87	88	84	79	77	78	93	94	83	83	87	92	88	73	74	36	41	2	1	1
1	1	-1	35	46	79	83	83	78	92	90	67	82	87	87	85	85	70	67	90	95	89	92	89	76	70	86	50	37	16	-1	1
1	0	7	52	44	73	87	82	83	88	91	81	87	94	94	91	92	91	81	69	93	96	86	81	75	86	85	63	35	31	-1	1
1	-1	15	63	26	75	88	88	86	85	90	90	85	92	95	93	91	90	89	84	93	89	81	86	90	90	81	83	38	37	1	1
1	0	7	61	20	27	58	77	83	83	81	79	84	84	84	91	92	87	80	88	85	86	86	88	90	88	85	46	39	4	1	1
1	1	0	38	74	21	19	41	43	46	51	67	67	68	66	79	90	73	71	79	78	83	86	89	90	90	87	85	53	36	3	1
2	1	4	17	36	30	27	10	7	10	12	44	106	96	74	71	83	77	82	86	81	78	78	80	87	82	82	83	44	33	1	1
0	4	25	41	40	30	36	13	5	3	5	12	53	106	87	55	85	89	87	94	88	85	82	78	75	74	76	57	39	21	-1	1
13	29	49	57	37	38	35	26	12	14	29	16	20	64	68	48	67	85	86	92	92	91	87	83	80	73	40	26	48	3	1	3
50	50	66	77	50	39	21	25	27	15	33	50	50	57	39	37	39	53	76	87	89	89	88	85	84	56	28	55	27	-1	0	19
26	20	23	46	44	38	40	39	47	50	40	35	72	77	62	54	41	52	83	85	86	86	85	75	52	37	76	57	1	0	1	6
-1	29	58	64	39	37	43	36	40	41	30	19	59	76	68	55	44	58	72	73	72	65	54	52	52	73	76	20	-1	1	1	1
-1	27	73	73	29	47	50	47	72	66	36	45	58	69	65	45	51	61	89	61	60	62	75	71	75	73	58	1	1	2	2	1
-1	27	72	36	26	54	69	75	77	79	79	73	70	70	62	46	48	63	64	57	70	72	69	64	73	77	28	-2	1	1	1	1
0	23	46	46	70	74	71	71	72	73	74	73	66	64	61	42	38	59	65	57	69	67	63	69	68	67	5	1	1	1	1	1
-1	44	68	39	61	62	61	63	65	69	75	73	73	67	62	47	45	46	49	60	69	61	63	67	69	40	-2	1	1	1	1	1
-1	20	82	33	44	64	63	62	61	62	73	73	62	67	62	50	57	53	46	56	60	56	61	65	67	20	-1	1	1	1	1	1
1	5	56	42	20	60	70	64	60	61	64	68	41	50	63	47	43	45	38	53	50	51	61	61	55	8	0	1	1	1	1	1
-1	18	58	39	27	51	67	60	55	60	59	60	35	29	58	48	44	46	40	54	53	49	59	56	42	2	1	1	1	1	1	1
-1	24	57	35	23	46	57	57	56	58	55	55	37	31	57	47	39	39	40	51	52	49	56	53	30	-1	1	1	1	1	1	1
0	8	44	30	48	50	49	50	49	43	49	47	26	49	37	33	40	34	45	46	46	47	44	22	0	2	1	1	1	1	1	1
0	0	5	17	21	24	25	24	27	26	26	23	13	12	27	23	21	24	26	28	28	30	25	25	13	0	1	0	0	0	0	0

No spatial information




## Using spatial features



**Idea:** connect patches of input to neurons in hidden layer




## Using spatial features



- Slide patch window across input image
- Weight pixels inside the patch
- Apply weighted summation

→ Convolution


## Using spatial features



- Slide patch window across input image
- Weight pixels inside the patch
- Apply weighted summation

→ Convolution

## Using spatial features



- Slide patch window across input image
  - Weight pixels inside the patch
  - Apply weighted summation
- Convolution

## The convolution operation

- Slide the filter over the input image
- Element-wise multiply
- Add the outputs

1	1	1	0	0
0	1	1	1	0
0	0	1	1	1
0	0	1	1	0
0	1	1	0	0

Image



1	0	1
0	1	0
1	0	1

Filter

## The convolution operation

- Slide the filter over the input image
- Element-wise multiply
- Add the outputs

$$\begin{aligned} & 1 \times 1 + 1 \times 0 + 1 \times 1 \\ & + 0 \times 0 + 1 \times 1 + 1 \times 0 \\ & + 0 \times 1 + 0 \times 0 + 1 \times 1 \\ & = 4 \end{aligned}$$

1 <sub>x1</sub>	1 <sub>x0</sub>	1 <sub>x1</sub>	0	0
0 <sub>x0</sub>	1 <sub>x1</sub>	1 <sub>x0</sub>	1	0
0 <sub>x1</sub>	0 <sub>x0</sub>	1 <sub>x1</sub>	1	1
0	0	1	1	0
0	1	1	0	0

Image



1	0	1
0	1	0
1	0	1

Filter



4		

Feature map

## The convolution operation

- Slide the filter over the input image
- Element-wise multiply
- Add the outputs

$$\begin{aligned} & 1 \times 1 + 1 \times 0 + 0 \times 1 \\ & + 1 \times 0 + 1 \times 1 + 1 \times 0 \\ & + 0 \times 1 + 1 \times 0 + 1 \times 1 \\ & = 3 \end{aligned}$$

1	1 <sub>x1</sub>	1 <sub>x0</sub>	0 <sub>x1</sub>	0
0	1 <sub>x0</sub>	1 <sub>x1</sub>	1 <sub>x0</sub>	0
0	0 <sub>x1</sub>	1 <sub>x0</sub>	1 <sub>x1</sub>	1
0	0	1	1	0
0	1	1	0	0

Image



1	0	1
0	1	0
1	0	1

Filter



4	3	

Feature map

## The convolution operation

- Slide the filter over the input image
- Element-wise multiply
- Add the outputs

$$\begin{aligned} & 1 \times 1 + 1 \times 0 + 1 \times 1 \\ & + 1 \times 0 + 1 \times 1 + 0 \times 0 \\ & + 1 \times 1 + 0 \times 0 + 0 \times 1 \\ & = 4 \end{aligned}$$

1	1	1	0	0
0	1	1	1	0
0	0	1 <sub>x1</sub>	1 <sub>x0</sub>	1 <sub>x1</sub>
0	0	1 <sub>x0</sub>	1 <sub>x1</sub>	0 <sub>x0</sub>
0	1	1 <sub>x1</sub>	0 <sub>x0</sub>	0 <sub>x1</sub>

Image



1	0	1
0	1	0
1	0	1


Filter




4	3	4
2	4	3
2	3	4

Feature map


Different filters = different feature maps



Original image




Vertical edge detection



Horizontal edge detection




## Using spatial features




- Apply a set of weights - a filter - to extract **local features**
- Use **multiple filters** to extract different features

## CNNs for classification



## Convolutional layer




For a neuron in hidden layer:

- Take inputs from patch
- Compute weighted sum
- Apply bias
- Activate with non-linear function

Neuron  $(p, q)$  in hidden layer  
Filter size  $4 \times 4$   
Weights  $w_{i,j}$

$$h \left( \sum_{i=1}^4 \sum_{j=1}^4 w_{i,j} x_{i+p, j+q} + b \right)$$

## Spatial arrangement of output volume



### Layer Dimensions:

$$h \times w \times d$$

$h$  &  $w$  = spatial dimensions

$d$  = number of filters

### Stride:


Filter step size

### Receptive Field:

Locations in input image that a node is connected to


## Introducing non-linearity

### Rectified linear unit (ReLU)




$$h(z) = \max(0, z)$$

Input feature map



Rectified feature map



Black: negative values - White: positive values

## Pooling

- Reduce dimensionality while preserving spatial invariance

Input feature map

1	1	8	3
4	7	1	9
5	3	1	4
2	3	6	0

Max pooling with  
2×2 filter and stride 2




Pooled feature map

7	9
5	6

## Pooling

- Reduce dimensionality while preserving spatial invariance

Input feature map



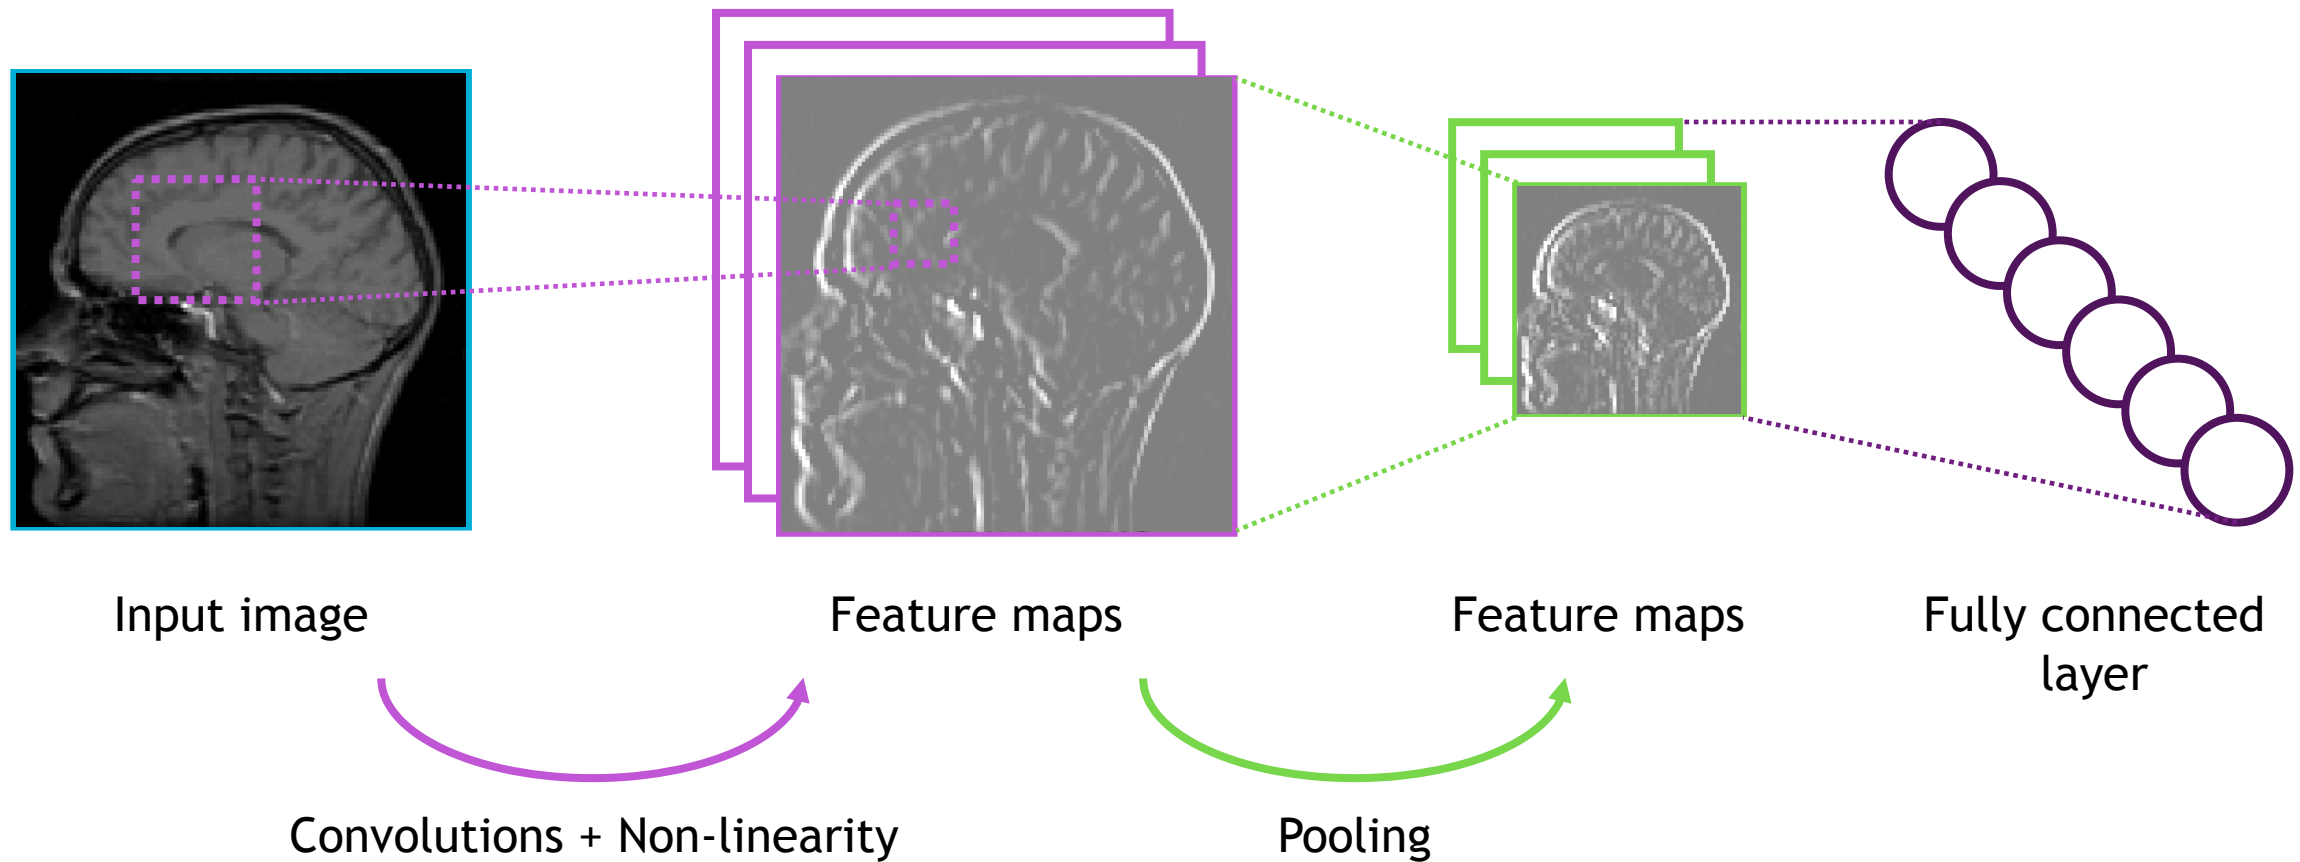
Max pooling with  
 $2 \times 2$  filter and stride 2



Pooled feature map

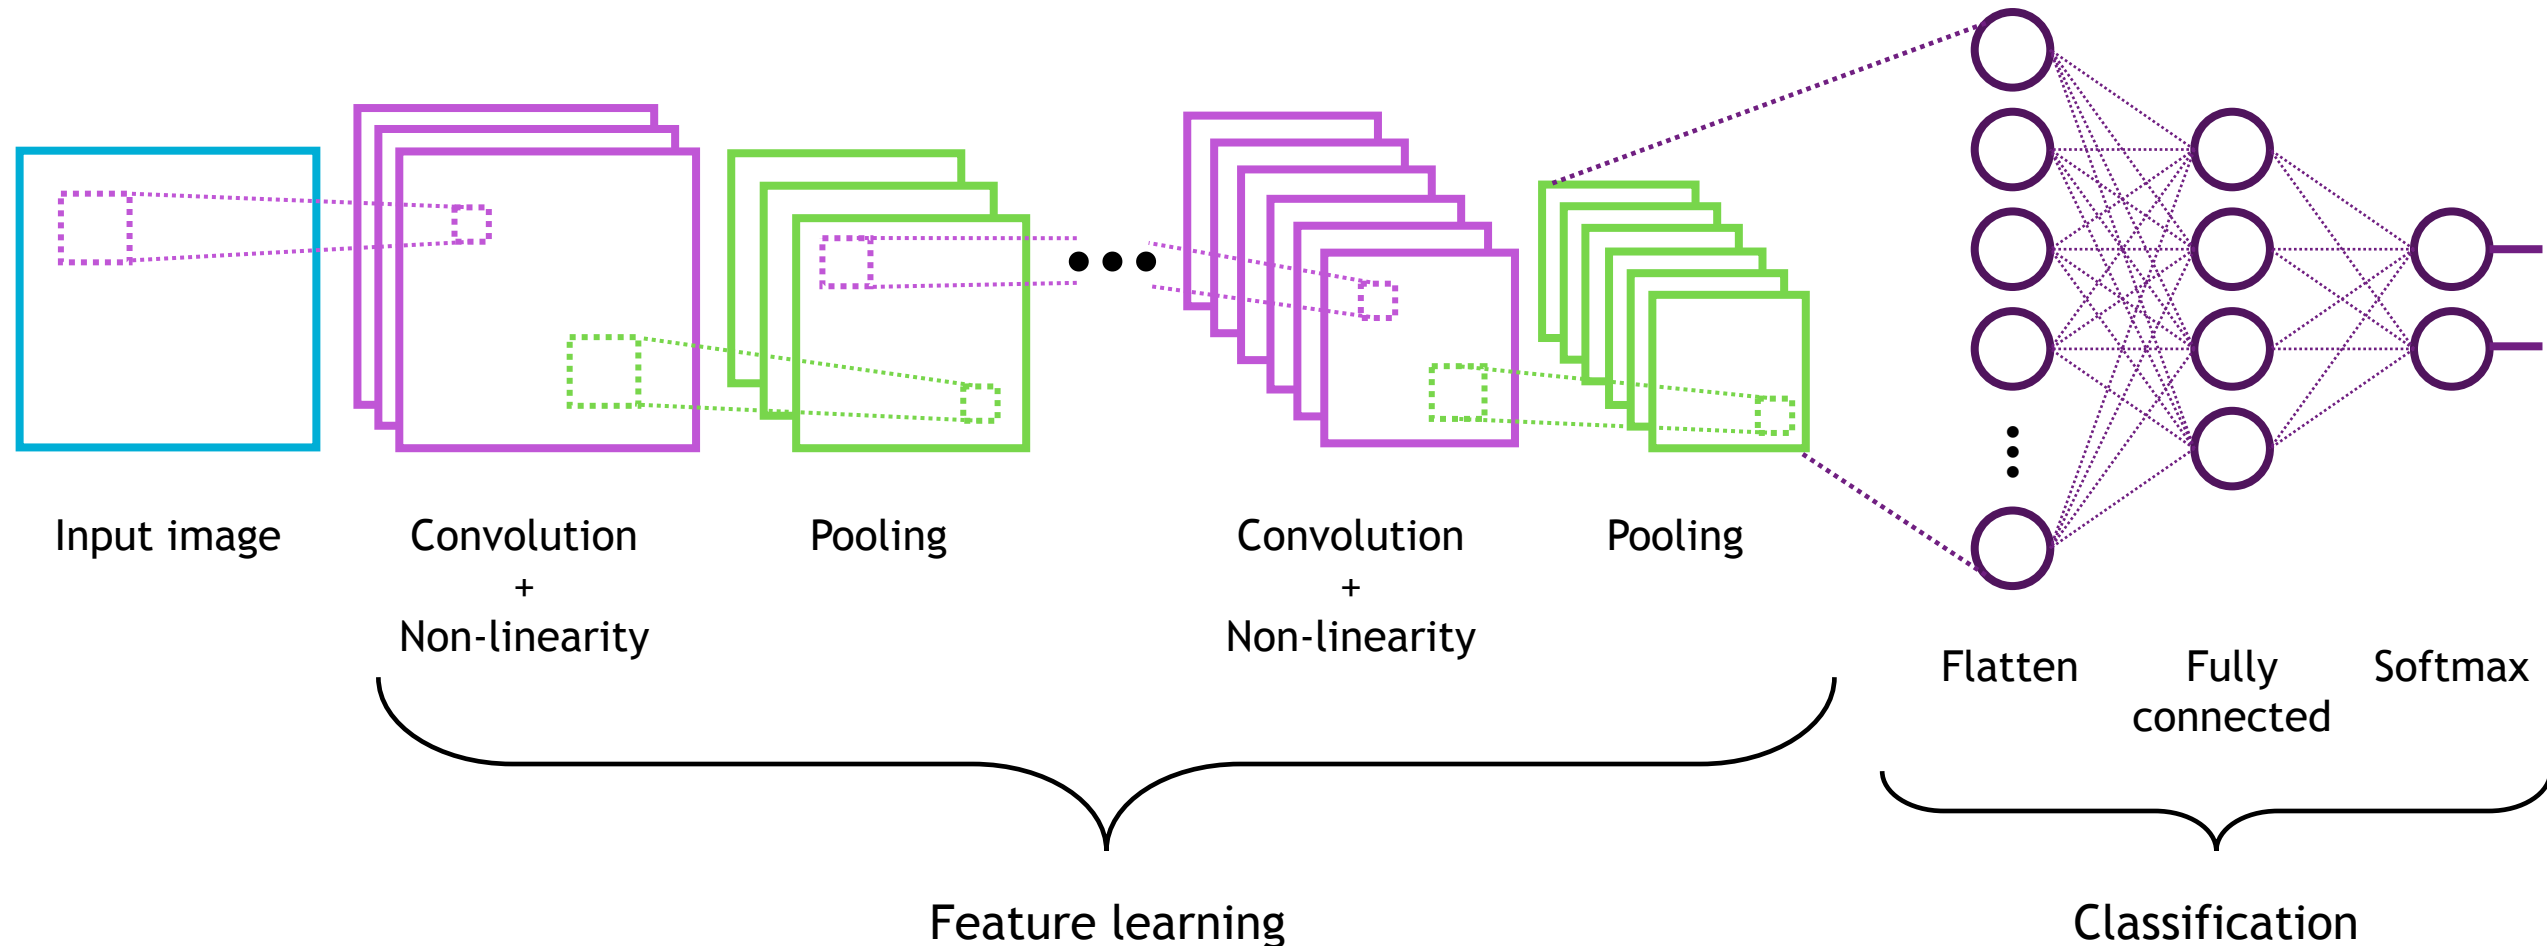


## CNNs for classification

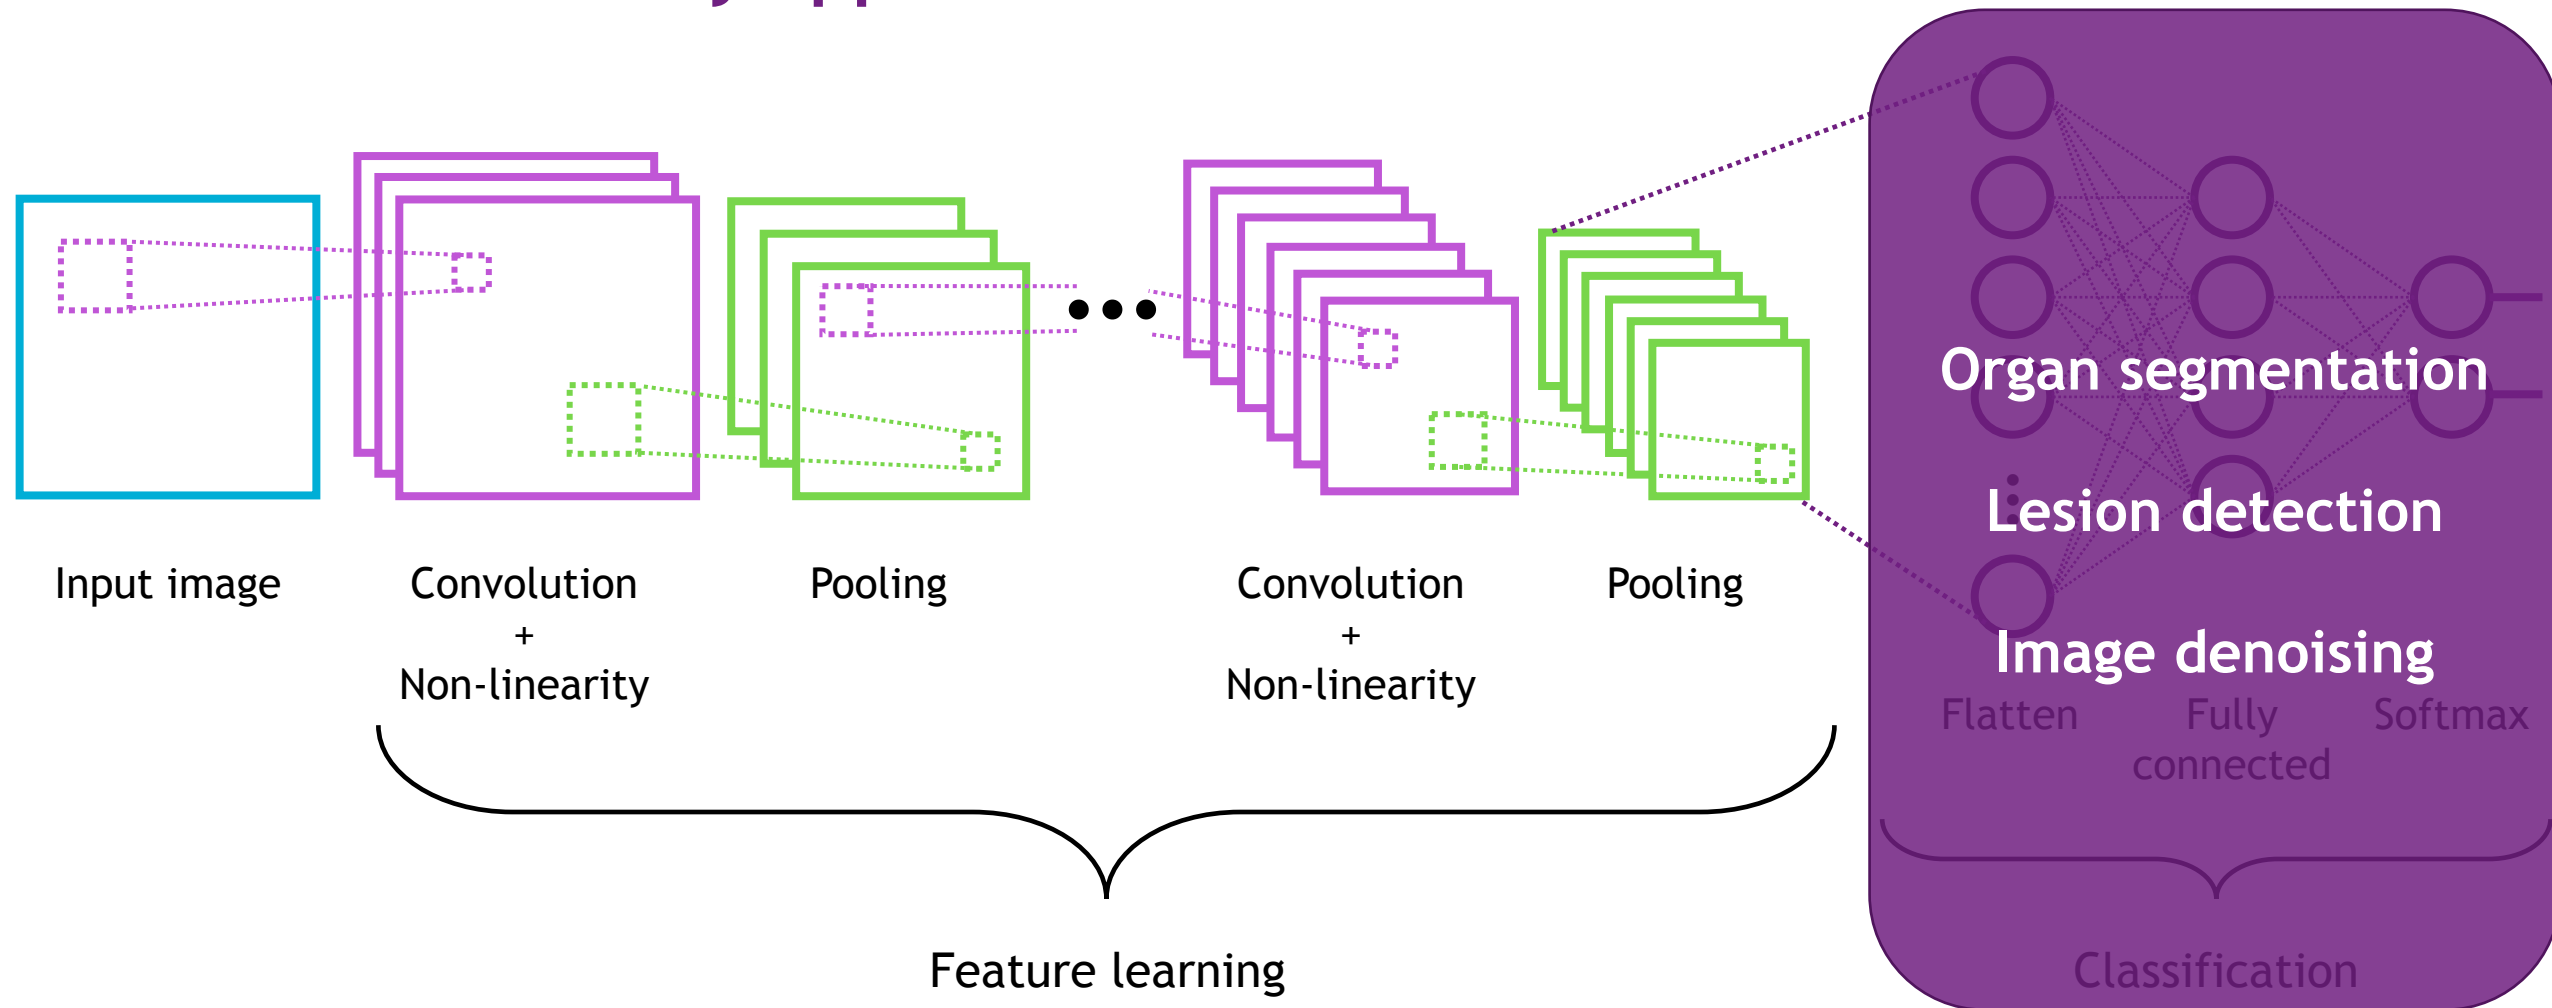




## CNNs for classification



## CNNs for many applications



## Images

- Representing images
- Convolutions for feature extraction




## CNNs

- Convolution → non-linearity → pooling
- Stacking layers

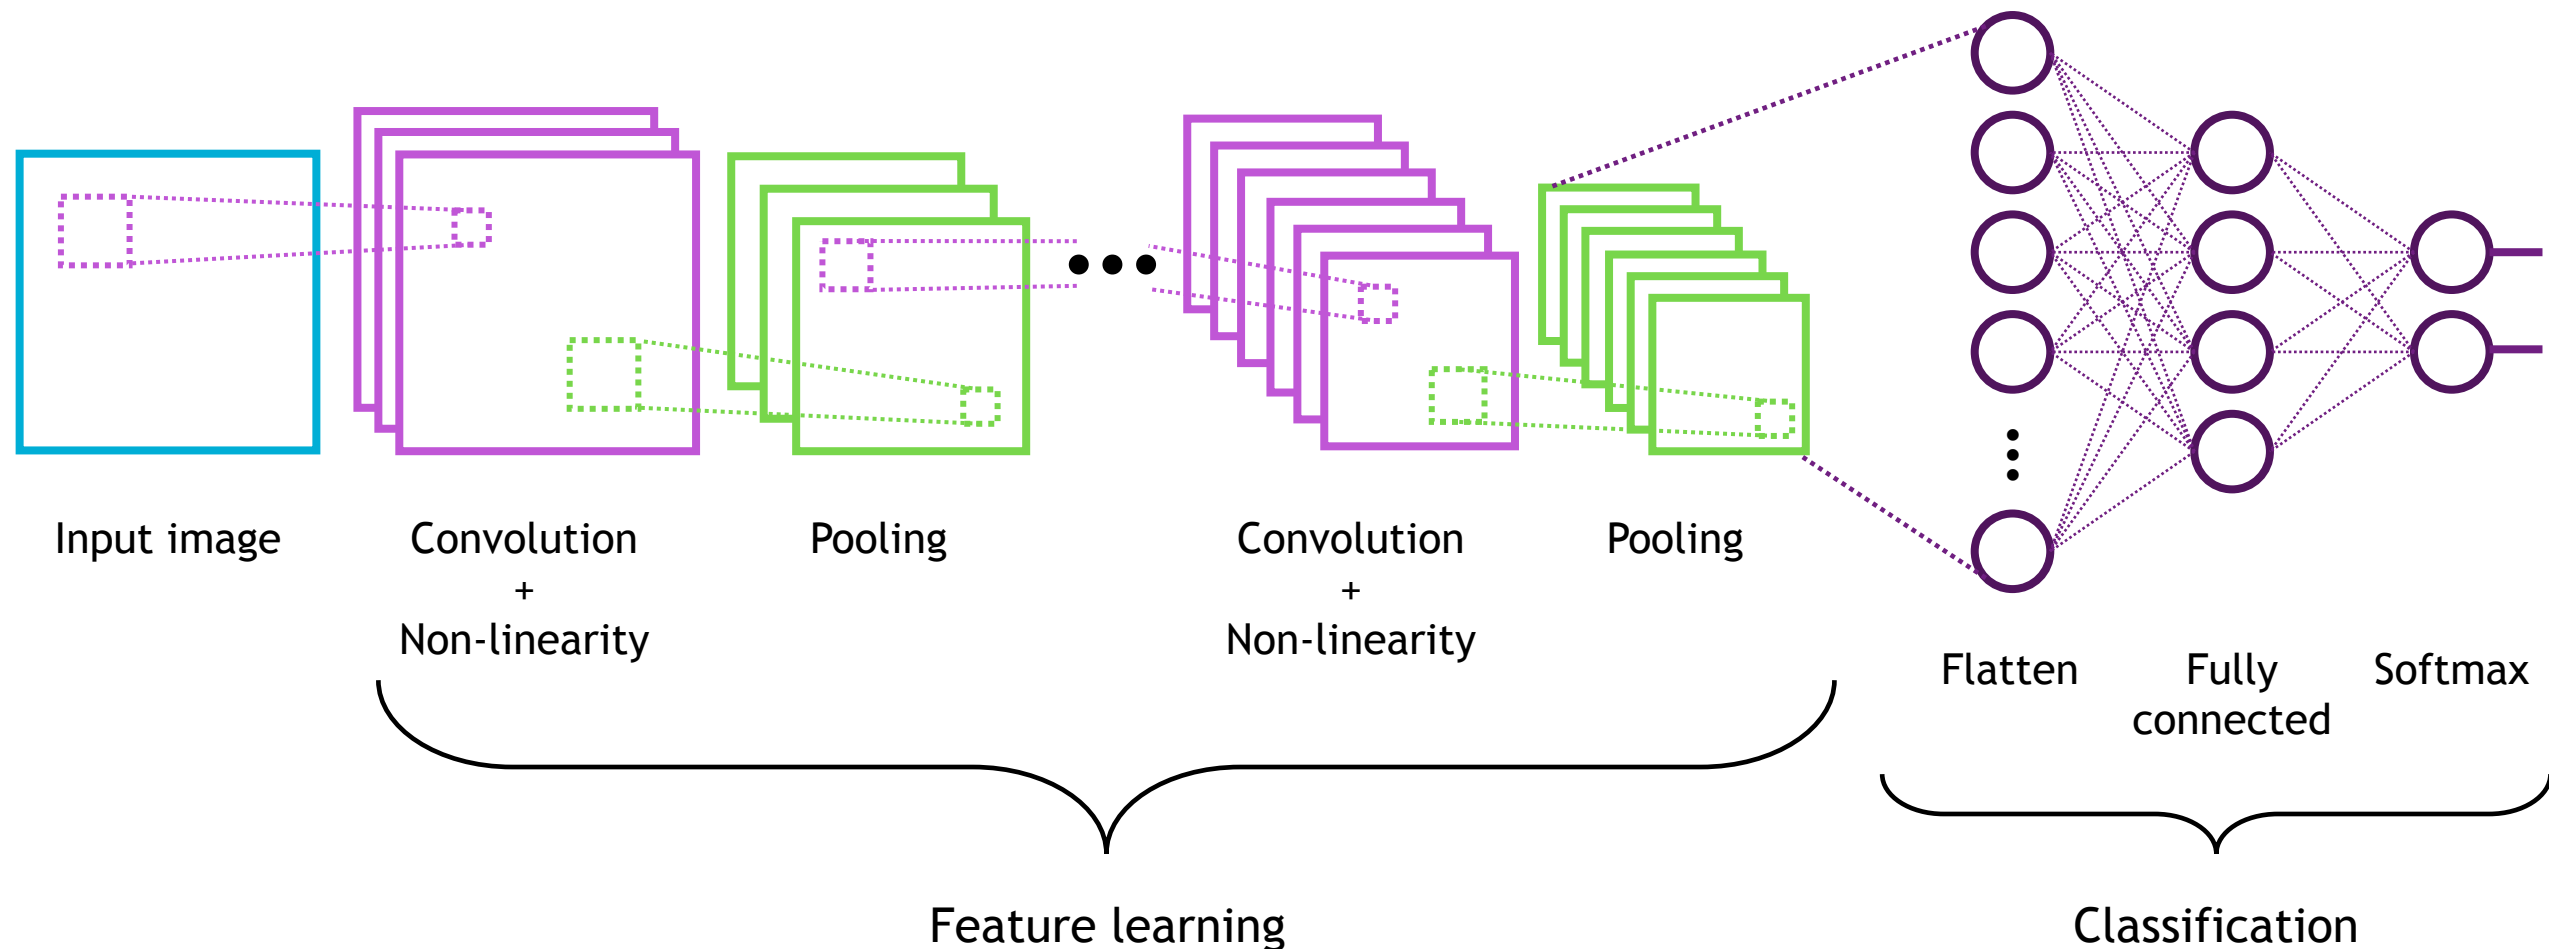
## Applications

- Classification
- Segmentation
- Detection




# Generative Deep Learning

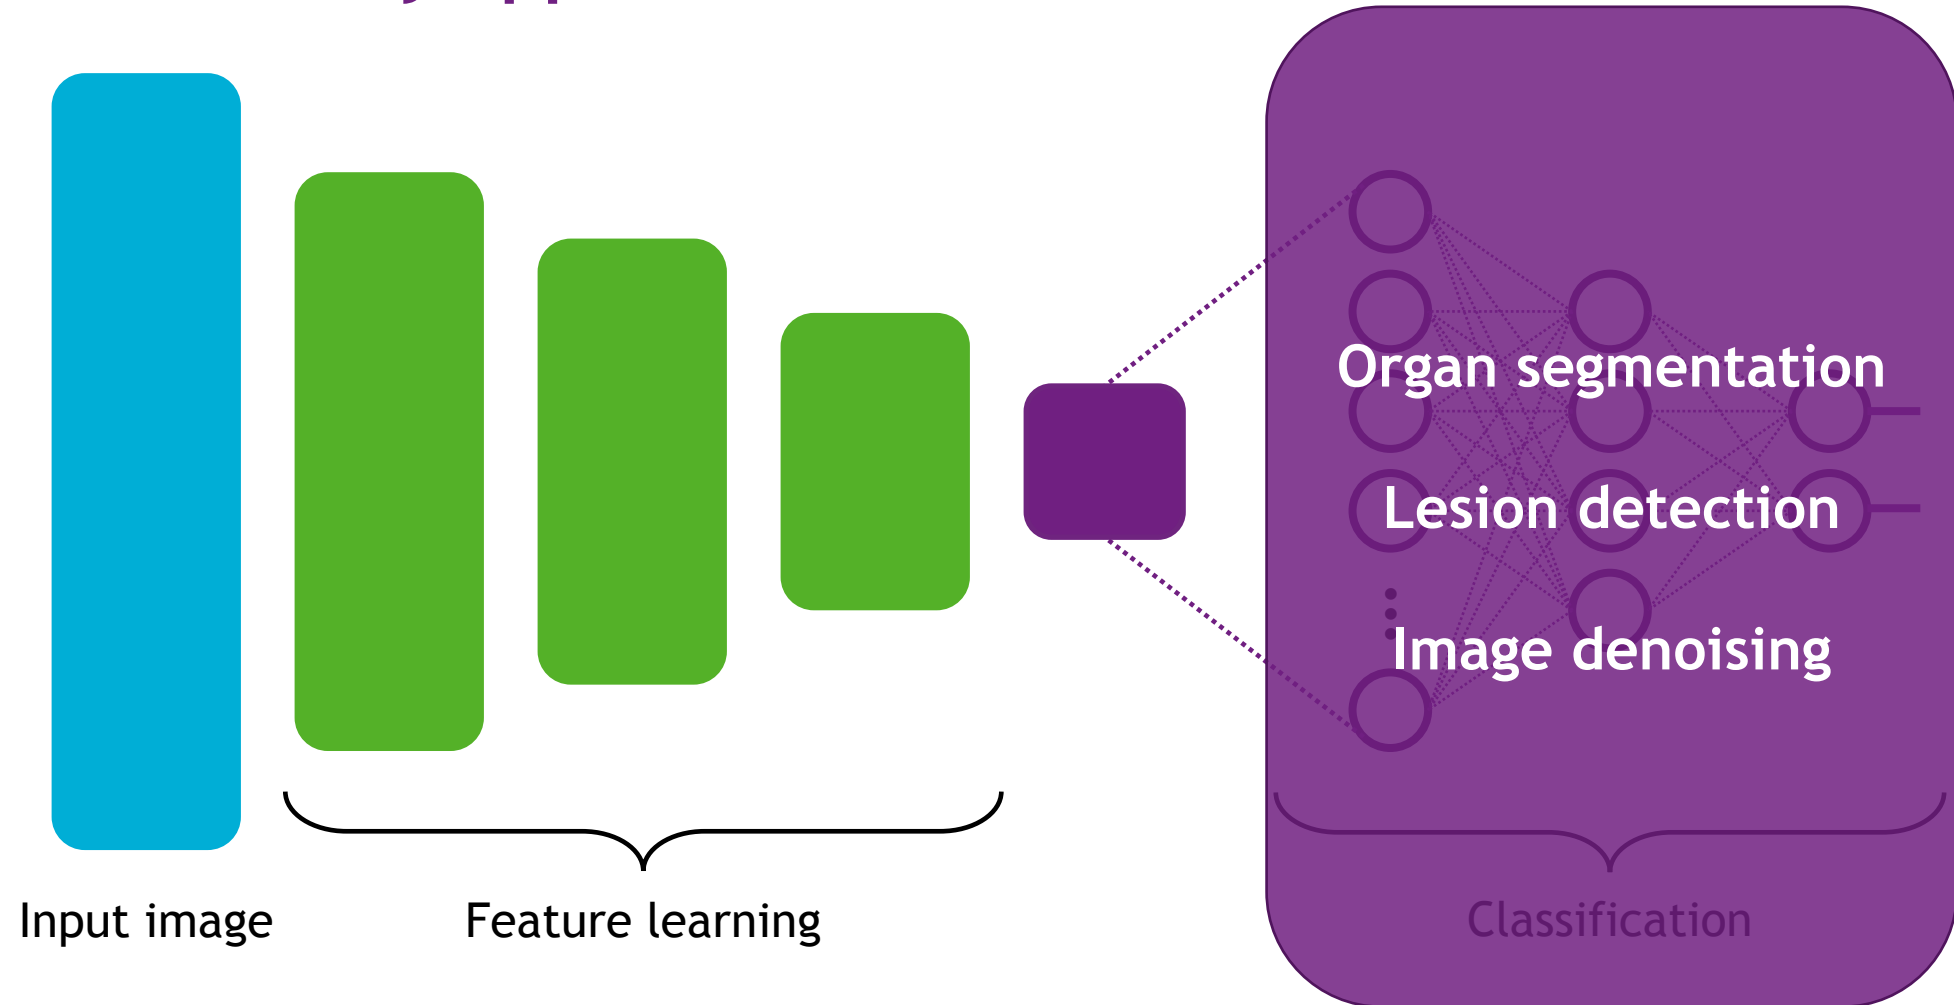
## CNNs for classification




## CNNs for classification



## CNNs for many applications



## CNNs for image generation







Observed variable



Latent variable




## Encoder




Input image  
=  
observed data

Latent space  
=  
low dimensional representation  
of the observed data


## Training autoencoders




$$\mathcal{L}(x, \hat{x}) = \|x - \hat{x}\|^2$$



# Autoencoders

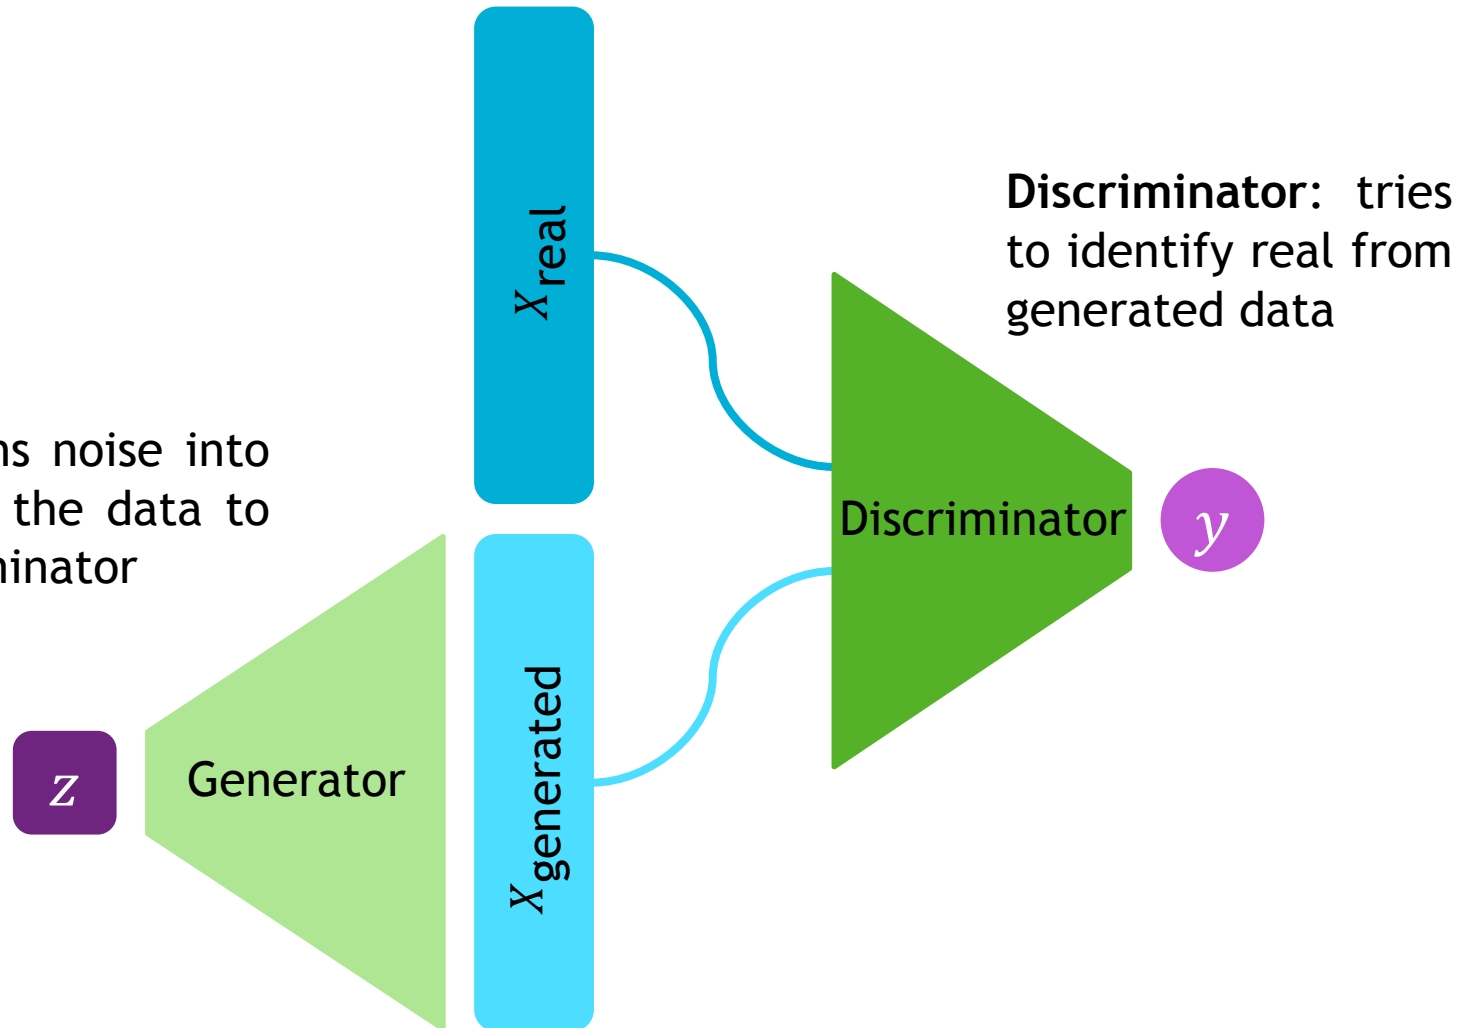


## Generating images from scratch

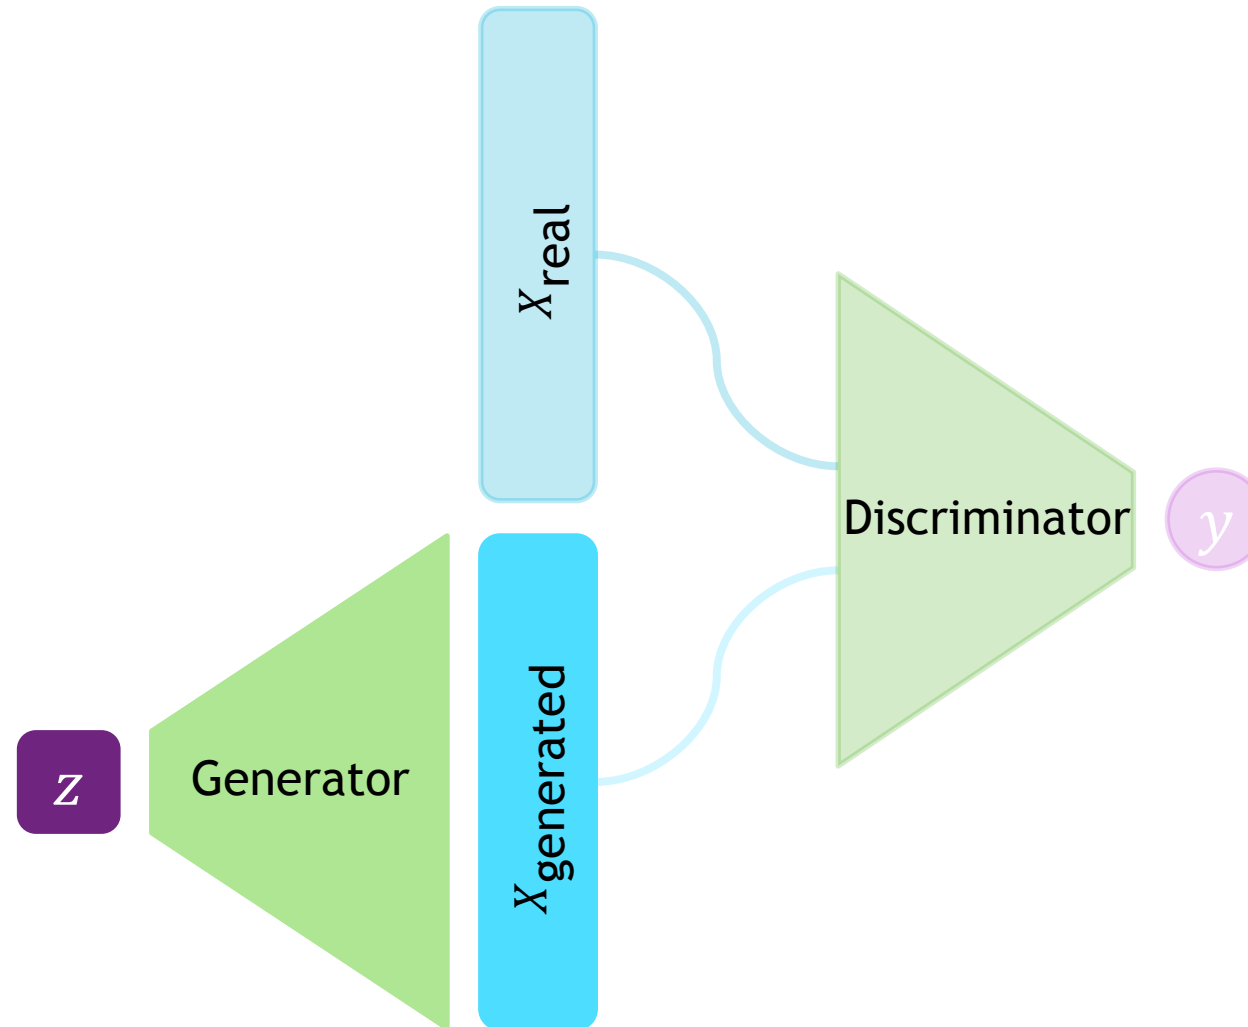


## Competing networks

**Generator:** turns noise into an imitation of the data to trick the discriminator




## Generating new images







## Image translation with conditional GANs



## Image translation with conditional GANs




## Image translation with conditional GANs




## Autoencoders

- Learn low dimensional latent space



## GANs

- Competing generator and discriminator networks



## Conditional GANs

- For image translation

