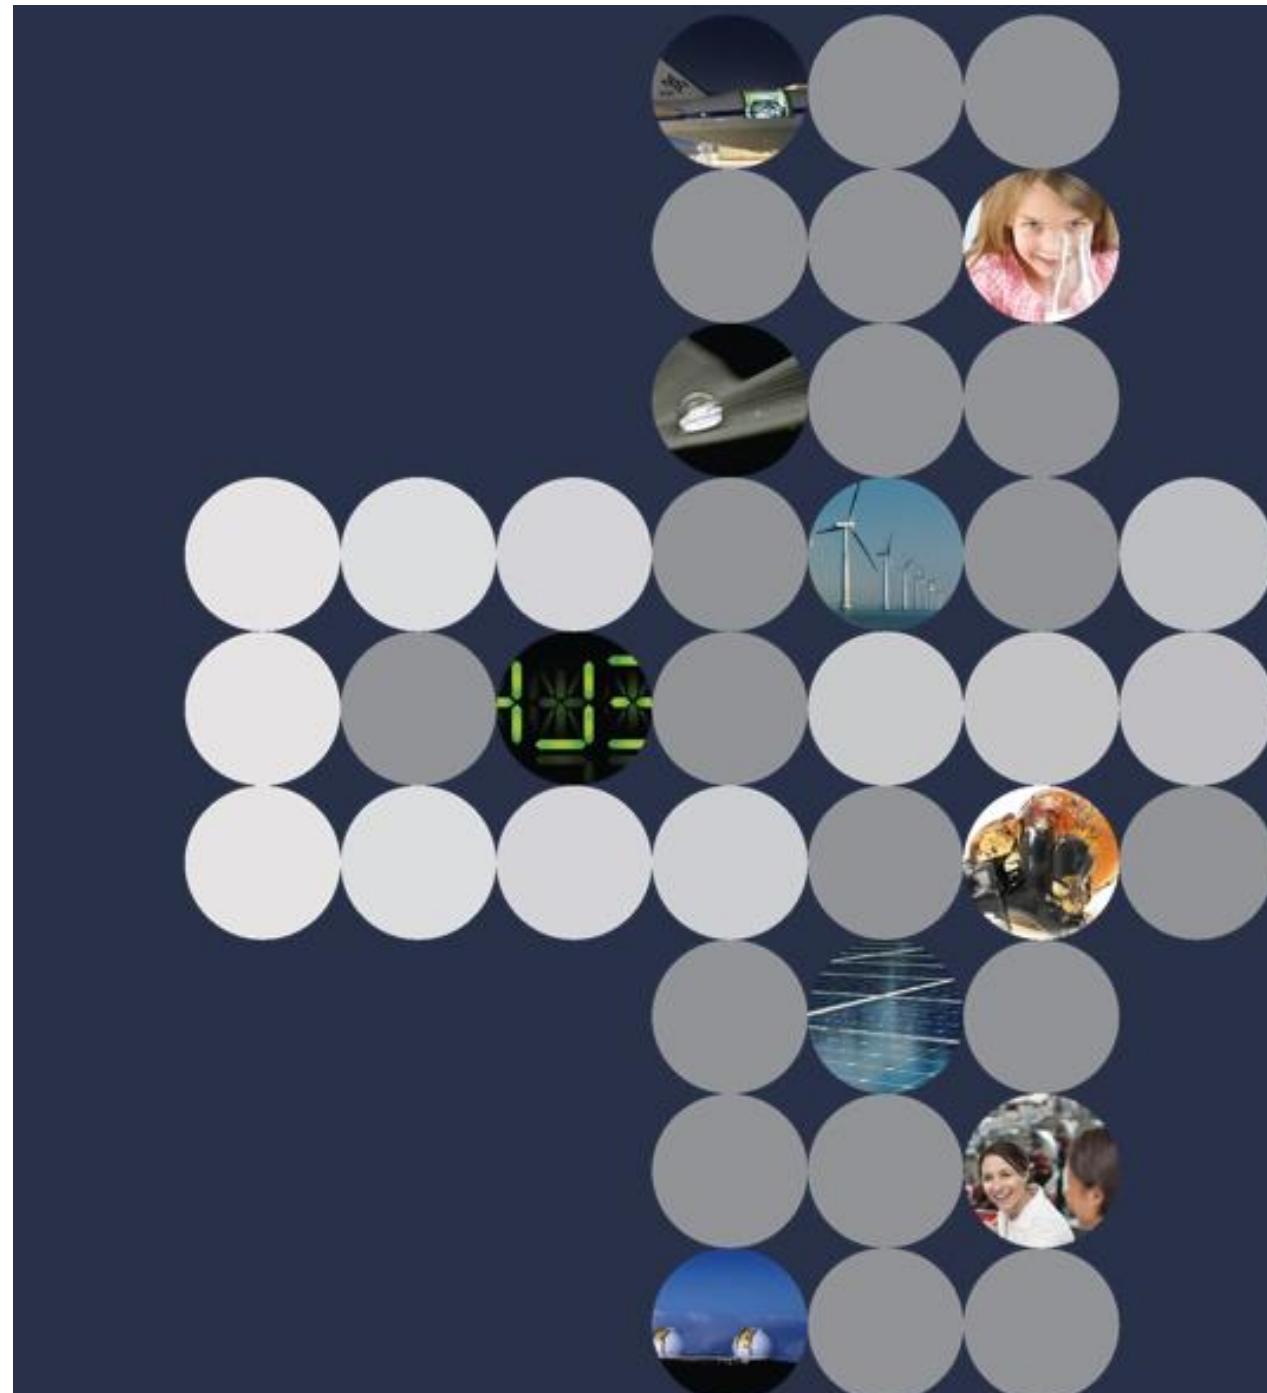


# EE512 – Applied Biomedical Signal Processing

## Neural network architecture

Clémentine AGUET  
CSEM Signal Processing & AI Group



# Communication

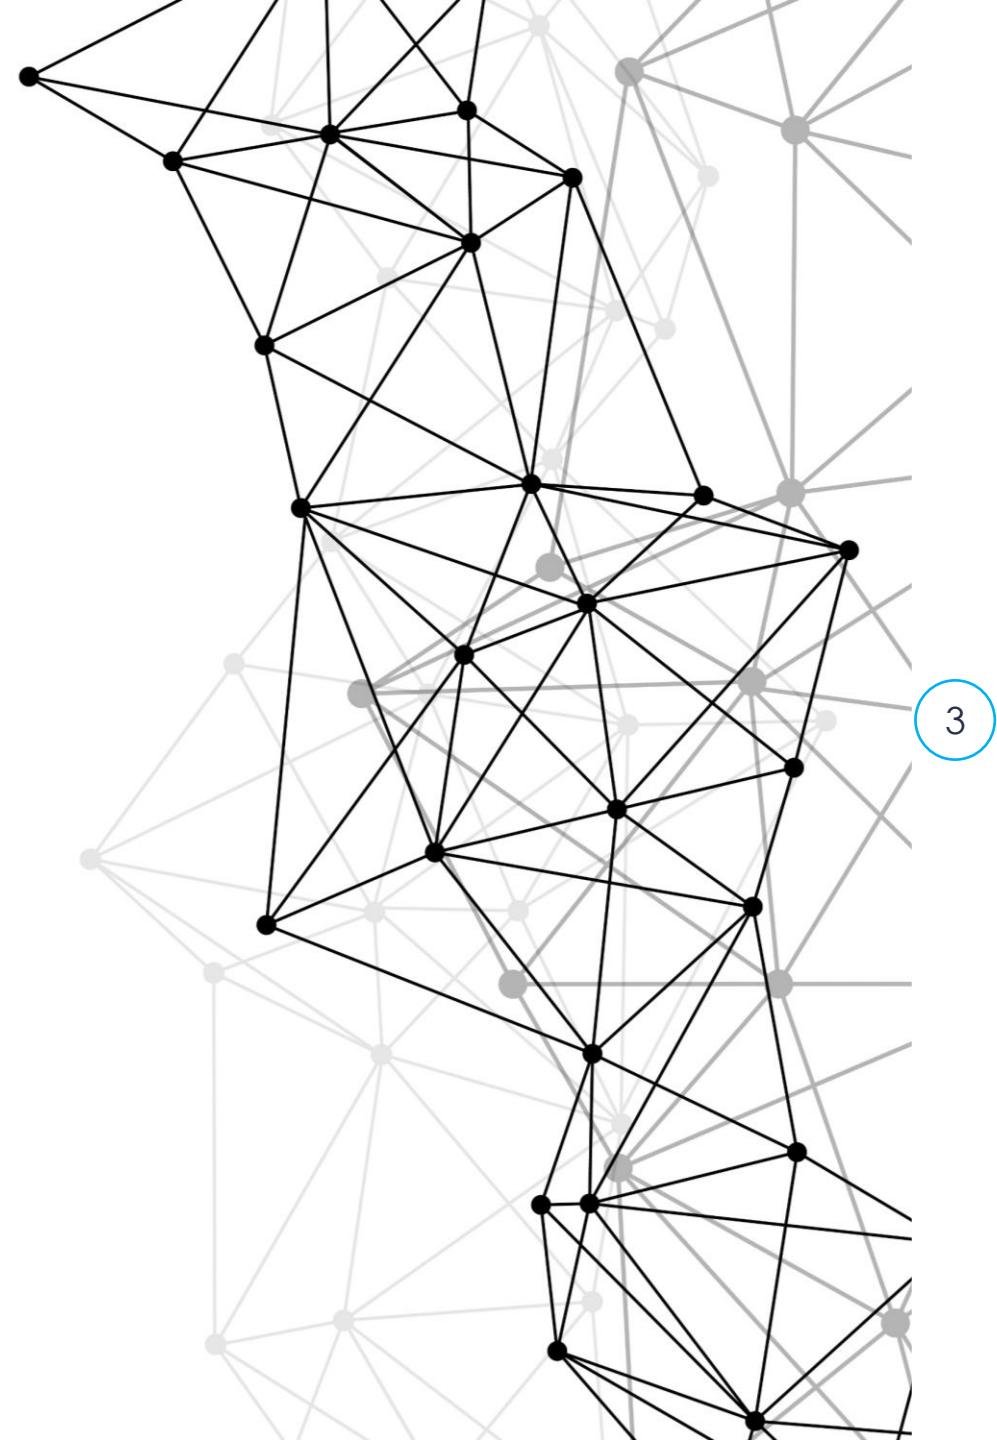
- Next week
  - No lecture nor lab but Q&A session.
  - We will be there to answer your questions on lectures, labs, etc.
- Exam
  - Question on each labs + 4 exercises (similar to mid-term).
  - Open book but no laptop nor smartphone.

2



# Content

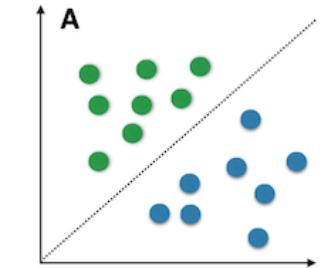
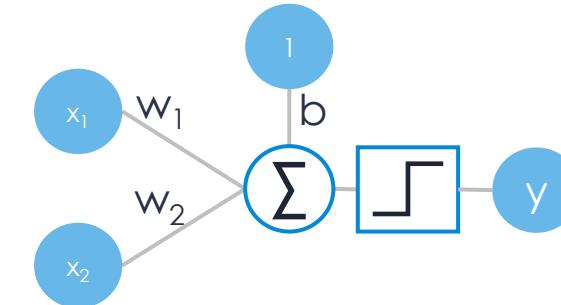
- Neural network recap
- Regularization
- Different architectures
- Labs



# Neural network recap

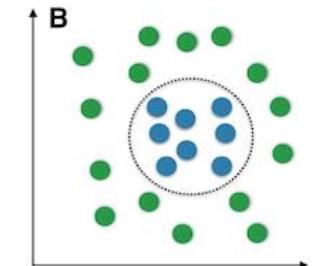
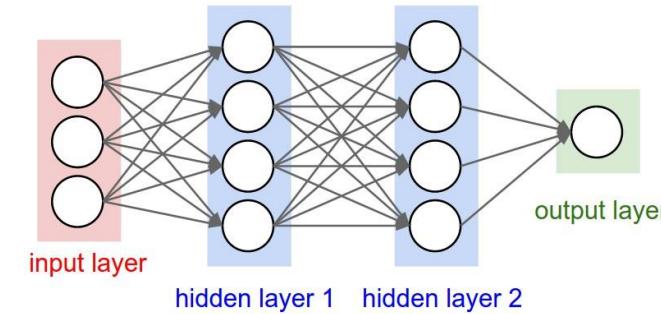
- **Perceptron**

- Neural network building block
- Linear binary classifier



- **Multilayer Perceptron (MLP)**

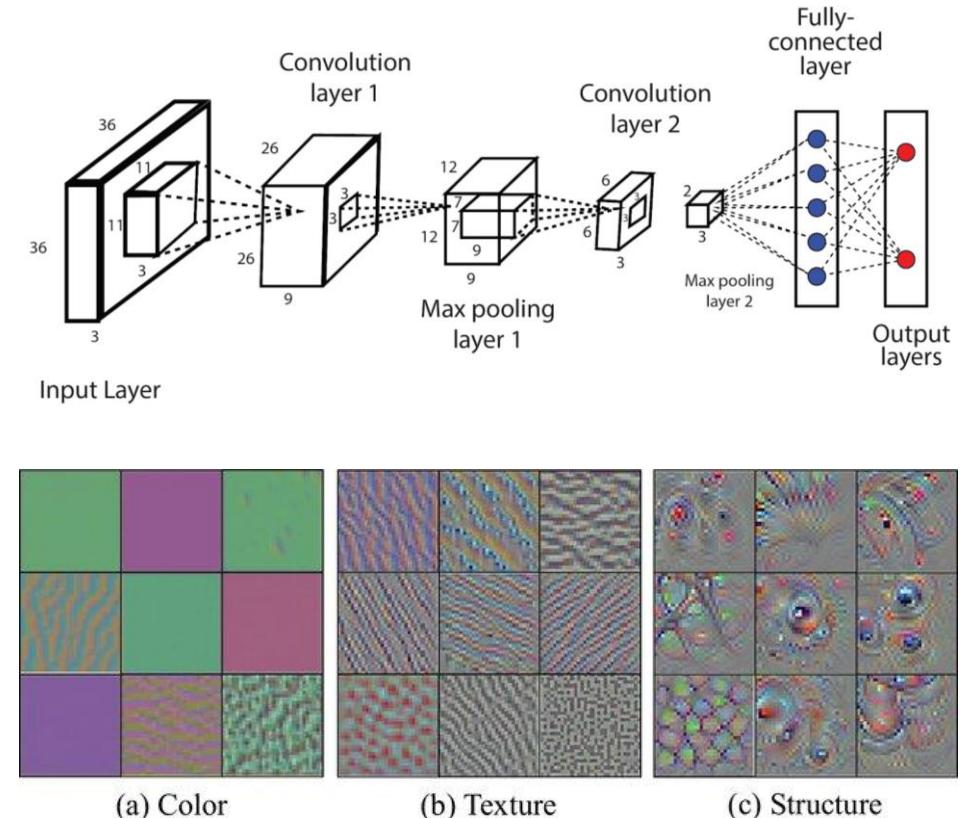
- Feedforward NN
- Solve non-linearly separable problems
- Each node in a layer is connected to all nodes in next layer
- Series of fully connected layers
- Each connection has a weight
- Deep models = MLP with many layers



4

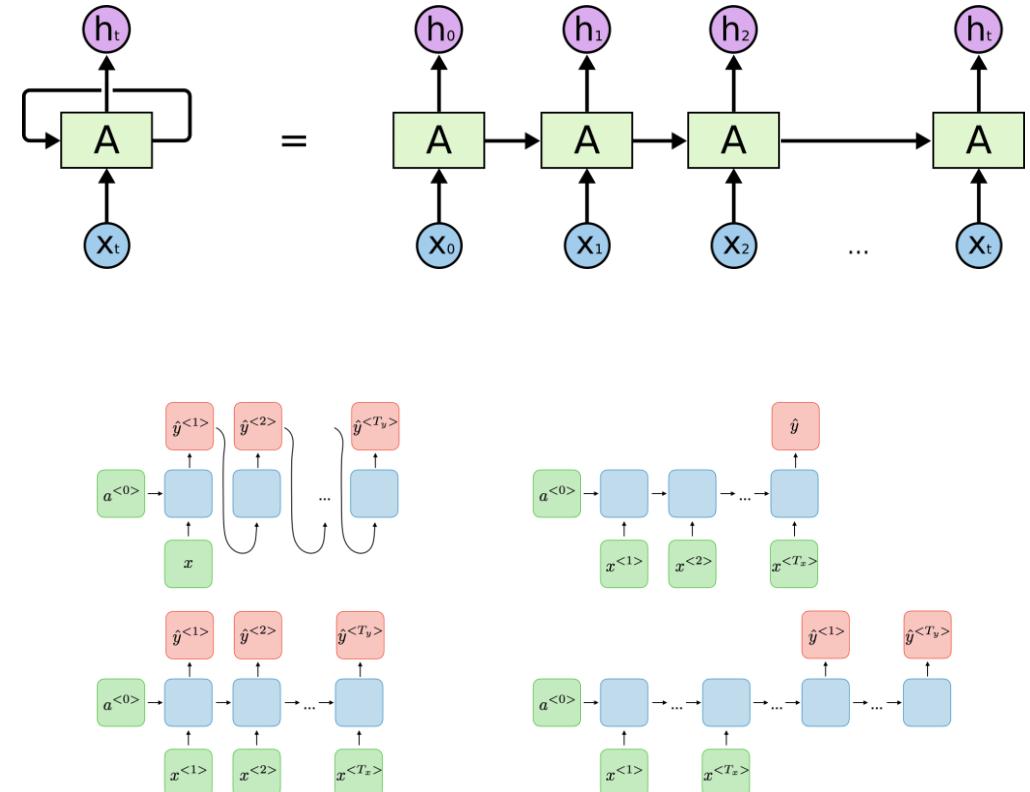
# Neural network recap

- **Convolutional neural network (CNN)**
  - Computer vision
  - Extract features hierarchically
  - **Convolutional layers**
    - Extract features
    - Apply set of learnable filters (kernel)
    - Local connectivity and parameter sharing
  - **Pooling layers**
    - Dimensionality reduction
      - Reduce number of parameters and memory usage
    - Invariance to transformations
  - **Fully connected layers**
    - Flattened feature maps to 1D
    - Final classification or regression task



# Neural network recap

- **Recurrent neural network (RNN)**
  - Natural language processing, speech recognition
  - Sequential input data
  - Capture temporal representation
  - Weights are shared across time
  - Each neuron has an internal memory
  - Long Short-Term Memory units (LSTM)
  - Gated Recurrent Unit (GRU)



# Neural network recap

How neural networks are trained?

- Find the network weights and bias that minimize the empirical loss

$$W^*, b^* = \operatorname{argmin}_{W,b} L(W, b)$$

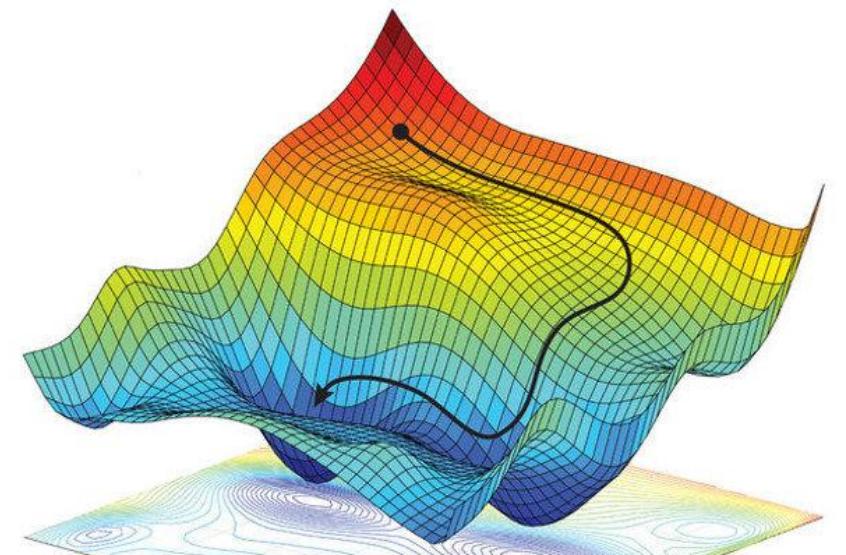
- Calculate cost function
- Compute gradient through backpropagation

$$\frac{\partial L(W,b)}{\partial W} \text{ and } \frac{\partial L(W,b)}{\partial b}$$

- Update parameters

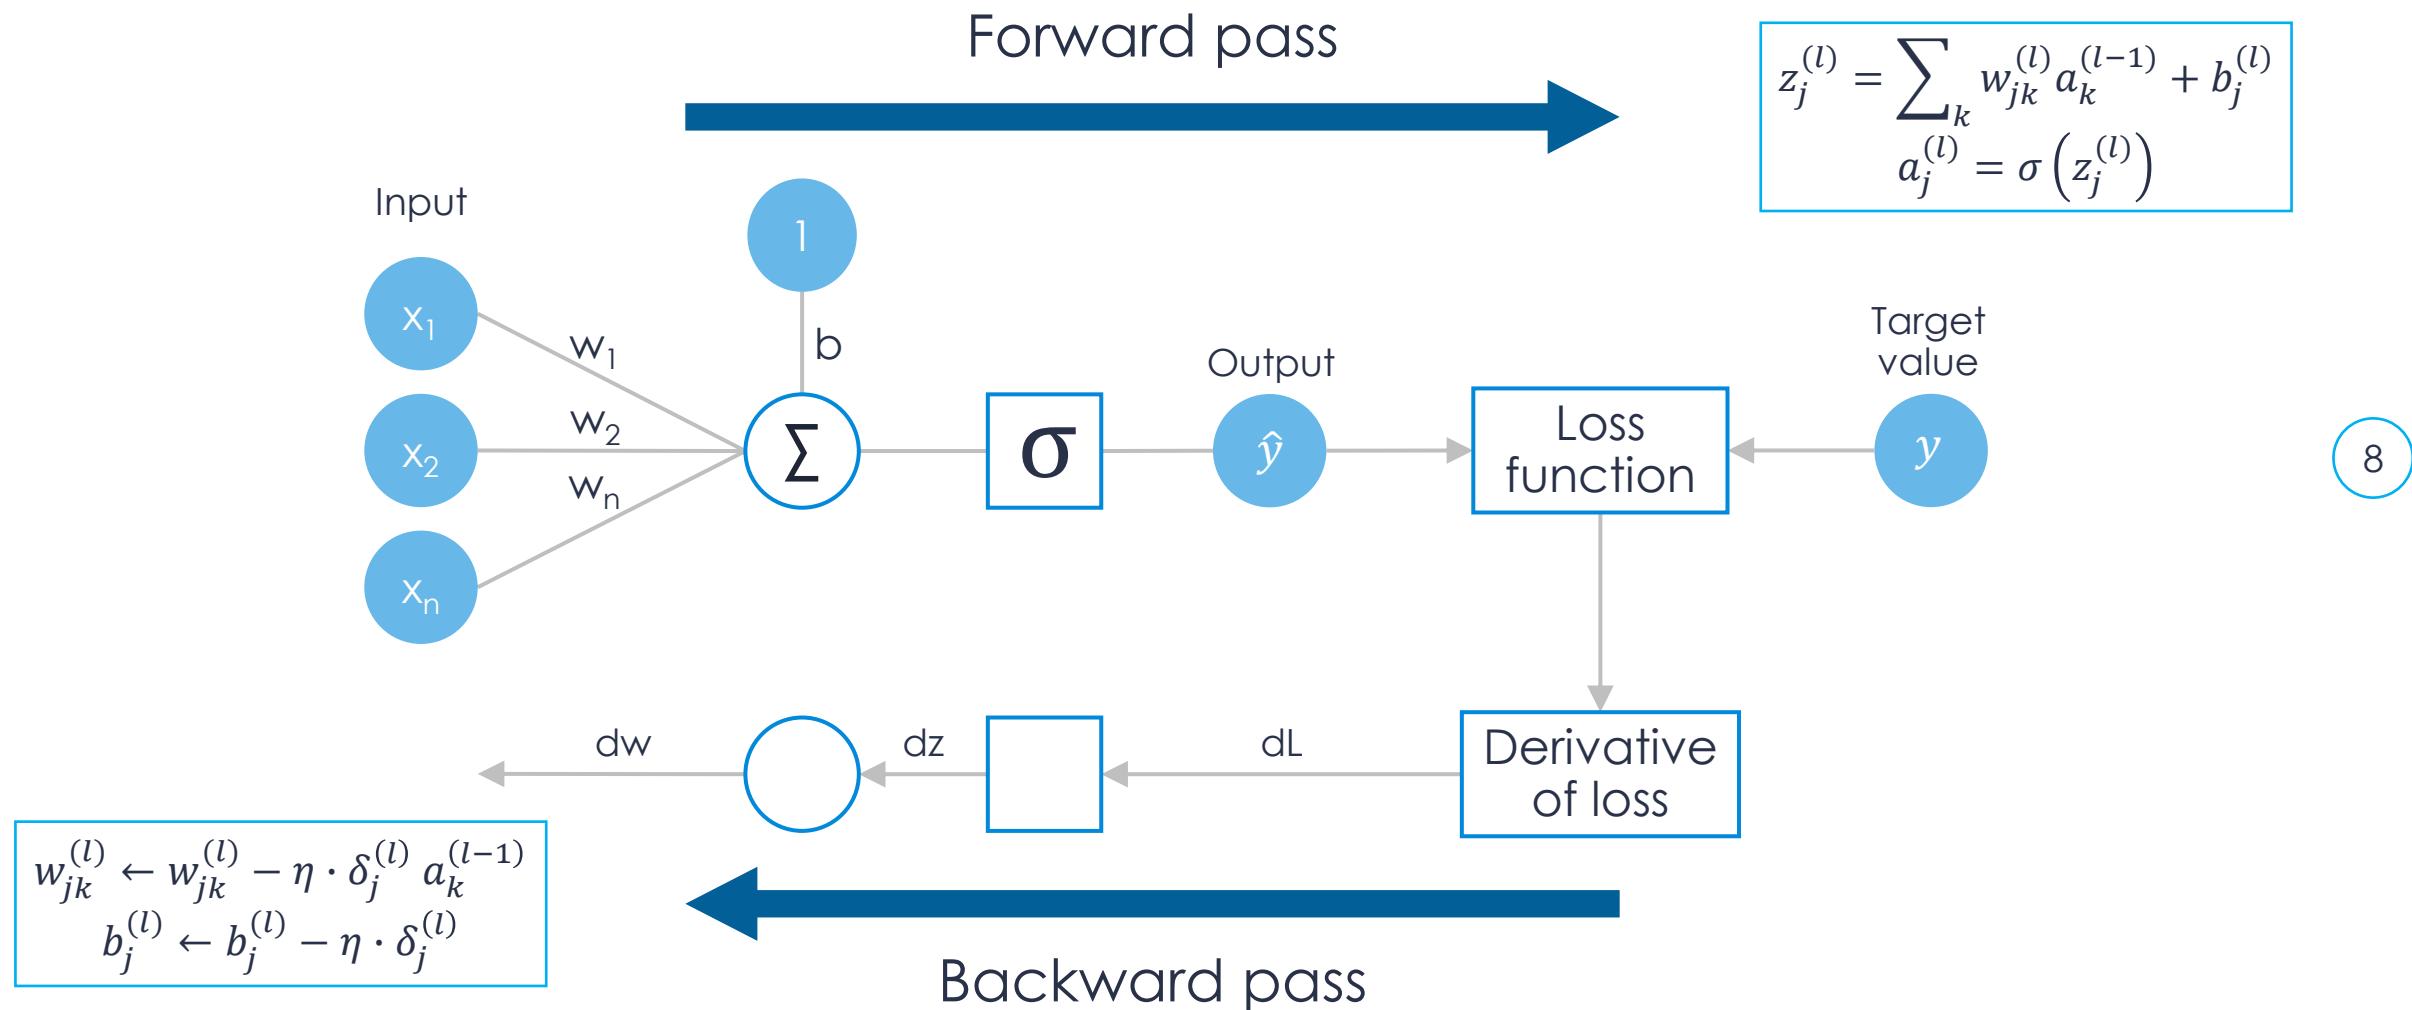
$$W \leftarrow W - \eta \frac{\partial L(W,b)}{\partial W}$$

$$b \leftarrow b - \eta \frac{\partial L(W,b)}{\partial b}$$



7

# Neural network recap – Compute gradient



# Neural network recap

- Applications



Speech recognition



Autonomous driving cars



Natural language processing

9



Stock market prediction



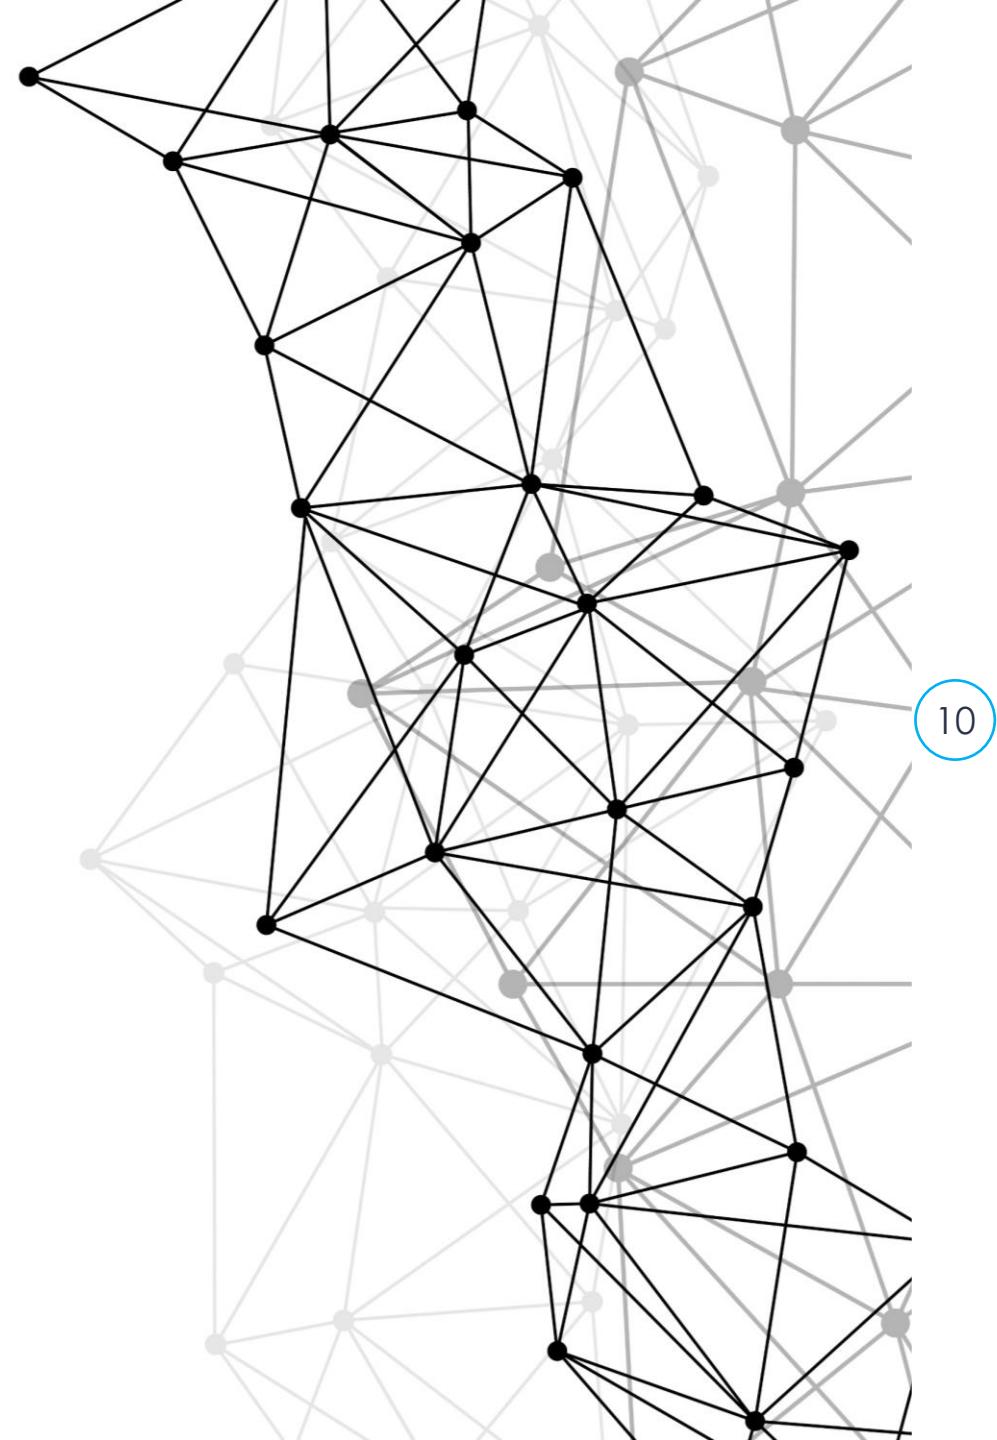
Fraud detection



Healthcare

# Content

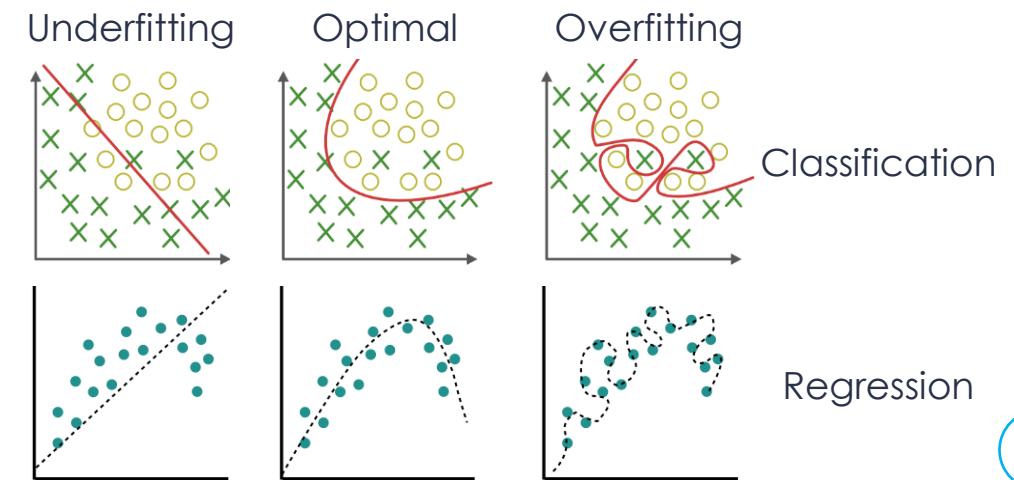
- Neural network recap
- Regularization
- Different architectures
- Labs



# Regularization

Why?

- Neural networks are prone to **overfitting**
- Model does not generalize well on unseen data
- Low training error but high validation error

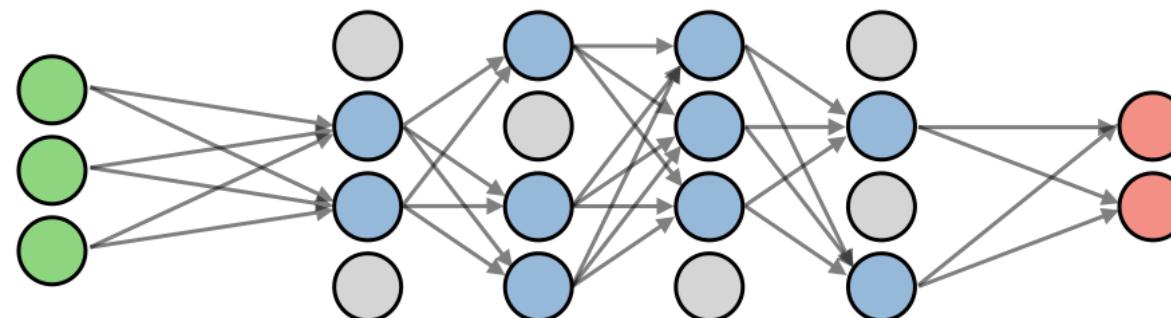


- **Regularization**
  - Slight modifications to learning algorithm such that the model generalizes better
  - Lower the complexity
  - Combine multiple techniques

# Regularization – Dropout

- Avoid relying too much on particular nodes → Learn robust and generalized features
- At every **training** iteration, randomly remove some nodes
- Each iteration has a different set of nodes
- Each node has a probability  $p > 0$  to be turned off
- $p$  is an hyperparameter (typically between 0.2 and 0.5)
- Full capacity during inference → All nodes used to make predictions
- More efficient in fully connected layers than convolutional layers

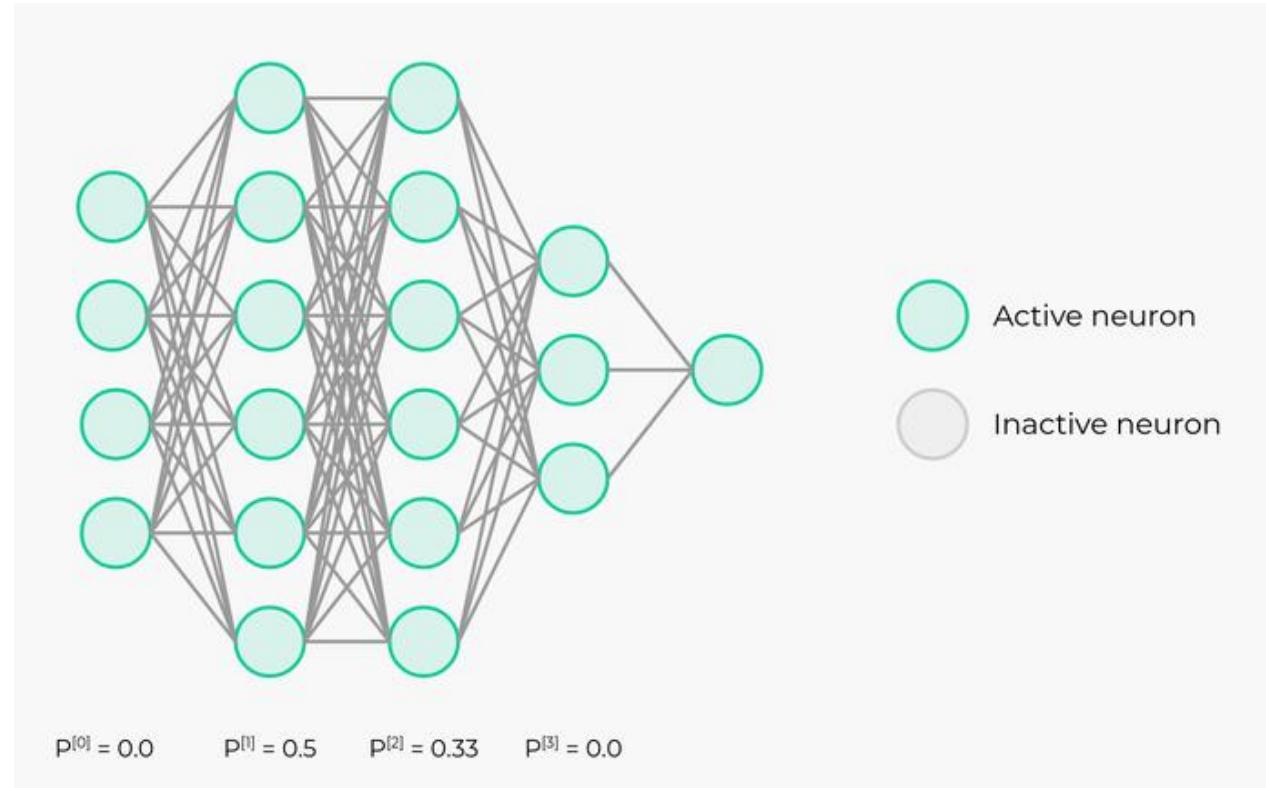
12



- PyTorch: `torch.nn.Dropout (p=0.5)`

# Regularization – Dropout

- + Efficient in large networks
- + Easy to implement
- + Computationally efficient
- + Robustness in feature learning
- Stochasticity → Slower convergence
- Not as effective in small networks
- Hyperparameter to tune



# Regularization – L1 & L2 regularization

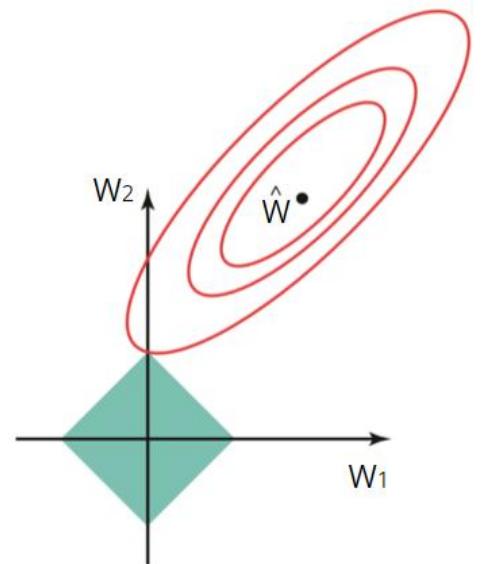
- Add regularization term to the cost function
- Used as penalty

$$\text{Loss} = \text{Loss} + \text{Regularization term}$$

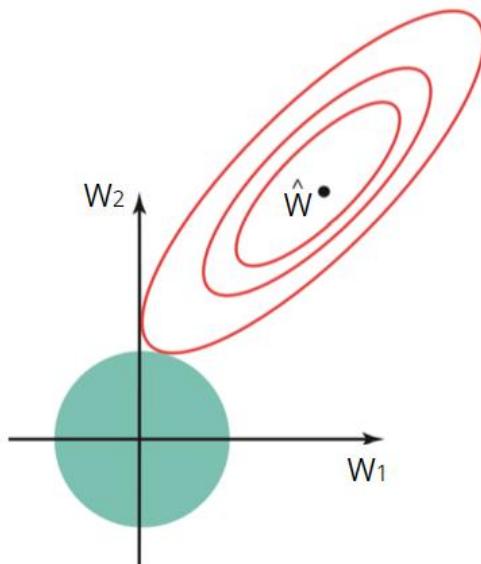
- Ensure that the weights are not too large
- Smaller weights lead to simpler models → Avoid overfitting

14

L1 regularization



L2 regularization



# Regularization – L1 & L2 regularization

- **L2 regularization**

- Weight decay or Ridge regression
- Regularization term  $\Omega(W)$  is defined as Euclidean norm or L2 norm

$$\Omega(W) = \|W\|_2^2 = \sum_i \sum_j w_{ij}^2$$

- Regularization term is weighted by a scalar  $\lambda$  and divided by 2
- $\lambda$  is the regularization rate

$$\hat{L}(W, b) = L(W, b) + \frac{\lambda}{2} \|W\|_2^2$$

- Compute gradient

$$\frac{\partial \hat{L}(W, b)}{\partial W} = \frac{\partial L(W, b)}{\partial W} + \lambda W$$

- Weights update

$$W \leftarrow W - \eta \frac{\partial \hat{L}(W, b)}{\partial W} = (1 - \eta \lambda)W - \eta \frac{\partial L(W, b)}{\partial W}$$

# Regularization – L1 & L2 regularization

- **L2 regularization**
  - + General regularization in most NN setups
  - + Promoting smoothness → Penalizes large weights and generalizable model
  - + Differentiable penalty
  - + Computational efficiency
  - + Handling multicollinearity
  - + Unique solution
  - Hyperparameter to tune (regularization rate)
  - Lack of sparsity → No feature selection
  - Sensitivity to irrelevant features

# Regularization – L1 & L2 regularization

- **L1 regularization**

- Lasso regression
- Regularization term  $\Omega(W)$  is defined as L1 norm (Manhattan distance)

$$\Omega(W) = \|W\|_1 = \sum_i \sum_j |w_{ij}|$$

- Regularization term is weighted by a scalar  $\lambda$
- $\lambda$  is the regularization rate

$$\hat{L}(W, b) = L(W, b) + \lambda \|W\|_1$$

- Compute gradient

$$\frac{\partial \hat{L}(W, b)}{\partial W} = \frac{\partial L(W, b)}{\partial W} + \lambda sign(W)$$

- Weights update

$$W \leftarrow W - \eta \frac{\partial \hat{L}(W, b)}{\partial W} = W - \eta \lambda sign(W) - \eta \frac{\partial L(W, b)}{\partial W}$$

# Regularization – L1 & L2 regularization

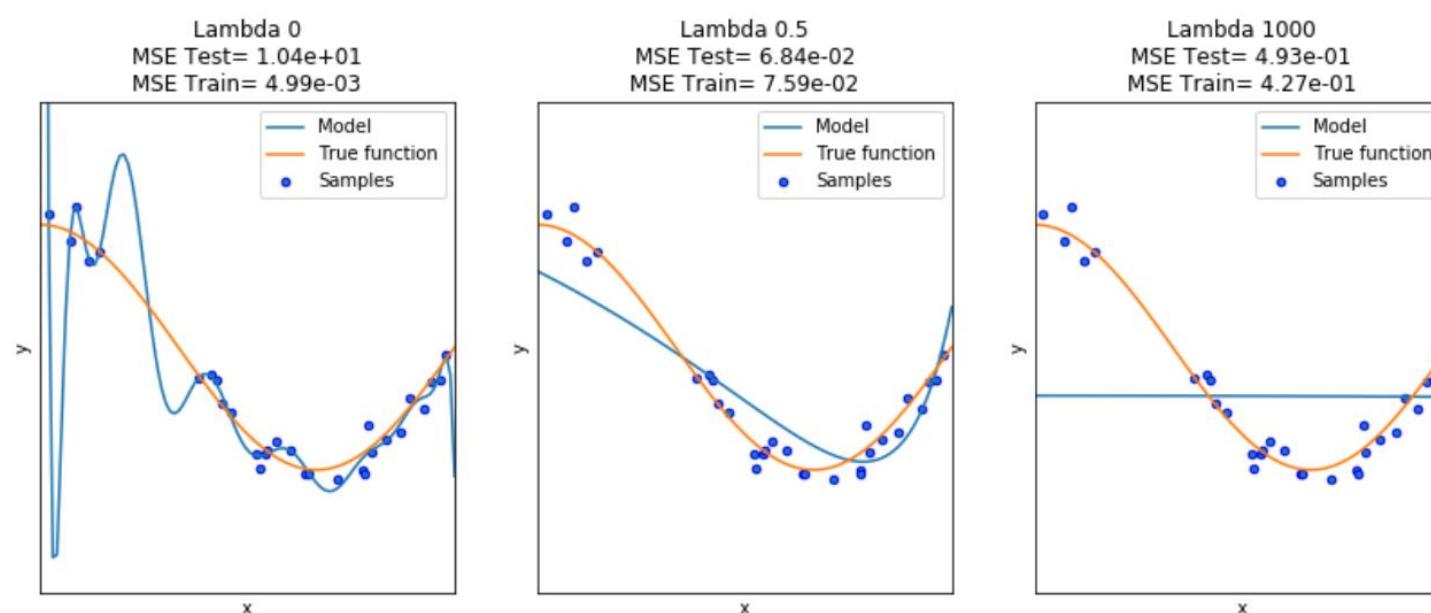
- **L1 regularization**

- + Feature selection through sparsity
- + Remove irrelevant features → Simpler model
- + Enhance interpretability
- + Robust to outliers
- + Handling high-dimensional data
- + Computational efficiency
- + Simple and widely applicable
- Hyperparameter to tune (regularization rate)
- Unstable feature selection
- Correlated features
- Non-differentiability → Complicate optimization (sub-gradient methods)

# Regularization – L1 & L2 regularization

- **Regularization rate ( $\lambda$ )**

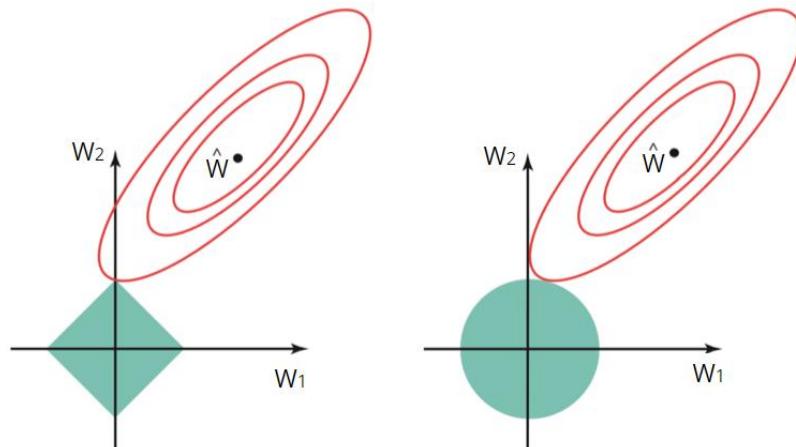
- Hyperparameter to tune
- Should be chosen carefully
  - Too high value → Model simpler but increased risk of underfitting
  - Too small value → Model more complex and increased risk of overfitting



# Regularization – L1 & L2 regularization

L1 regularization	L2 regularization
Sum of absolute value of weights	Sum of square of weights
Sparse solution	Non-sparse solution
Built in feature solution	No feature selection
Multiple solutions	One solution
Robust to outliers	Not robust to outliers

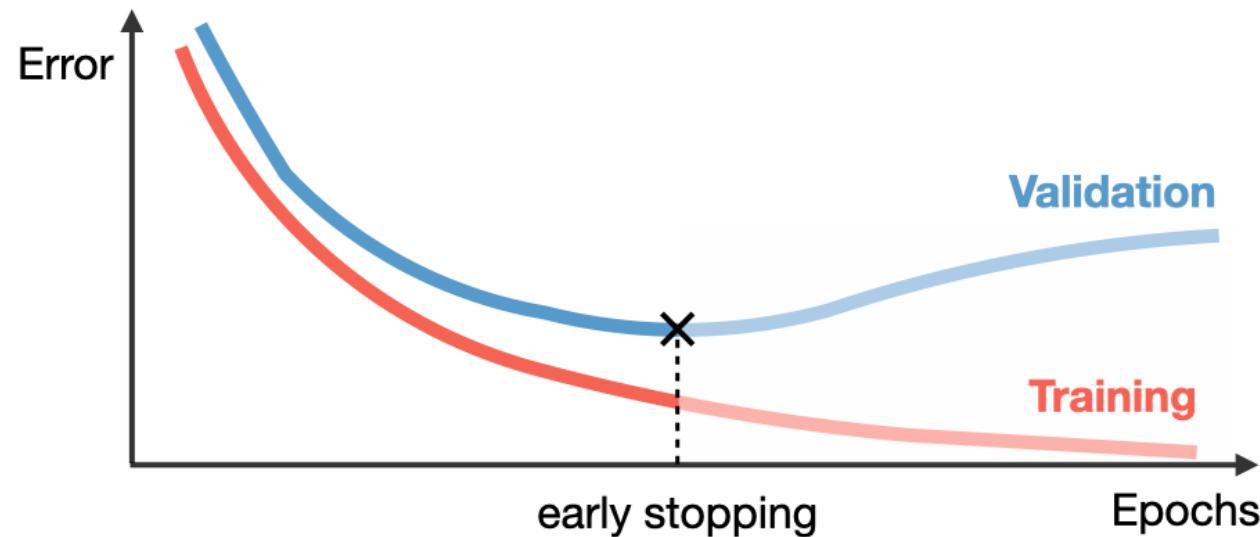
20



# Regularization – Early stopping

- Training for too many epochs can lead to overfitting
- Stop training at maximum generalization
- Make sure the model does not learn the noise in the training data
- Kind of cross-validation strategy
- After each epoch, compute error on unseen data (validation data)
- Stop training as soon as the validation loss reached a plateau or starts to increase

21

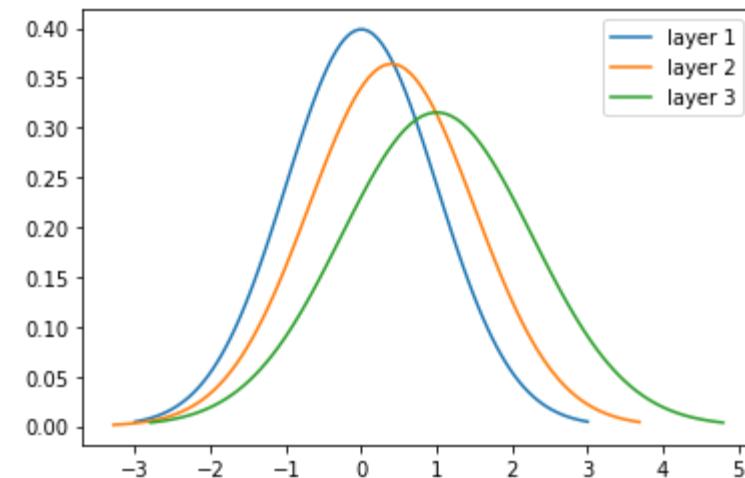


# Regularization – Early stopping

- + Computational efficiency
- + No architecture modification
- + Simplicity
- Dependency on validation set
- Hyperparameter to tune (patience)
- Potential underfitting

# Regularization – Batch normalization

- Internal covariate shift
  - Shift in input distribution over layers
  - Slow learning
  - Deeper network → Amplified effect

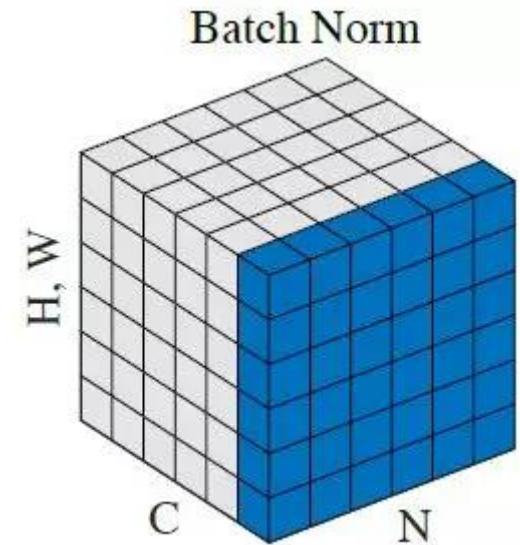


# Regularization – Batch normalization

- Reduce internal covariate shift
- Smooth the loss landscape
- Improve speed, performance and stability of NN
- Normalize inputs of each layer
  - Mean activation output zero and unit variance
- Run over batch axis

- After a fully connected or convolutional layer and before non-linearity layer (activation function)
- PyTorch: `torch.nn.BatchNorm1d(num_features, eps=1e-05)`



# Regularization – Batch normalization

**Input:** Values of  $x$  over a mini-batch:  $\mathcal{B} = \{x_1 \dots m\}$ ;

Parameters to be learned:  $\gamma, \beta$

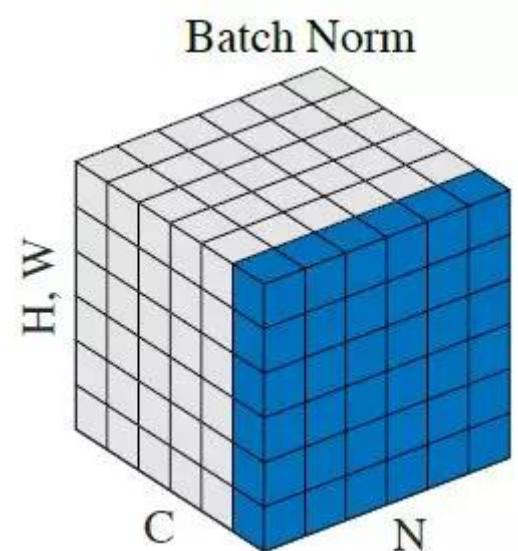
**Output:**  $\{y_i = \text{BN}_{\gamma, \beta}(x_i)\}$

$$\mu_{\mathcal{B}} \leftarrow \frac{1}{m} \sum_{i=1}^m x_i \quad // \text{mini-batch mean}$$

$$\sigma_{\mathcal{B}}^2 \leftarrow \frac{1}{m} \sum_{i=1}^m (x_i - \mu_{\mathcal{B}})^2 \quad // \text{mini-batch variance}$$

$$\hat{x}_i \leftarrow \frac{x_i - \mu_{\mathcal{B}}}{\sqrt{\sigma_{\mathcal{B}}^2 + \epsilon}} \quad // \text{normalize}$$

$$y_i \leftarrow \gamma \hat{x}_i + \beta \equiv \text{BN}_{\gamma, \beta}(x_i) \quad // \text{scale and shift}$$



Loffe and Szegedy (2015), Batch normalization: accelerating deep network training by reducing internal covariate shift

# Regularization – Batch normalization

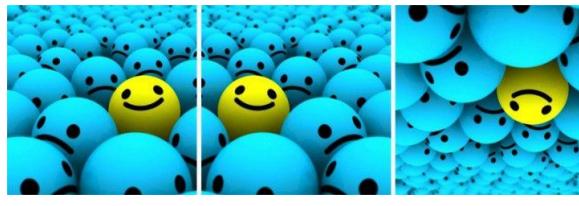
- + Accelerate training → Faster convergence
- + Higher learning rate
- + Improved gradient flow → Reduced vanishing/exploding gradients
- + Reduce sensitivity to initialization
  
- Dependency on mini-batch size
- Increase computational overhead
- Inconsistent behavior between training and inference
- Complexity in recurrent networks

# Regularization – Data augmentation

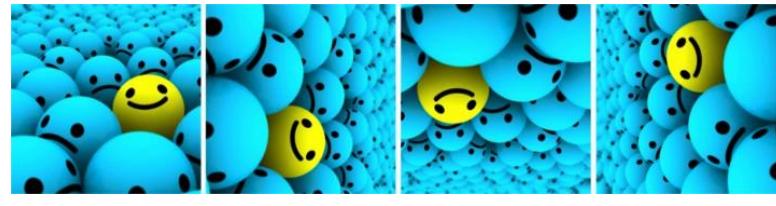
- NNs need a lot of data to be properly trained
- More information can be extracted from original data through augmentation
- Increase amount of data
  - Synthetic data → Generated artificially (GANs for example)
  - Augmented data → Derived from original images
- **Data augmentation**
  - Mostly used with images
  - Get more data from existing ones
  - Make minor changes

# Regularization – Data augmentation

- Data augmentation for **images**



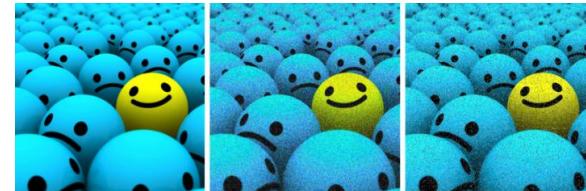
Flip



Rotation



Crop

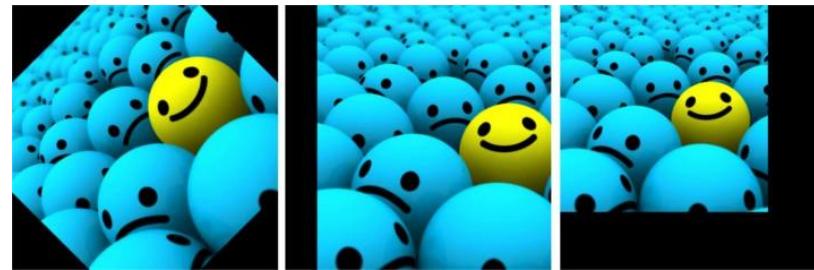


Add noise

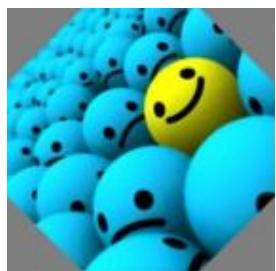
- Others
  - Color shift
  - Information loss

# Regularization – Data augmentation

- Data augmentation for **images**
  - Interpolation to preserve image size



29



Constant



Edge



Reflect



Symmetric



Wrap

# Regularization – Data augmentation

PyTorch function:

```
import torch
from torchvision import transform

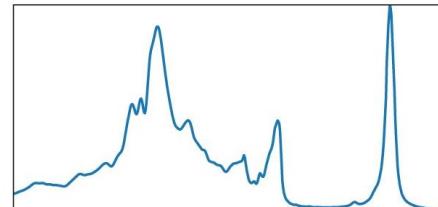
transform = transforms.Compose([
    transforms.RandomHorizontalFlip(p=0.5),
    transforms.RandomVerticalFlip(p=0.5),
    transforms.RandomRotation(degrees, interpolation),
    transforms.Resize(size, interpolation),
    transforms.RandomCrop(size, padding),
    transforms.ToTensor()
])
```

30

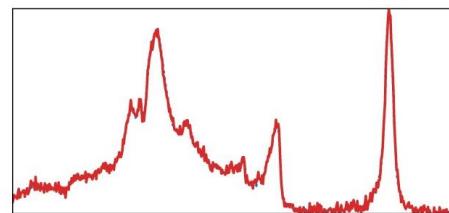


# Regularization – Data augmentation

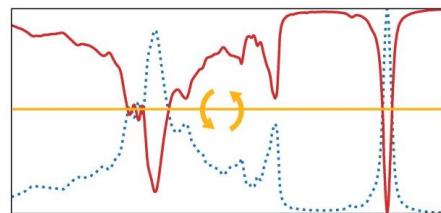
- Data augmentation for **time series**
  - Some geometric transformations might change the signal properties
  - Less investigations



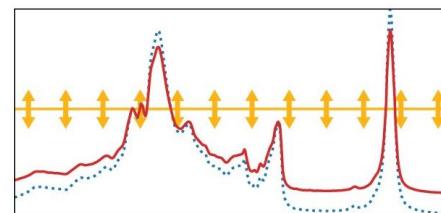
(a) Original



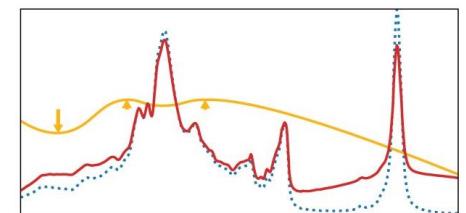
(b) Jittering



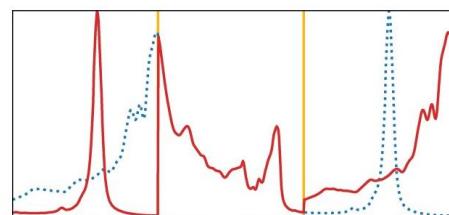
(c) Flipping



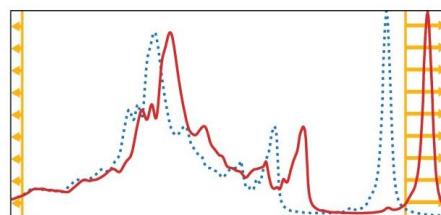
(d) Scaling



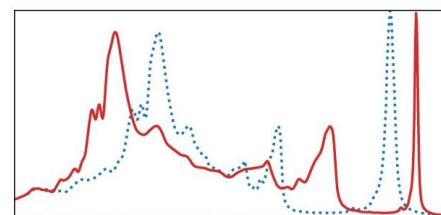
(e) Magnitude Warping



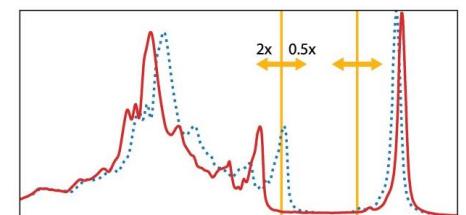
(f) Permutation



(g) Window Slicing



(h) Time Warping



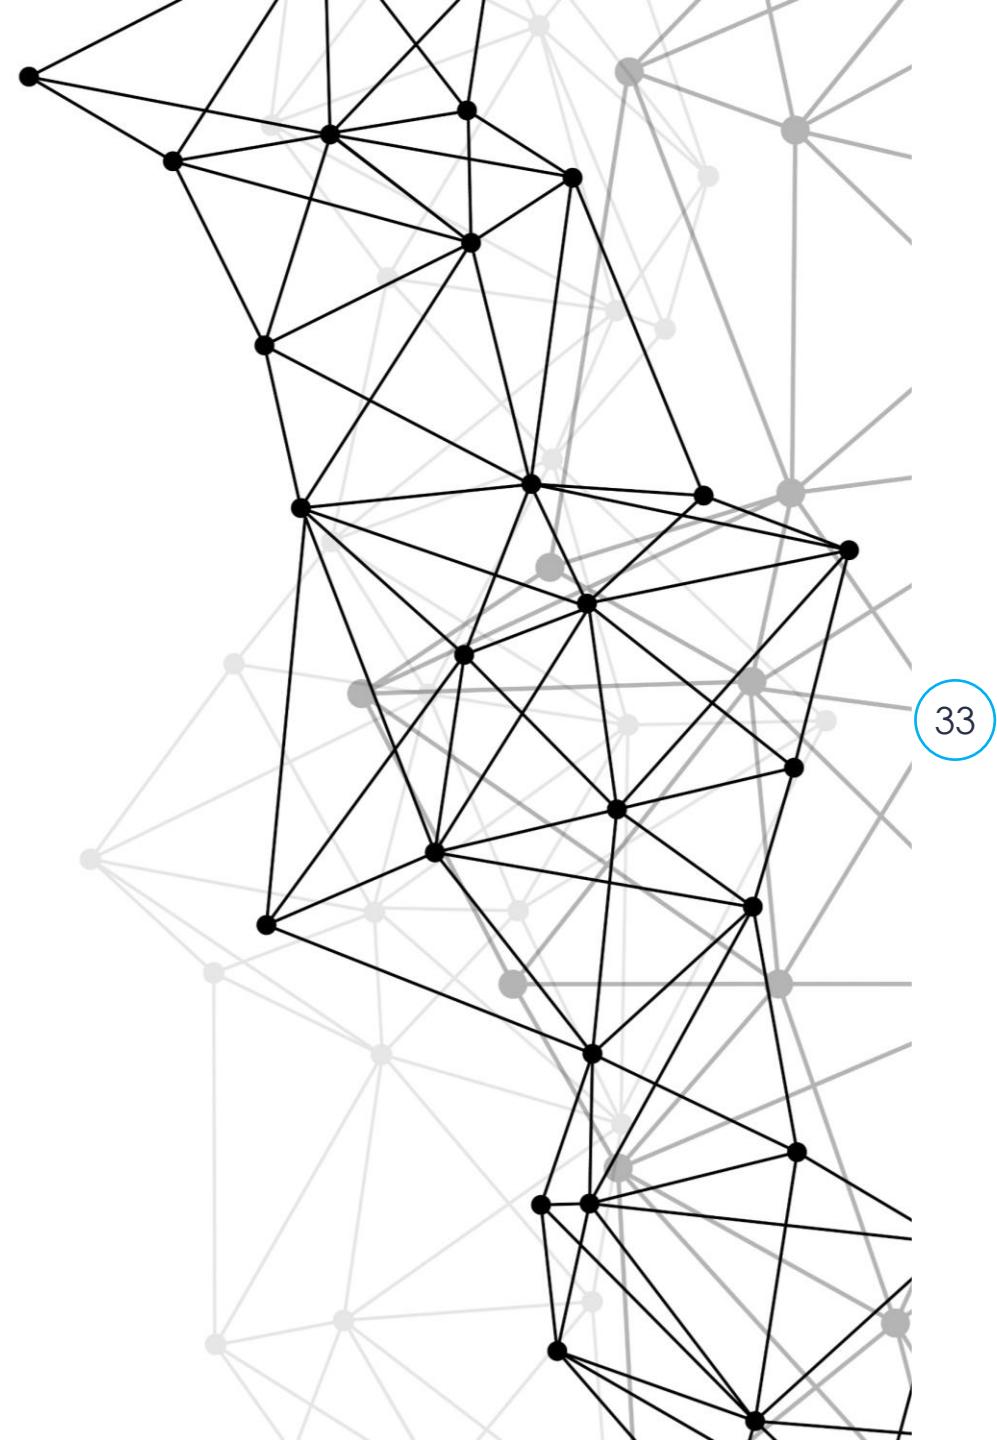
(i) Window Warping

# Regularization – Data augmentation

- + Introduction of invariance and robustness
- + Reduced need for costly data collection
- + No modification to the model architecture
  
- Potential introduction of noise and irrelevant variations
- Risk of label inconsistency
- Domain-specific design and expertise required
- Increased computational overhead
- Not universally applicable across all data types

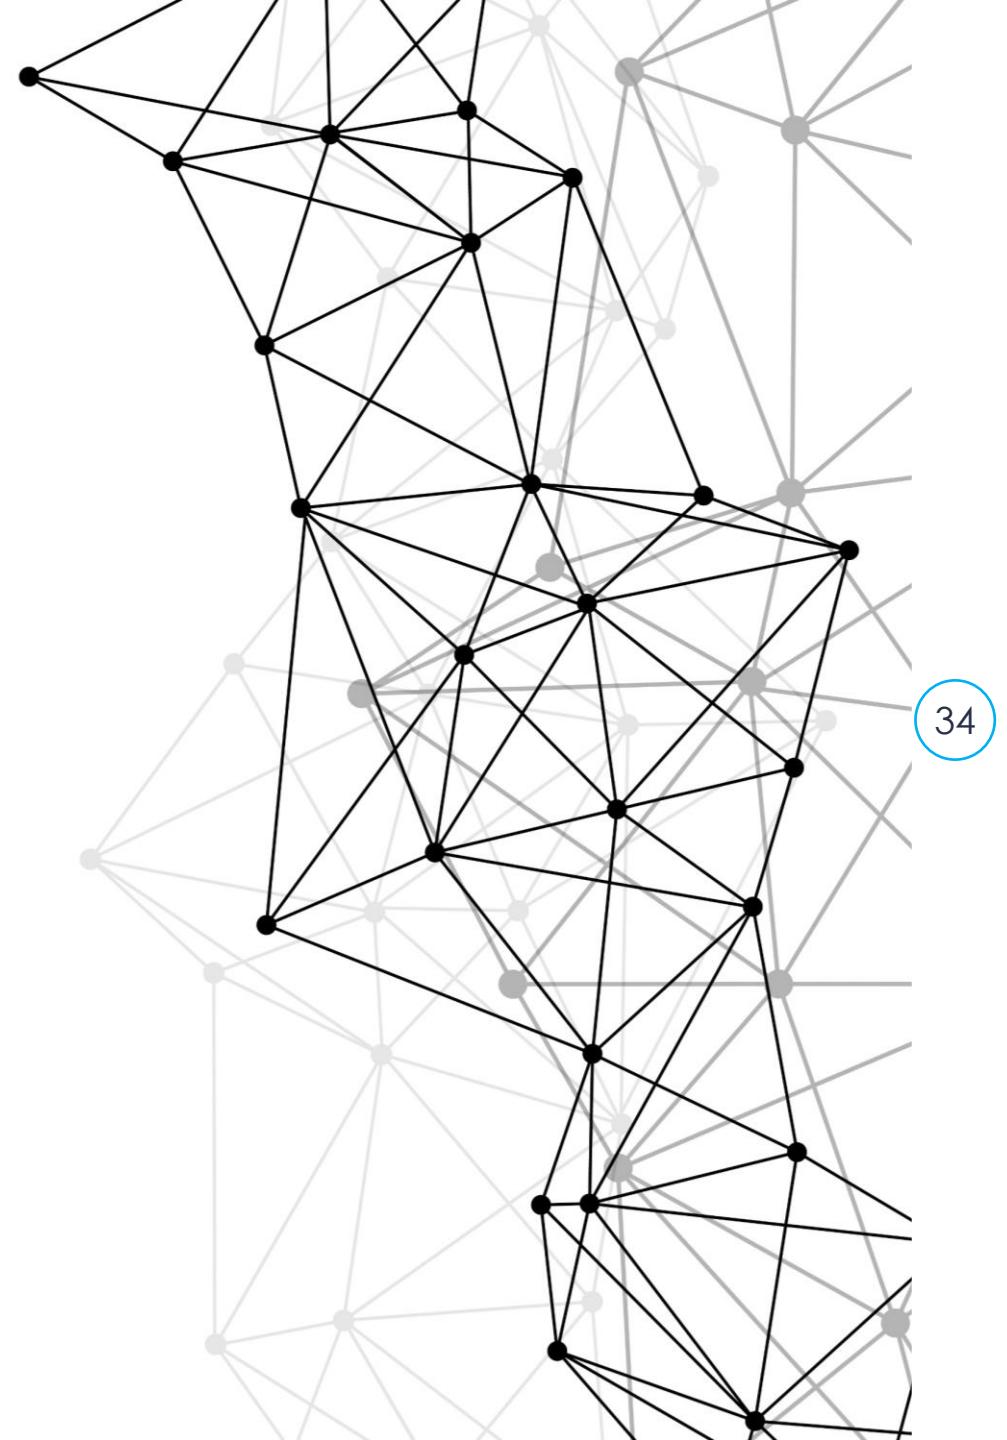
# Content

- Neural network recap
- Regularization
- Different architectures
- Labs



# Content

- Different architectures
  - Siamese network
  - Generative adversarial network
  - Explainable neural network

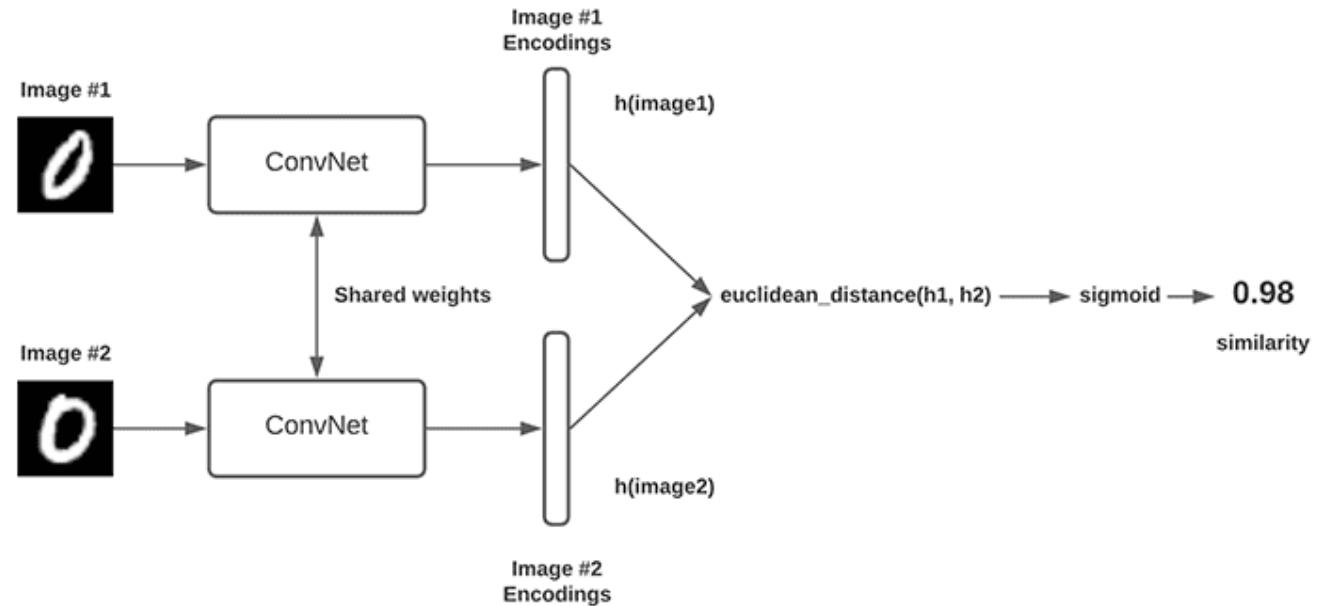


# Siamese network

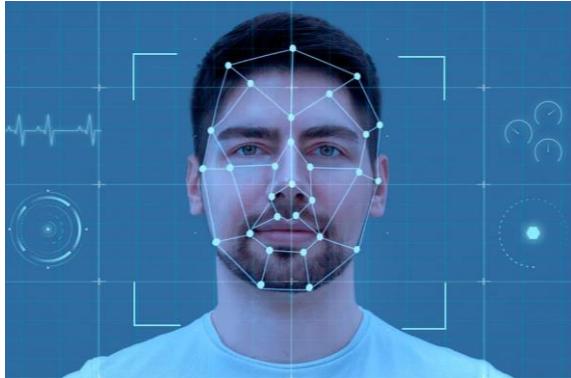
- Two or more identical subnetworks sharing the same parameters and weights
- Compare two inputs → How similar or dissimilar they are
- Learn representations to compare inputs similarity
  1. Inputs are passed into twin networks simultaneously
  2. Each subnetwork extracts features from its respective input
  3. Features are combined to measure distance/similarity between inputs

35

- + Learning from few examples
- + Weight sharing
- Training complexity
- Distance function sensitivity



# Siamese network - Applications



Face recognition



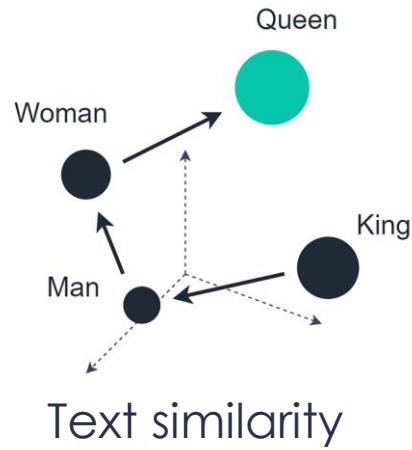
Signature verification

ए

ग घ र द न  
ङ ङ ङ ङ ङ<sup>ं</sup>  
ङ ङ ङ ङ ङ<sup>ं</sup>  
ङ ङ ङ ङ ङ<sup>ं</sup>

One-shot learning

36



Text similarity

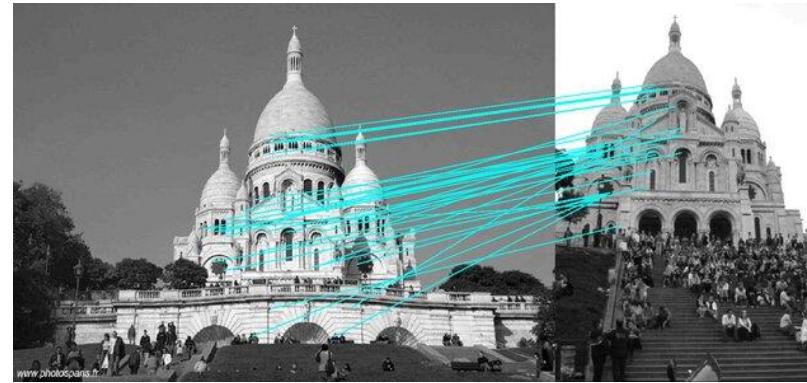


Image matching

# Siamese network

Schlesinger et al. 2020, Blood pressure estimation from PPG signals using convolutional neural networks and Siamese network  
IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)

- GOAL

- Estimate **blood pressure** (BP) from **photoplethysmography** (PPG) signal
- Feature learning approach  
→ Automatic feature extraction
- **Siamese** architecture

## BLOOD PRESSURE ESTIMATION FROM PPG SIGNALS USING CONVOLUTIONAL NEURAL NETWORKS AND SIAMESE NETWORK

Oded Schlesinger<sup>\*1</sup>, Nitai Vigderhouse<sup>\*1</sup>, Danny Eytan<sup>†</sup>, Yair Moshe<sup>\*</sup>

<sup>\*</sup>Signal and Image Processing Laboratory (SPL)  
Andrew and Erna Viterbi Faculty of Electrical Engineering  
Technion – Israel Institute of Technology  
<http://spl.technion.ac.il/>

<sup>†</sup>Ruth and Bruce Rappaport Faculty of Medicine  
Technion – Israel Institute of Technology

### ABSTRACT

Blood pressure (BP) is a vital sign of the human body and an important parameter for early detection of cardiovascular diseases. It is usually measured using cuff-based devices or monitored invasively in critically-ill patients. This paper presents two techniques that enable continuous and noninvasive cuff-less BP estimation using photoplethysmography (PPG) signals with Convolutional Neural Networks (CNNs). The first technique is calibration-free. The second technique achieves a more accurate measurement by estimating BP changes with respect to a patient's PPG and ground truth BP values at calibration time. For this purpose, it uses Siamese network architecture. When trained and tested on the MIMIC-II dataset, it achieves mean absolute difference of 4.41 mmHg and a BP of 5.95 mmHg and 3.41 mmHg respectively. These results almost comply with the AAMI recommendation and are as accurate as the values estimated by many home BP measuring devices.

**Index Terms**— Blood pressure, convolutional neural network (CNN), noninvasive, photoplethysmography (PPG), Siamese network

### 1. INTRODUCTION

Blood pressure (BP) is the result of force exerted in the arteries by blood as it circulates. It is usually expressed in terms of systolic pressure (when the heart beats and BP is at its highest) and diastolic pressure (between heart beats, when BP is at its lowest) and measured in millimeters of mercury (mmHg). Normal resting BP in an adult is approximately 120 mmHg systolic and 80 mmHg diastolic. BP is an important parameter of the human body whose measurement allows for the early detection of medical issues, especially cardiovascular diseases, which are a leading cause of mortality and morbidity worldwide. High BP, hypertension, is a major risk for dangerous health conditions such as stroke or heart attack. Low BP, hypotension, can cause dizziness and fainting or may indicate serious heart, endocrine or neurological disorders. Therefore, it is highly important to measure BP accurately. Continuous monitoring of BP, along with monitoring of other vital signs, allows an accurate evaluation of the patient's physiological state, prompt detection of deteriorations and their prediction. The current widespread BP monitoring methods are divided into invasive and noninvasive methods. Invasive arterial line is a clinical standard for continuous high accuracy BP measurement. However, it has adverse effects associated with invasive measurements, such as potential infection, all of them are associated with an increased morbidity. Noninvasive BP measurement methods typically use an oscillometry inflation arm or wrist cuff. These methods are not feasible for long-term monitoring of BP, mainly due to discomfort caused by repeated inflation and deflation and mobility limitations caused by the measuring device.

The Association for the Advancement of Medical Instrumentation (AAMI) recommends that the mean absolute difference (MAD) of noninvasive BP measurement technologies should not be greater than 5 mmHg and the standard deviation should not be greater than 8 mmHg compared to a reference method [1]. Home BP measuring devices may be inaccurate in 3% to 15% of patients, and a difference of 5 mmHg or higher is common [2, 3]. Therefore, BP estimation using other noninvasive technologies are expected to have at least similar accuracy, as they are at least as reliable as home BP measuring devices.

Photoplethysmography (PPG) is an optically obtained signal that can be used to detect blood volume changes in the microvascular bed of tissue. It is obtained by illuminating the skin and measuring changes in light absorption. In a clinical setting, this signal is often obtained by a pulse oximeter while in an ambulatory setting, this signal can be obtained by a smartwatch or other mobile devices. PPG signals contain information on cardiovascular parameters such as heart rate, blood oxygen saturation and BP. Since PPG is noninvasive, simple and low-cost, it has wide potential for clinical

# Siamese network – Blood pressure (BP)

- **#** risk factor of cardiovascular diseases
- 1.28 billion people affected
- 46% no noticeable symptoms

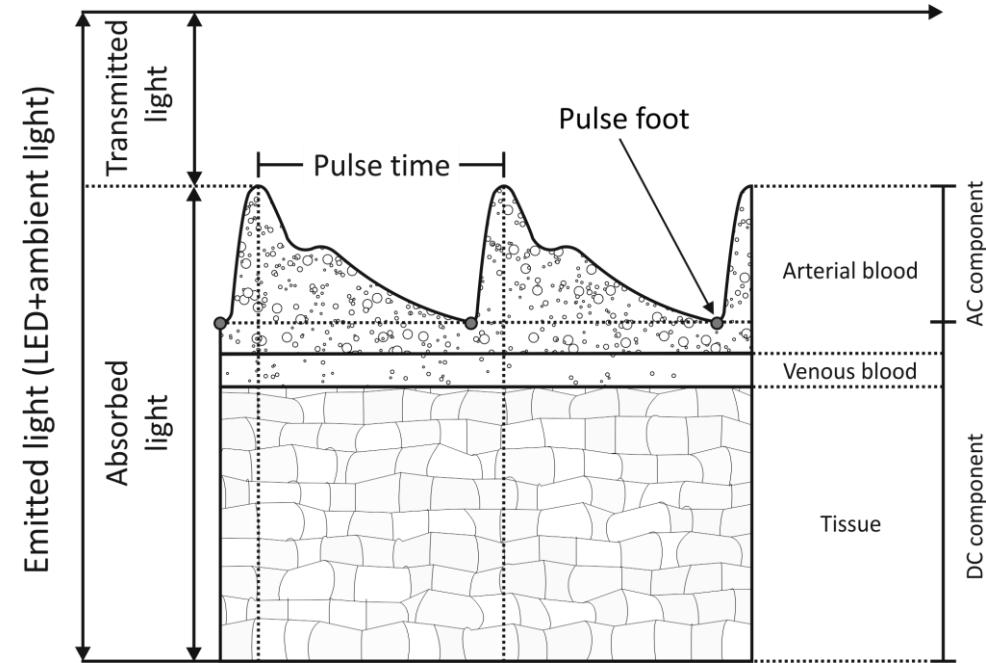


Cuff-based BP measurement

- Occlusive
- Intermittent
- Uncomfortable

# Siamese network – Photoplethysmography (PPG)

- Measure of changes in light absorption
- Related to blood volume variations
- **Light source** to illuminate tissue
- **Photodetector** to measure changes in light intensity
- Information on **cardiovascular system**
  - Blood oxygen saturation
  - Heart rate
  - **Blood pressure**



39



# Siamese network – Data

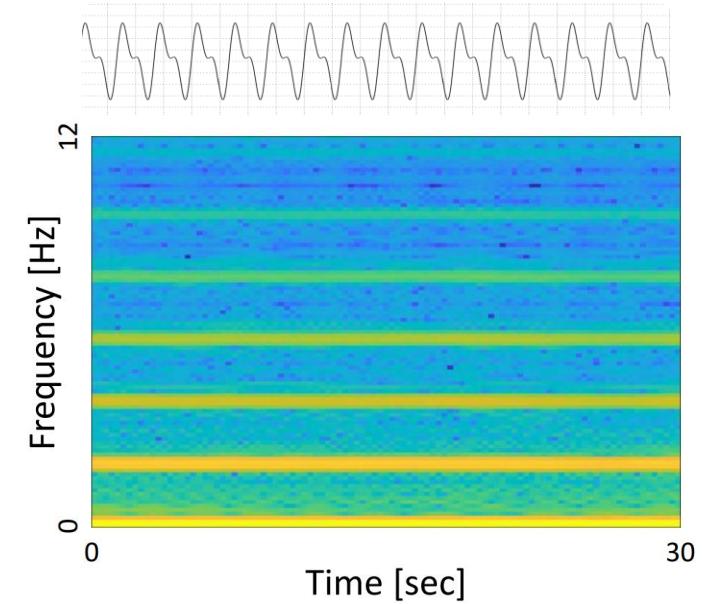
- MIMIC-II database [1]
  - 1459 patients in intensive care unit (ICU)



PPG



BP arterial line



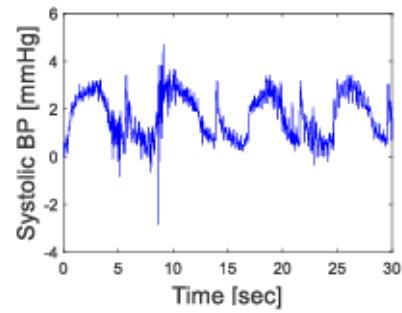
**spectrogram**  
of 30-second PPG window

[1] M. Saeed et al., "Multiparameter Intelligent Monitoring in Intensive Care II (MIMIC-II): a public-access intensive care unit database," Critical Care Medicine, vol. 39, no. 5, p. 952, 2011.  
<https://archive.physionet.org/mimic2/>

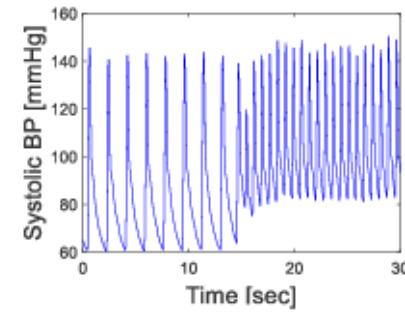
# Siamese network

- Preprocessing
  - Remove unreliable windows
  - Remove unreliable patients
  - Remove outliers

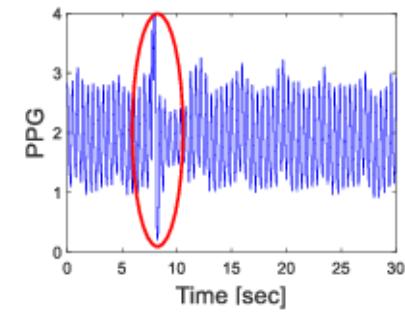
$\sim 10^5$  30s windows  
from 304 patients



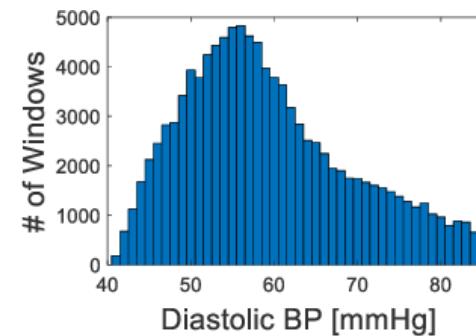
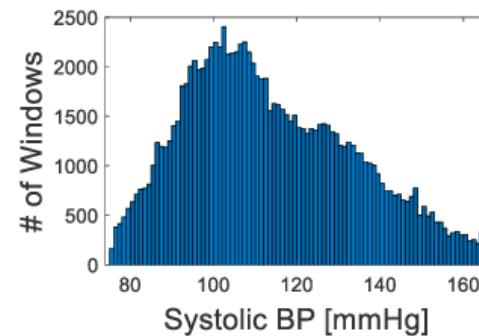
No physiological BP



BP fluctuation  
within 30s window

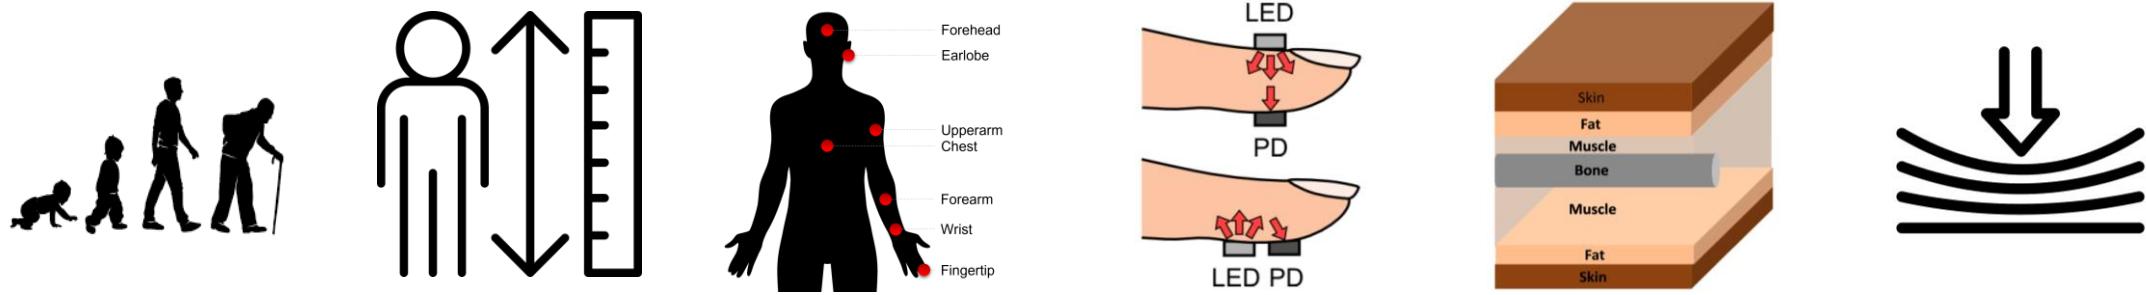


Noisy PPG and  
ABP signals



# Siamese network

- Calibration
  - PPG waveform variability due to individual-specific characteristics or external factors

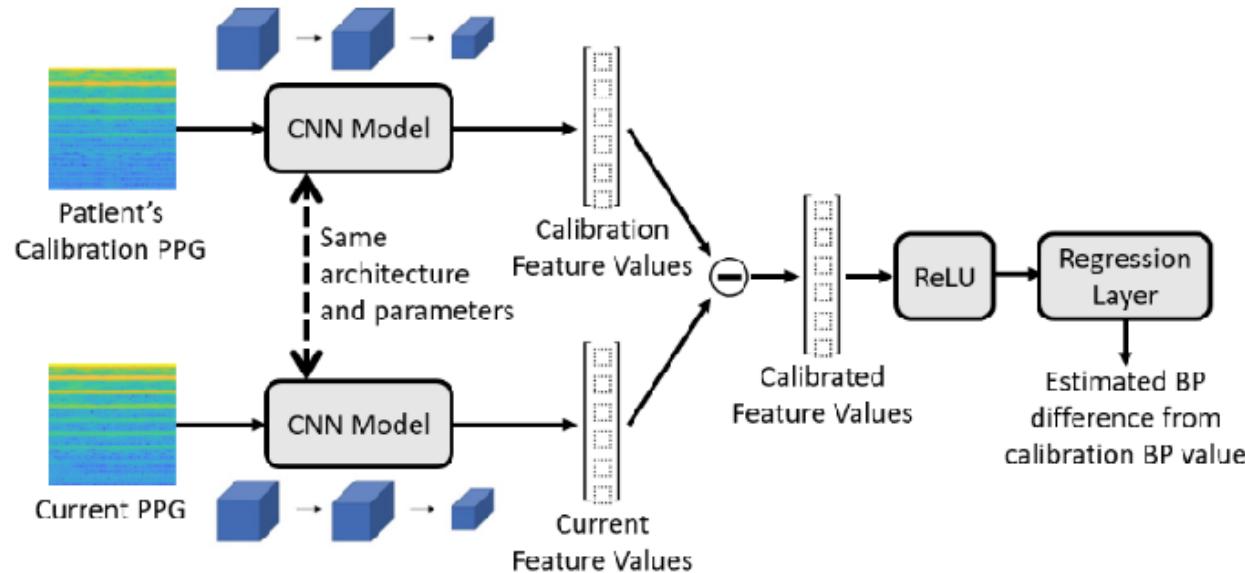


- Measure taken at the doctor office
- Using subject mean value
- Partly or completely retraining the model

# Siamese network

- Calibration
  - First available 30s window
  - Using Siamese network
  - 2 identical subnetworks with same architecture and parameters
  - Working in parallel on two different inputs (calibration and current PPG) to compute feature vectors
  - Compare inputs by computing the difference between feature vectors.

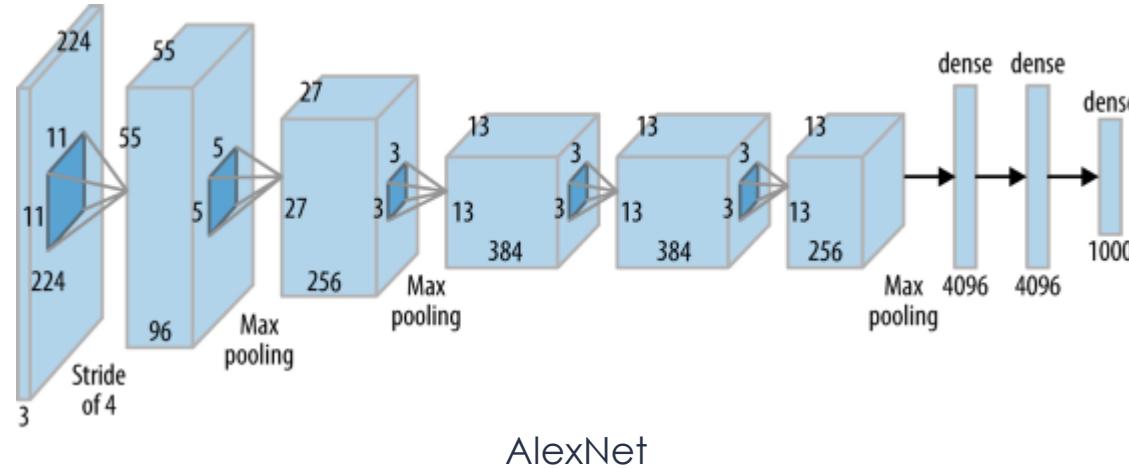
43



# Siamese network

- CNN architecture
  - Inspired by AlexNet
  - Regularization with **batch normalization**
  - L1 loss function
  - Regression problem → Linear layer at the end

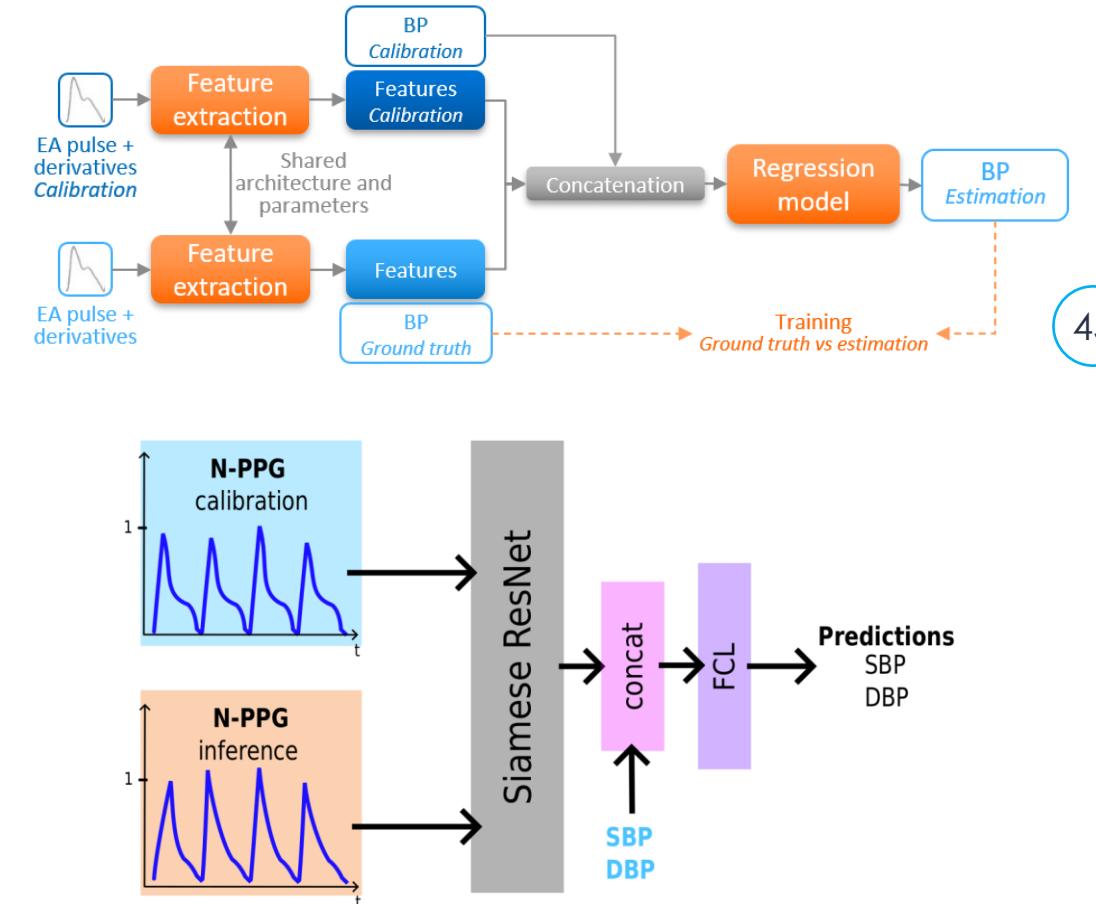
44



# Siamese network

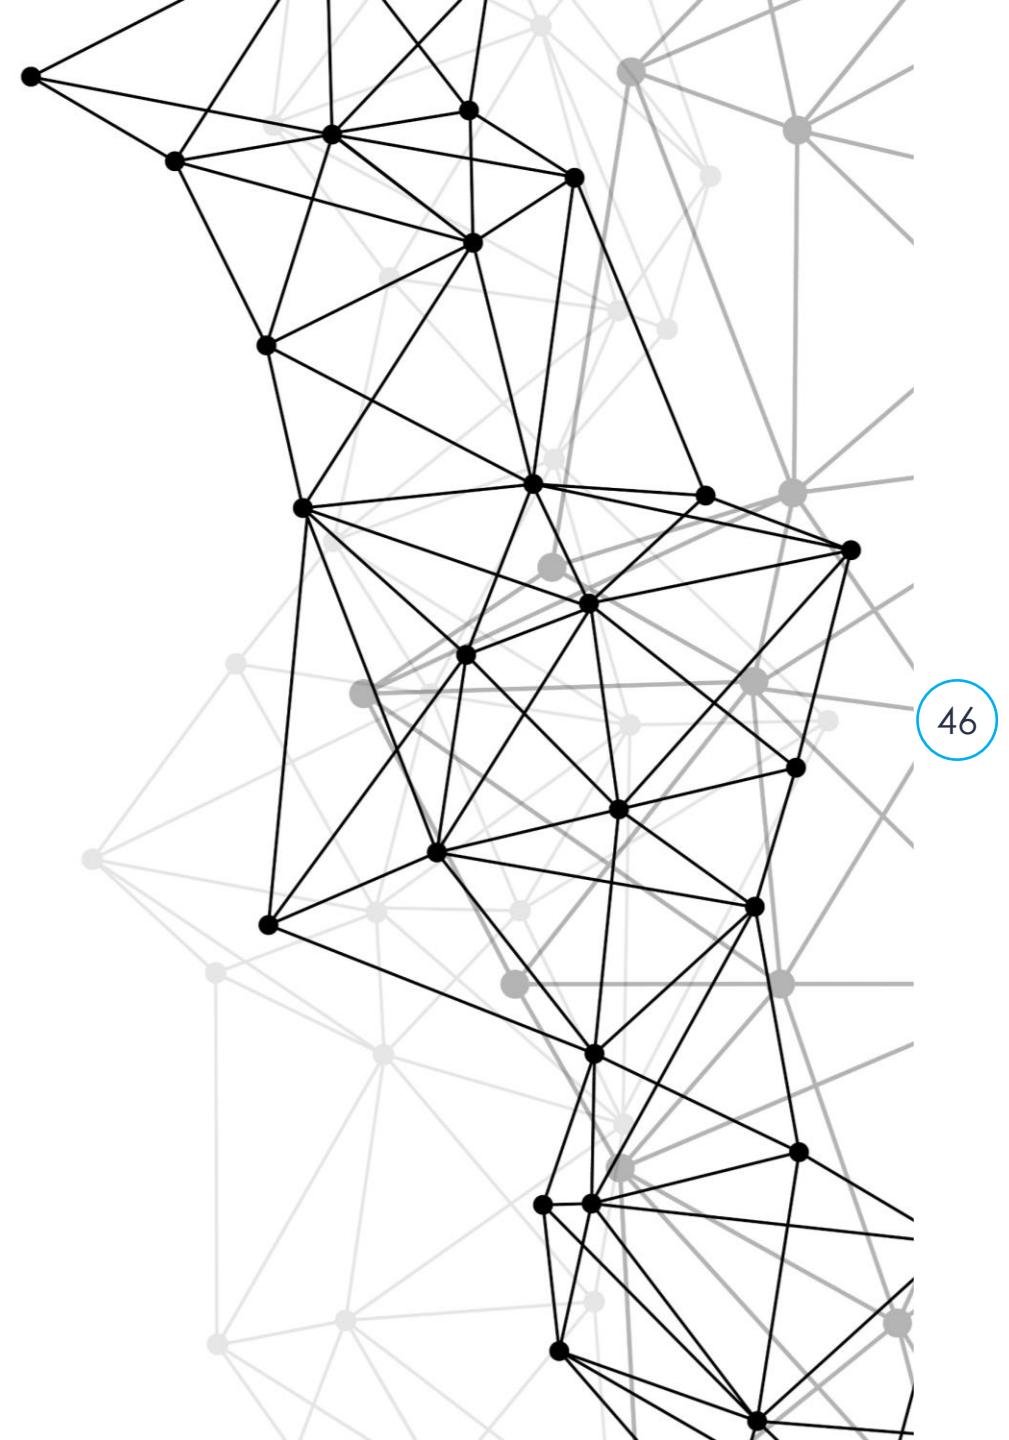
- **Similar approach**

- C. Aguet et al. - Feature learning for blood pressure estimation from photoplethysmography (2021)
  - Ensemble average pulse over 20s window of PPG and its derivatives
  - Siamese CNN
- F.M. Dias et al. - Exploring the limitations of blood pressure estimation using the photoplethysmography signal (2024)
  - 24s PPG segments
  - Siamese ResNet



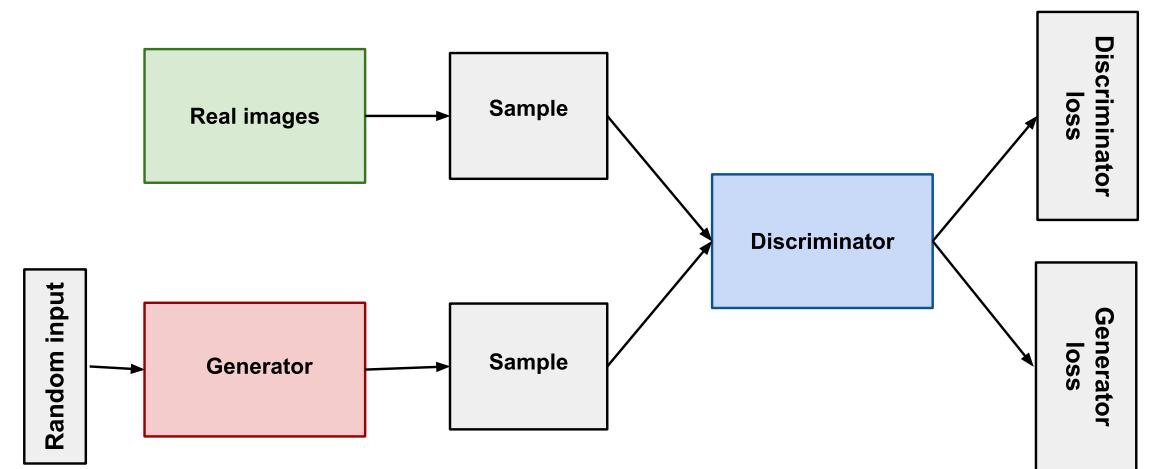
# Content

- Different architectures
  - Siamese network
  - Generative adversarial network
  - Explainable neural network

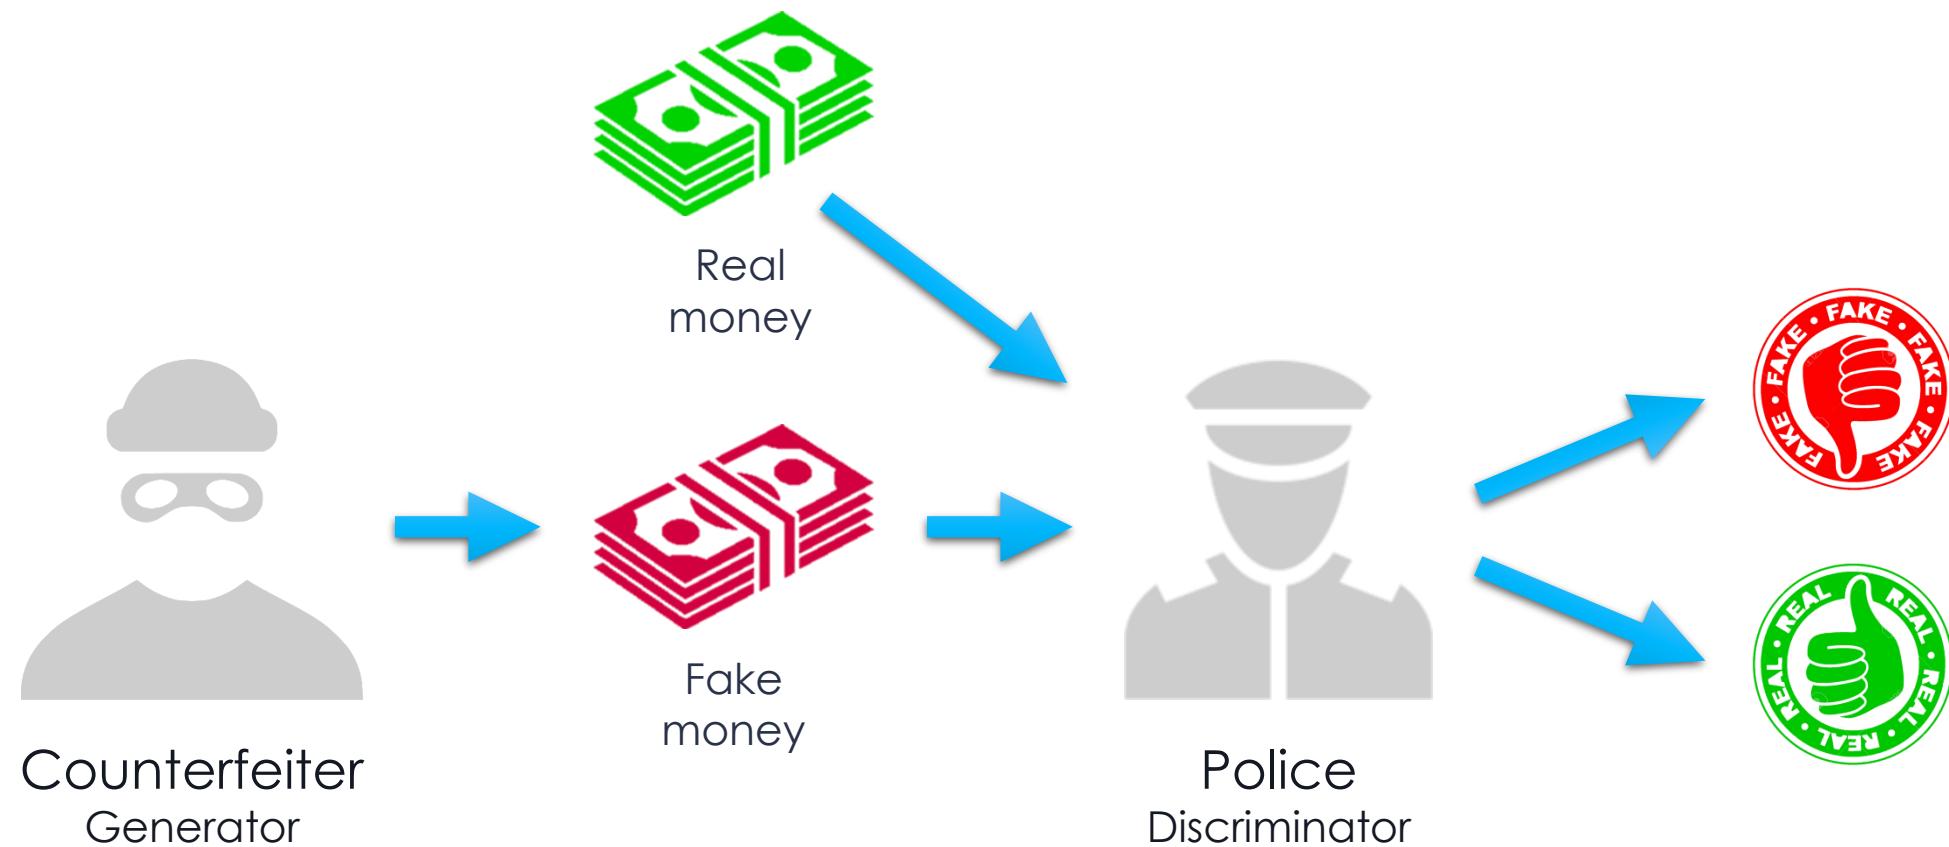


# Generative adversarial network

- Proposed by Ian Goodfellow in 2014
- Generating new samples from scratch
- Learning by comparison
- **Generator**
  - Generate realistic-looking synthetic data
  - Try to fool discriminator
  - From noise to sample
  - Low to high dimensional space
- **Discriminator**
  - Distinguish generated and real samples
  - Binary classifier
  - From sample to decision
  - High to low dimensional space



# Generative adversarial network



# Generative adversarial network

- GAN training
  - Approximating data distribution and sampling from such approximation
  - Adversarial learning → Two network fight against each other
  - Simultaneous training → Alternating Generator and Discriminator updates
  - Discriminator
    - Maximize probability of assigning correct class → Maximize  $\log D(X)$
  - Generator
    - Maximize Discriminator uncertainty → Minimize  $\log(1 - D(G(X)))$

49

$$\begin{aligned} & \min_G \max_D \mathbb{E}_{X \sim \mu} [\log D(X)] + \mathbb{E}_{X \sim \mu_Z} [\log (1 - D(G(X)))] \\ & \min_G \max_D \mathbb{E}_{X \sim \mu} [\log D(X)] + \mathbb{E}_{X \sim \mu_G} [\log (1 - D(X))] \end{aligned}$$

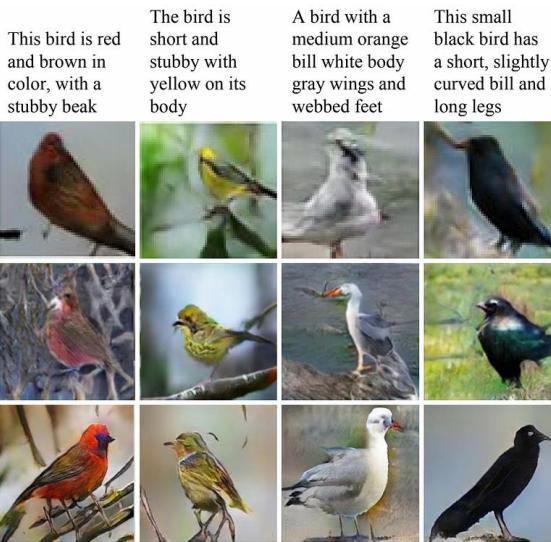
# Generative adversarial network

- Applications

Image-to-image  
translation

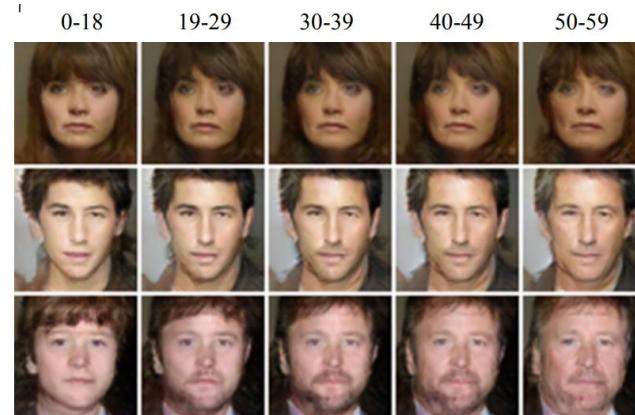


New poses  
generation



Text-to-image  
synthesis

50



Face aging

# Generative adversarial network

Sarkar and Etemad 2020, CardioGAN: Attentive generative adversarial network with dual discriminators  
AAAI conference on artificial intelligence

CardioGAN: Attentive Generative Adversarial Network with Dual Discriminators for Synthesis of ECG from PPG

Pritam Sarkar, Ali Etemad  
Dept. of ECE & Ingenuity Labs Research Institute  
Queen's University, Kingston, Canada  
{pritam.sarkar, ali.etemad}@queensu.ca

51

- GOAL
  - Use of ECG device for continuous heart rate (HR) monitoring
  - PPG to ECG translation

## Abstract

Electrocardiogram (ECG) is the electrical measurement of cardiac activity, whereas Photoplethysmogram (PPG) is the optical measurement of volumetric changes in blood circulation. While both signals are used for heart rate monitoring, from a medical perspective, ECG is more accurate as it carries additional cardiac information. Despite many attempts toward incorporating ECG sensing in smartwatches or similar wearable devices for continuous and reliable cardiac monitoring, it remains a challenging and multi-faceted sensing solution available. In order to tackle this problem, we propose CardioGAN, an adversarial model which takes PPG as input and generates ECG as output. The proposed network utilizes an attention mechanism to focus on local salient features, as well as dual discriminators to prevent the integrity of generated data in both time and frequency domains. Our experiments show that the ECG generated by CardioGAN provides more reliable heart rate measurements compared to the original input PPG, reducing the error from 9.74 beats per minute (measured from the PPG) to 2.89 (measured from the generated ECG).

## 1 Introduction

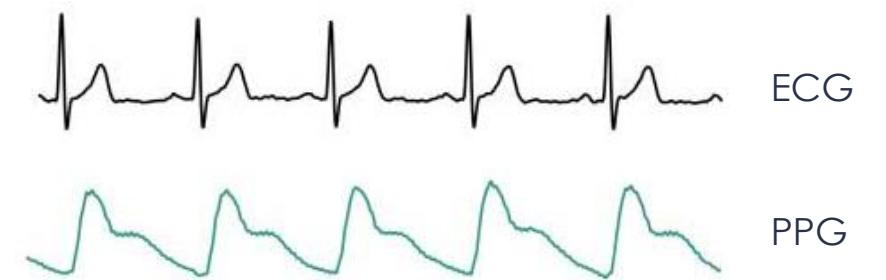
According to the World Health Organization (WHO) in 2017, Cardiovascular Diseases (CVDs) are reported as the leading causes of death worldwide (WHO 2017). The report indicates that CVDs cause 31% of global deaths, out of which 80% are due to cardiovascular diseases in the low or medium-income countries. One of the primary reasons behind this is the lack of primary healthcare support and the inaccessible on-demand health monitoring infrastructure. Electrocardiogram (ECG) is considered as one of the most important attributes for continuous health monitoring required for identifying potential serious risk of future cardiovascular events to health. Yet, a significant work is being conducted with the goal of developing wearable devices capable of continuous ECG monitoring and feasible for daily life use, largely to no avail. Currently, very few wearable devices provide wrist-based ECG monitoring, and those who do, require the user to stand still and touch the watch with both hands in order to *close the circuit* in order to record an ECG segment of limited duration (usually 30 seconds), making these solutions non-continuous and sporadic.

Photoplethysmogram (PPG), an optical method for measuring blood volume changes at the surface of the skin, is considered as a close alternative to ECG. It provides valuable cardiovascular information (Ghil et al. 2010; Schäfer and Yagci 2013). For instance, studies have shown that a number of features extracted from PPG (e.g., pulse rate variability) are highly correlated with corresponding metrics extracted from ECG (e.g., heart rate variability) (Gil et al. 2019), further suggesting that PPG can be used for monitoring heart rate variability. Yet, through recent advancements in smartwatches, smartphones, and other similar wearable and mobile devices, PPG has become the industry standard as a simple, wearable-friendly, and low-cost solution for continuous heart rate (HR) monitoring for everyday use. Nonetheless, PPG suffers from inaccurate HR estimation and several other limitations in comparison to conventional ECG monitoring (Korpi et al. 2017). These include, but not limited to, diverse skin types, motion artifacts, and signal crossovers among others. Moreover, the ECG waveform carries important information about cardiac activity. For instance, the P-wave indicates the sinus rhythm, whereas a long PR interval is generally indicative of a first-degree heart blockage (Ashley and Niebauer 2003). As a result, ECG is consistently being used by cardiologists for assessing the condition and performance of the heart.

Based on the above observations, there is a clear discrepancy between the need for continuous wearable ECG monitoring and the available solutions in the market. To address this, we propose CardioGAN, a generative adversarial network (GAN) (Goodfellow et al. 2014), which takes PPG as inputs and generates ECG. Our model is based on the CycleGAN architecture (Zhu et al. 2017) which is trainable and able to be trained in an unsupervised manner. Unlike CycleGAN, CardioGAN is designed with attention-based generators and equipped with multiple discriminators. We utilize attention mechanisms in the generators to better learn to focus on specific local regions such as the QRS complexes of ECG. To generate high fidelity ECG signals in terms of both time and frequency information, we utilize a dual discriminator

# Generative adversarial network

- **Electrocardiogram (ECG)**
  - Electrical measurement of cardiac activity
  - Continuous health monitoring
  - Complicated integration into wearables  
→ Non-continuous and sporadic measure
- **Photoplethysmogram (PPG)**
  - Optical measurement of volumetric changes in blood circulation
  - Simple, wearable-friendly, and low-cost



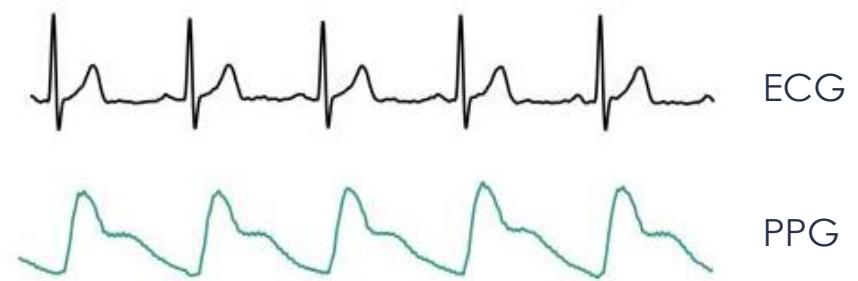
52



# Generative adversarial network

- Popular ECG-PPG **datasets**
  - BIDMC database – 53 ICU subjects, 8min
  - CAPNO database – 42 subjects, 8min
  - DALIA database – 15 subjects, ~2h during daily life activities
  - WESAD database – 15 subjects, >1h performing specific tasks
- ECG and PPG **preprocessing**
  1. Resampling to 128 Hz
  2. Filtering – band-pass FIR filter
  3. Z-score normalization on recording (zero mean and unit variance)
  4. Segmentation into 4s windows (512 data points)
  5. Min-max normalization on window  $[-1, 1]$

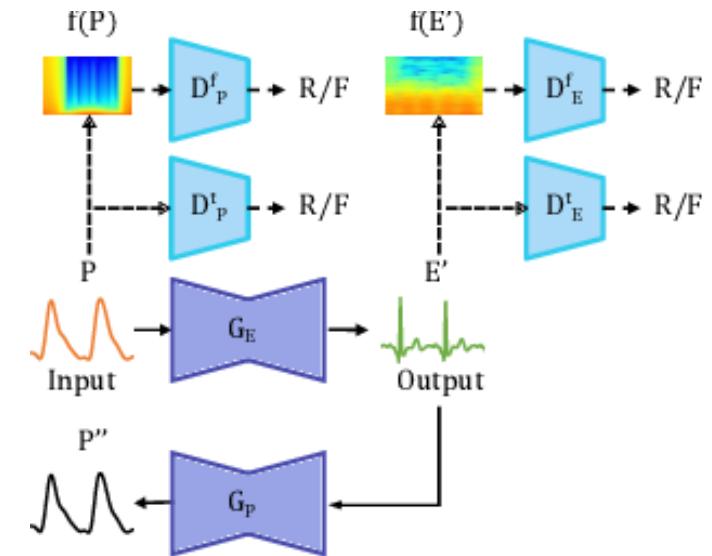
53



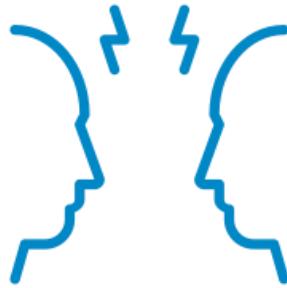
# Generative adversarial network

Learn the mapping between PPG (P) and ECG (E)

- Generator (G)
  - Forward mapping:  $G_E: P \rightarrow E$
  - Fake ECG:  $E' = G_E(P)$
  - Reconstructed PPG:  $P'' = G_P(G_E(P))$
- Discriminator (D)
  - Time-domain:  $D_E^t: E \text{ vs } E'$
  - Frequency-domain:  $D_E^f: f(E) \text{ vs } f(E')$
  - With  $f(x) = STFT(x)$
  - Short-Time Fourier Transform



# Generative adversarial network



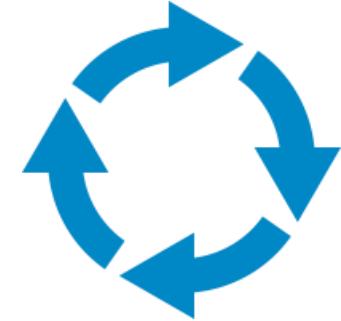
- Adversarial loss
  - Forward
    - Time-domain  $\rightarrow L_{adv}(G_E, D_E^t) = E_{e \sim E} \left[ \log(D_E^t(e)) \right] + E_{p \sim P} \left[ \log(1 - D_E^t(G_E(p))) \right]$
    - Frequency-domain  $\rightarrow L_{adv}(G_E, D_E^f) = E_{e \sim E} \left[ \log(D_E^f(f(e))) \right] + E_{p \sim P} \left[ \log(1 - D_E^f(f(G_E(p)))) \right]$
  - Backward
    - Time-domain  $\rightarrow L_{adv}(G_P, D_P^t)$
    - Frequency-domain  $\rightarrow L_{adv}(G_E, D_E^f)$

# Generative adversarial network

- Cycle loss

$$L_{cycle}(G_E, G_P) = E_{e \sim E} \left[ \|G_E(G_P(e)) - e\|_1 \right] + E_{p \sim P} \left[ \|G_P(G_E(p)) - p\|_1 \right]$$

- $p \rightarrow G_E(p) \rightarrow G_P(G_E(p)) \approx p$
- $e \rightarrow G_P(e) \rightarrow G_E(G_P(e)) \approx e$

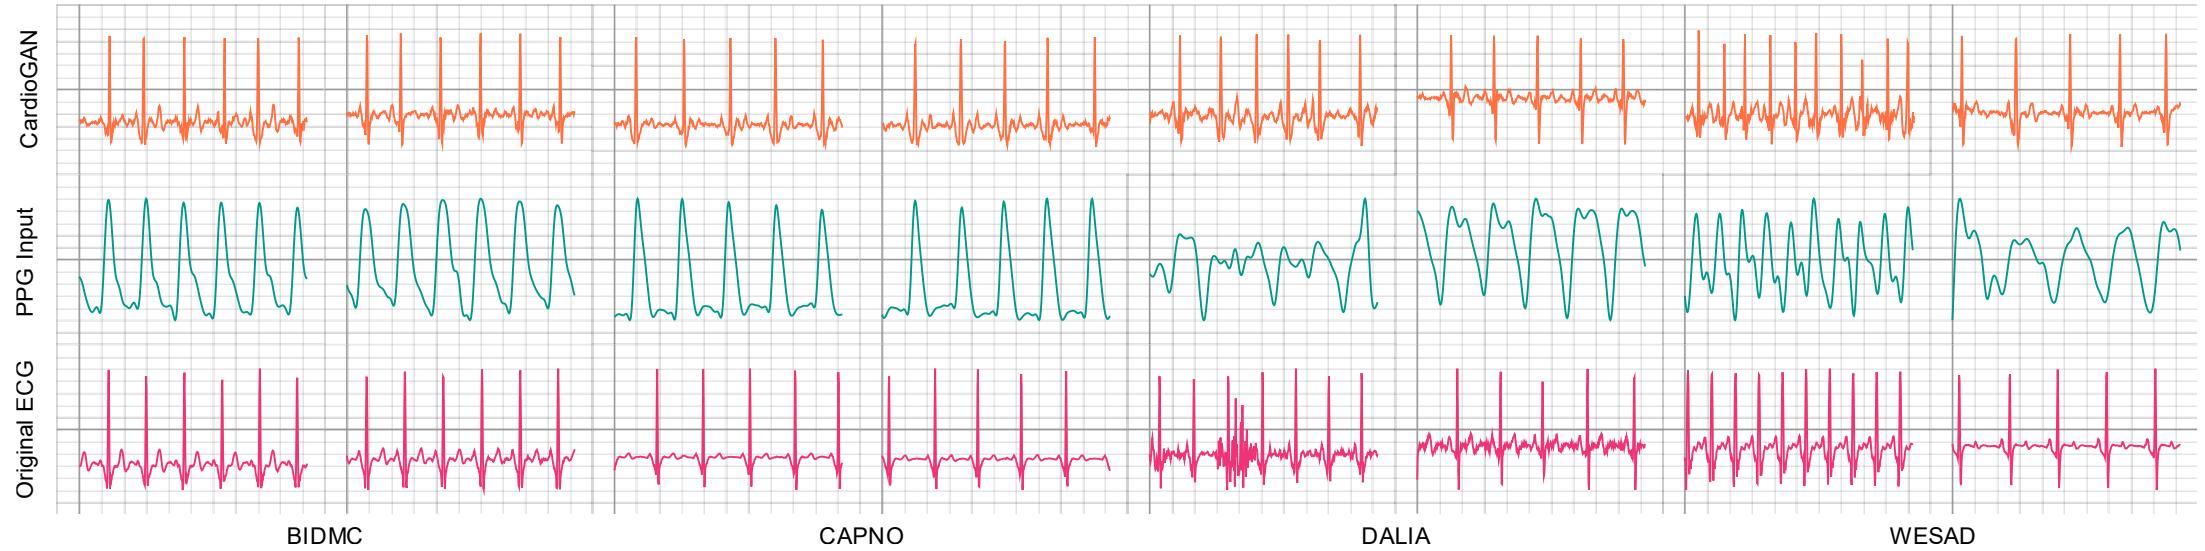


- Final loss

$$L_{CardioGAN} = \alpha L_{adv}(G_E, D_E^t) + \alpha L_{adv}(G_P, D_P^t) + \beta L_{adv}(G_E, D_E^f) + \beta L_{adv}(G_P, D_P^f) + \lambda L_{cycle}(G_E, G_P)$$

- $\alpha$  and  $\beta$  are adversarial loss coefficients
- $\lambda$  cycle consistency loss coefficient

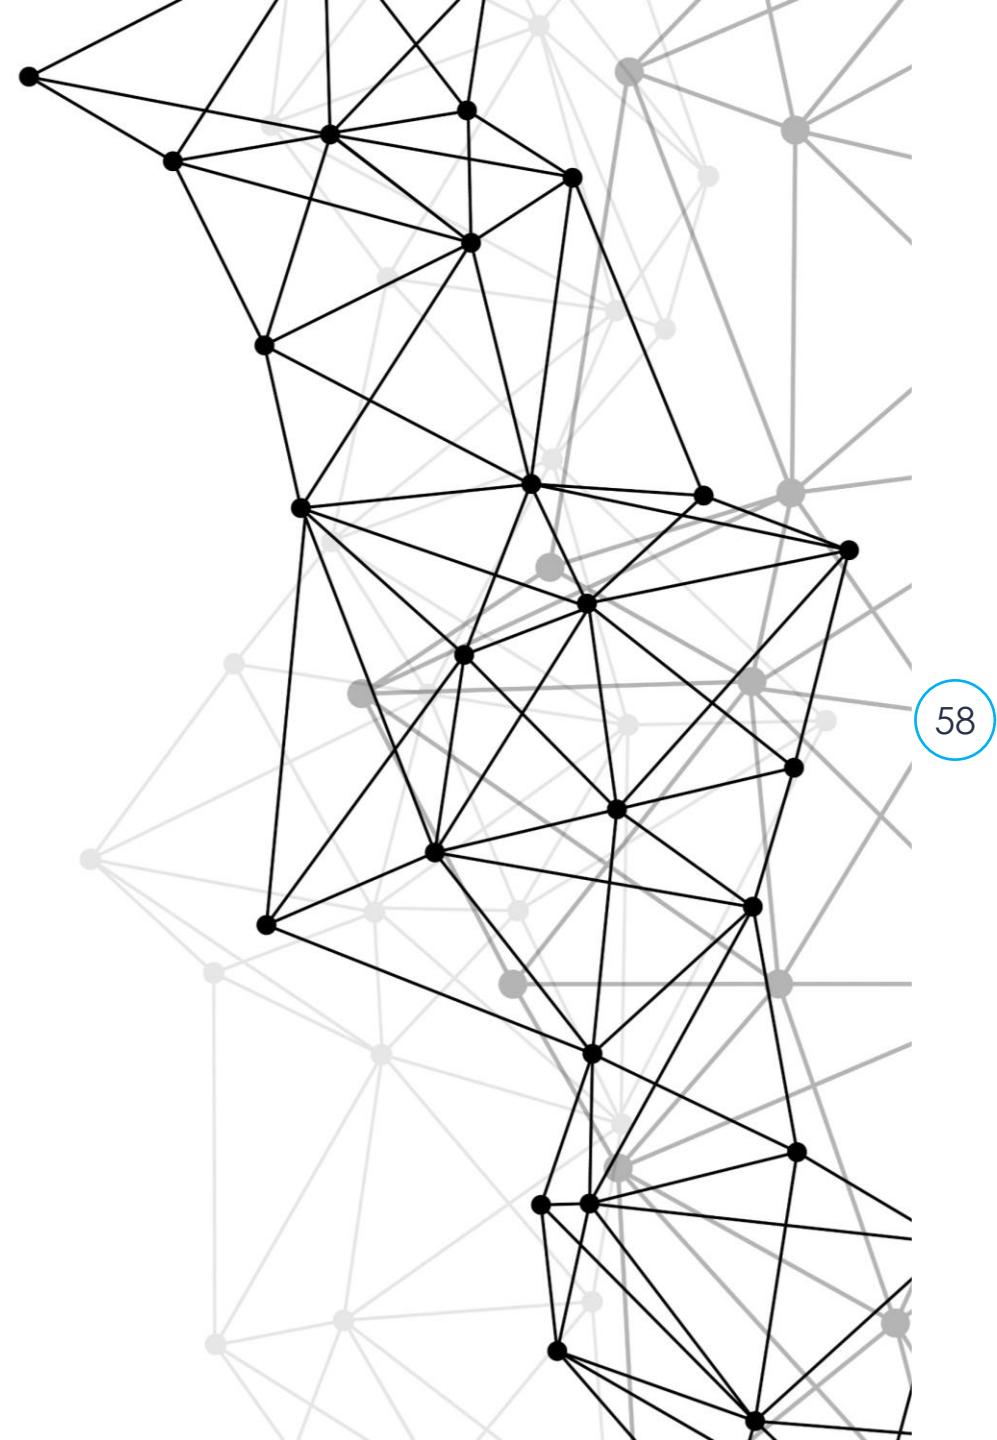
# Generative adversarial network



ECG generated by CardioGAN provides more reliable heart rate measurements compared to the original input PPG

# Content

- Different architectures
  - Siamese network
  - Generative adversarial network
  - Explainable neural network



# Explainable neural network



# Explainable neural network

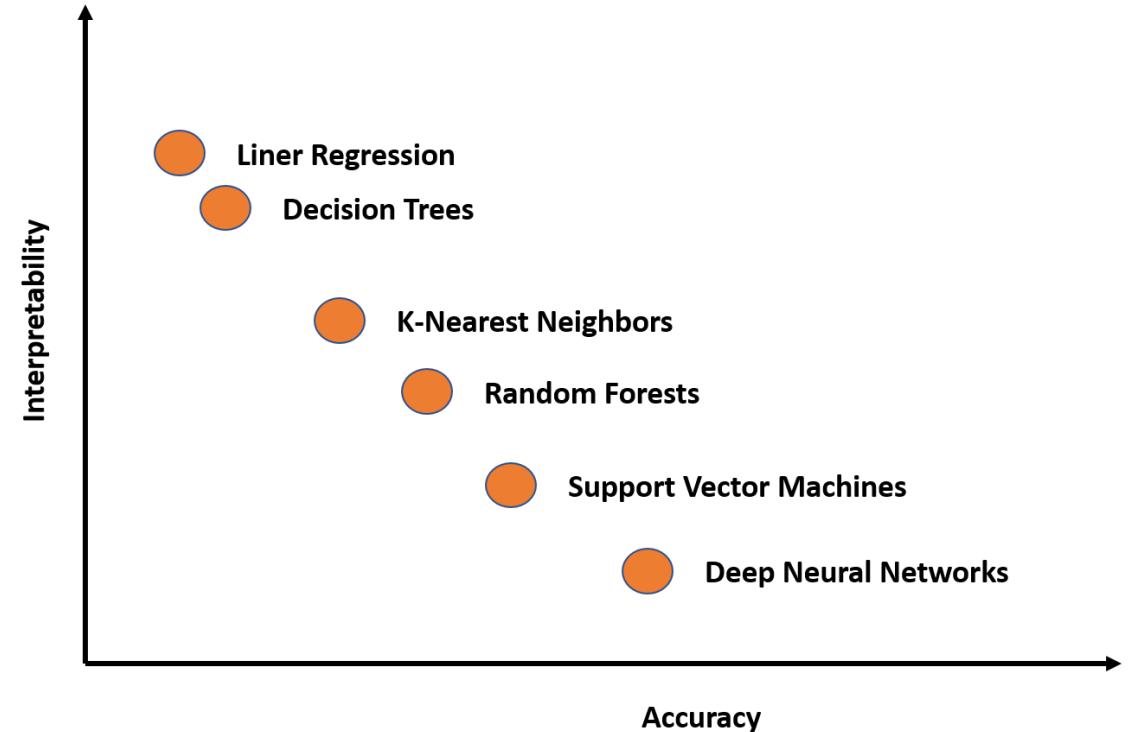


60

- Critics of ML → “black-box” model
- System that can produce valuable output, but which human might not understand
- But carefully constructed ML model can be verifiable and understandable
- Interpretable model = can tell how the model came to a decision

# Explainable neural network

- Some ML methods are more interpretable than others
- Linear regression, decision trees, generative additive models are inherently interpretable
- Interpretability often comes at the expense of power and accuracy
- Trade-off between interpretability and accuracy



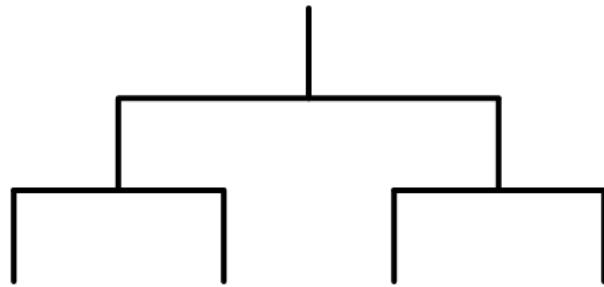
# Explainable neural network

- Intrinsic
  - Explainability built into the model
  - Achieved by restricting the model complexity
  - Examples: linear regression, decision tree, etc.
- Post-hoc
  - Interpretable techniques help to reveal how models make predictions
  - After training
  - Examples: permutation feature importance

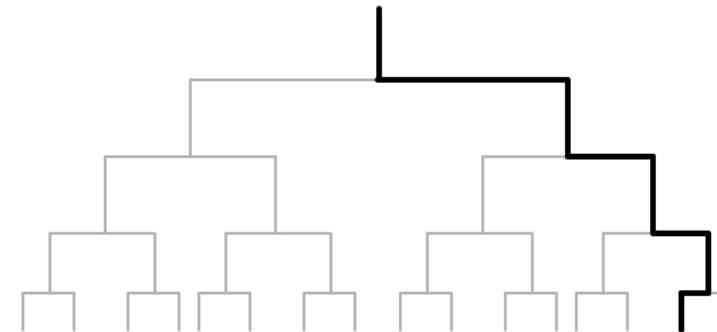
# Explainable neural network

- Global
  - Understanding how the complete model works
  - Examples: Partial dependence plot (PDP), permutation feature importance, etc.
- Local
  - Understanding how a single decision was reached
  - Examples: Local surrogate model (LIME), SHAP

63



Global

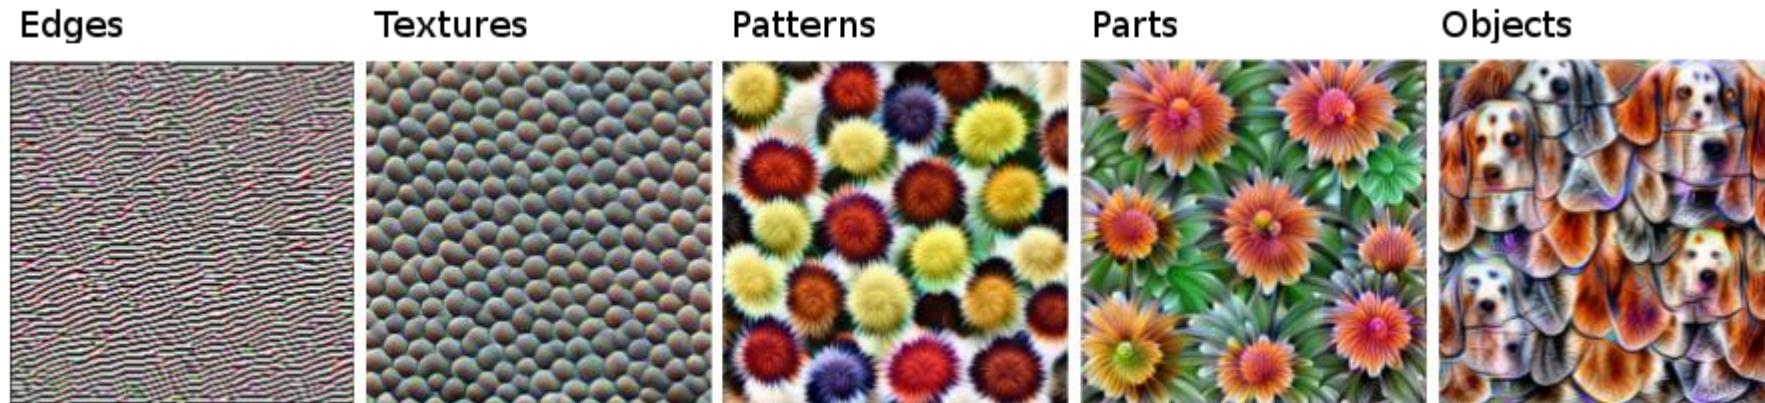


Local

# Explainable neural network

- CNN for image classification
  - Visualize features to understand what the network is seeing
  - Each layer learns a specific feature
  - Higher layers learn more complex features based on simpler features learned by lower layers

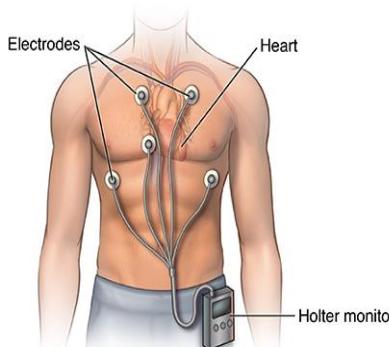
64



# Explainable neural network

Liu et al. 2022, Multiclass arrhythmia detection and classification from photoplethysmography signals using a deep convolutional neural network  
Journal of the American Heart Association

- GOAL
  - Classify **multiclass arrhythmia** types using PPG and deep CNN



## Multiclass Arrhythmia Detection and Classification From Photoplethysmography Signals Using a Deep Convolutional Neural Network

Zengding Liu , MPhil<sup>1</sup>; Bin Zhou, MD, PhD<sup>2</sup>; Zhiming Jiang, MPhil<sup>3</sup>; Xi Chen, BE<sup>4</sup>; Ye Li, PhD<sup>5</sup>; Min Tang , MD, PhD<sup>1</sup>; Fen Mao , PhD<sup>1</sup>

**BACKGROUND:** Studies have reported the use of photoplethysmography signals to detect atrial fibrillation; however, the use of photoplethysmography signals in classifying multiclass arrhythmias has rarely been reported. Our study investigated the feasibility of using photoplethysmography signals and a deep convolutional neural network to classify multiclass arrhythmia types.

**METHODS AND RESULTS:** ECG and photoplethysmography signals were collected simultaneously from a group of patients who underwent radiofrequency ablation for arrhythmias. A deep convolutional neural network was developed to classify multiple rhythms based on 10-second photoplethysmography waveforms. Classification performance was evaluated by calculating the area under the microaverage receiver operating characteristic curve, overall accuracy, sensitivity, specificity, and positive and negative predictive values against annotations on the rhythm of arrhythmias provided by 2 cardiologists consulting the ECG results. A total of 228 patients were included; 118 217 pairs of 10-second photoplethysmography and ECG waveforms were used. When validated against an independent test data set (23 384 photoplethysmography waveforms from 45 patients), the DCNN achieved an overall accuracy of 85.0% for 6 rhythm types (sinus rhythm, premature ventricular contraction, premature atrial contraction, ventricular tachycardia, supraventricular tachycardia, and atrial fibrillation); the microaverage area under the microaverage receiver operating characteristic curve was 0.978; the average sensitivity, specificity, and positive and negative predictive values were 75.8%, 96.9%, 75.2%, and 97.0%, respectively.

**CONCLUSIONS:** This study demonstrated the feasibility of classifying multiclass arrhythmias from photoplethysmography signals using deep learning techniques. The approach is attractive for population-based screening and may hold promise for the long-term surveillance and management of arrhythmia.

**REGISTRATION:** URL: [www.chictr.org.cn](http://www.chictr.org.cn); Identifier: ChiCTR2000031170.

**Key Words:** arrhythmias ■ deep convolutional neural networks ■ photoplethysmography

**A**rrhythmias affect the quality of life of tens of millions of people worldwide, with as many as one-quarter of Americans over 40 years old

developing cardiac arrhythmias.<sup>1</sup> Arrhythmias are associated with high risks of stroke,<sup>2</sup> heart failure,<sup>3</sup> and even sudden cardiac death. More than 80% of sudden

Correspondence to: Fen Mao, PhD, Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences, Shenzhen, Guangdong 518055, People's Republic of China. Email: [fenmao@szit.ac.cn](mailto:fenmao@szit.ac.cn) and Min Tang, MD, PhD, State Key Lab of Cardiovascular Disease, Fuwai Hospital, National Center for Cardiovascular Disease, Chinese Academy of Medical Sciences & Peking Union Medical College, No. 167 North 3rd Road, Xicheng District, Beijing 100037, People's Republic of China. Email: [mtang@cmifw.hscs.ac.cn](mailto:mtang@cmifw.hscs.ac.cn)

12. Li et al. *Am J Cardiol* 2018;121:780-786.

13. Mao and Tang contributed equally as senior authors.

Supplementary Material for this article is available at <https://www.ahajournals.org/doi/suppl/10.1161/JAHF.121.202365>

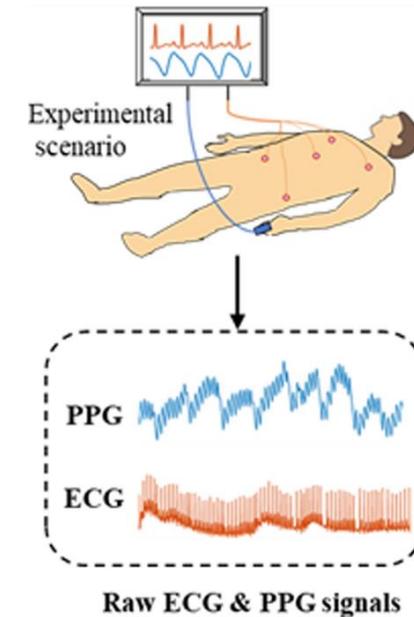
For Sources of Funding and Disclosures, see page 12.

© 2022 The Authors. Published on behalf of the American Heart Association, Inc., by Wiley. This is an open access article under the terms of the Creative Commons Attribution Non-Commercial License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non commercial and no modifications or adaptations are made.

JAHF is available at: [www.ahajournals.org/jahf](http://www.ahajournals.org/jahf)

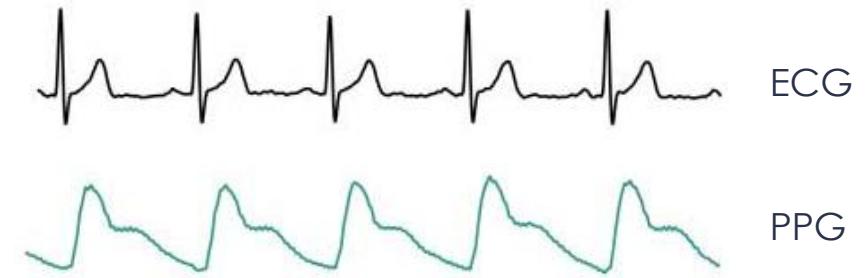
# Explainable neural network

- DATA
  - 242 patients with arrhythmia receiving radiofrequency catheter ablation
  - At Fuwai Hospital, Chinese Academy of Medical Sciences
  - fingertip **PPG**
  - 3-lead **ECG** for reference



# Explainable neural network

- Preprocessing
  1. Signal downsampling (250 to 100 Hz)
  2. Denoising with bandpass filtering
    - ECG from 0.5 to 50 Hz
    - PPG from 0.5 to 10 Hz
  3. Segmentation
    - Split into 10s non-overlapping segments
  4. Segment normalization
  5. Signal quality exclusion
    - Remove segments with artifacts in PPG
    - Remove segments with noisy ECG reference
  6. Labels
    - Annotation of ECG segments from 2 cardiologists
    - Create labels according to ECG annotations



67

118'217 labeled 10s  
PPG segments from  
228 patients

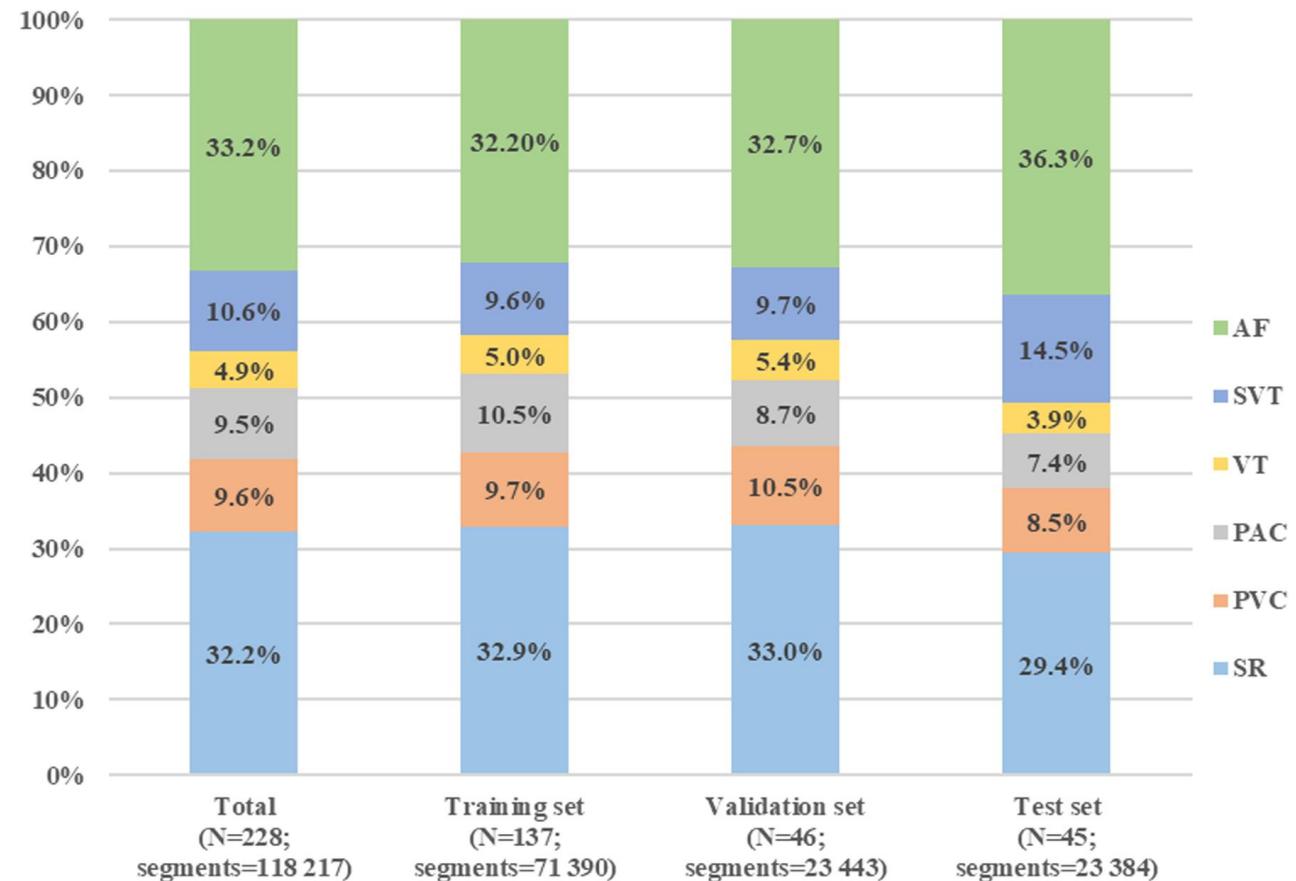
# Explainable neural network

- Each segment has only 1 identified rhythm type

<b>Rhythm type</b>	<b>Number of segments</b>
Sinus rhythm	38081
Premature ventricular contraction (PVC)	11372
Premature atrial contraction (PAC)	11248
Ventricular tachycardia (VT)	5783
Supraventricular tachycardia (SVT)	12539
Atrial fibrillation (AF)	39194

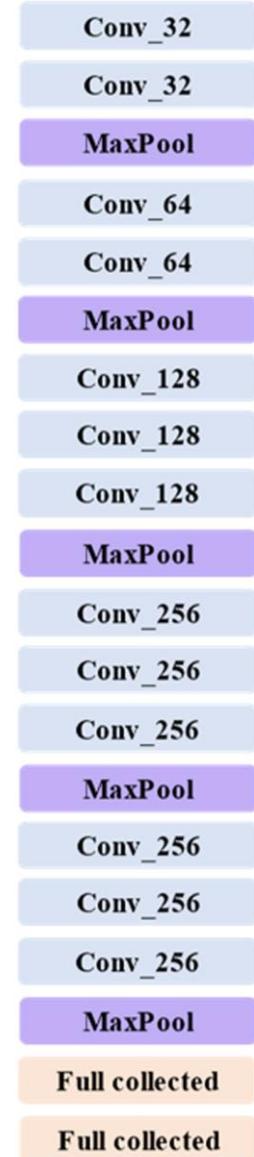
# Explainable neural network

- Split data
- 60% of patients in training set
- 20% of patients in validation set
- 20% of patients in test set

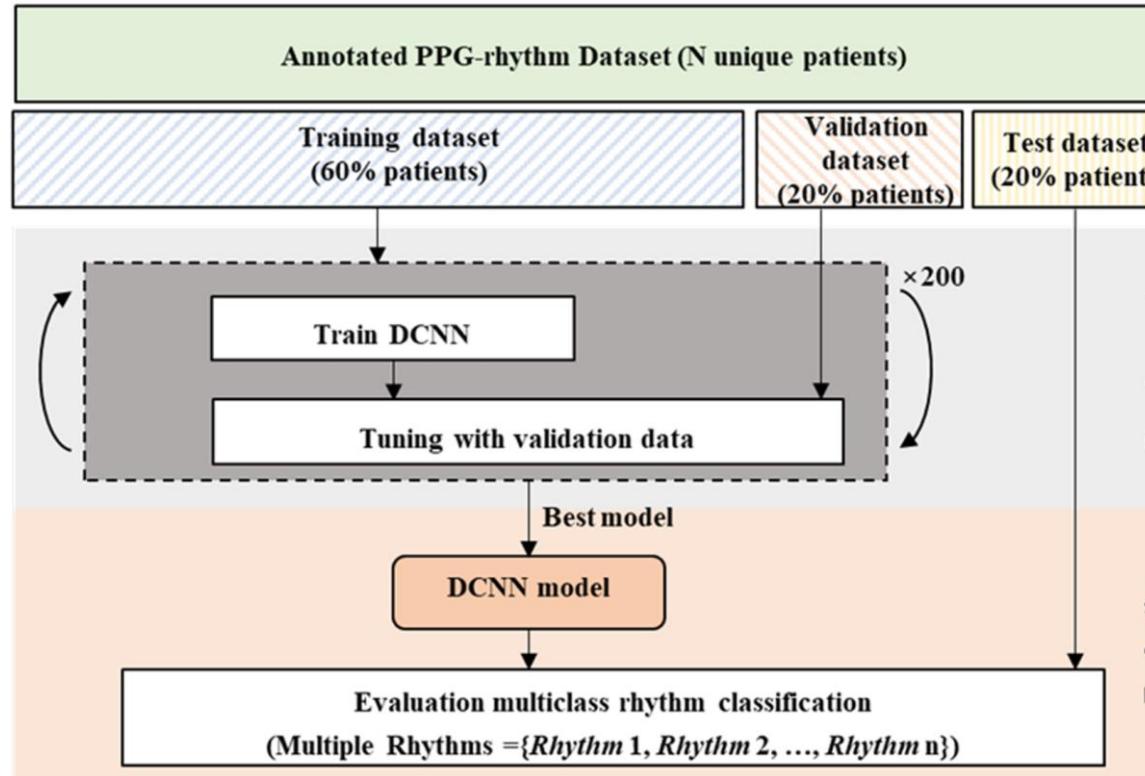


# Explainable neural network

- Multiclass rhythm classification from PPG
  - Model
    - Based on **VGGNet-16** (deep CNN)
    - Adapted for 1-dimensional input signal
    - 13 convolutional layers
    - 5 max-pooling layers
    - 3 fully connected layers
  - Input: 10-second PPG segment
  - Output: label prediction of one of the rhythm type



# Explainable neural network

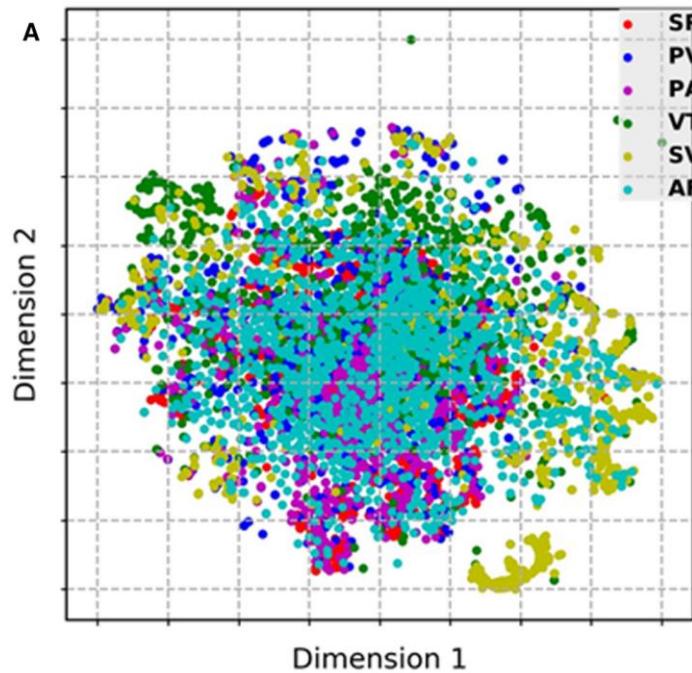


# Explainable neural network

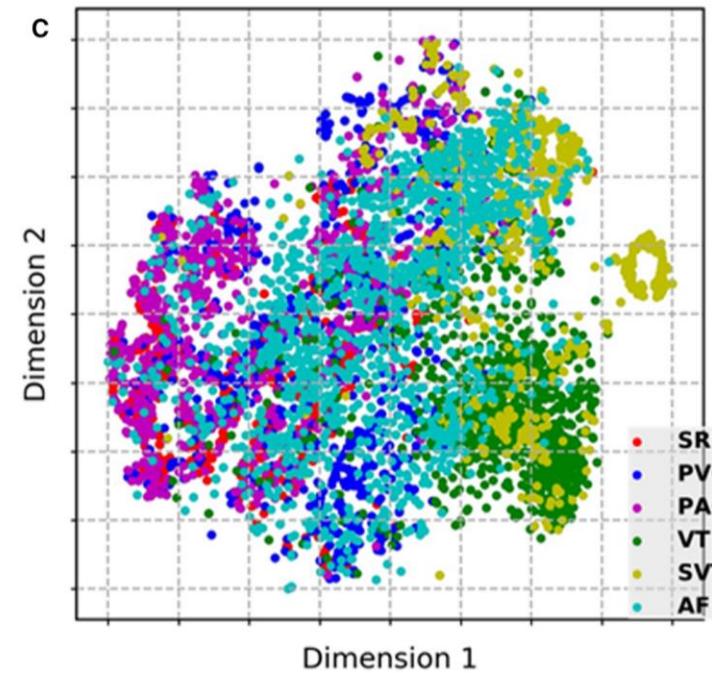
- Visualization of proposed DCNN
  - Improve understanding of model decision
  - Highlight regions in PPG important for the model prediction
  - **t-distributed stochastic neighbor embedding** (t-SNE)
    - Non-linear dimensionality reduction technique
    - Embedding high-dimensional data for visualization in low dimensional space
  - **Guided gradient-weighted class activation mapping** (Grad-CAM)
    - Fine-grained guided backpropagation
    - Gradient-weighted class activation mapping
    - Create high-resolution class-discriminative heatmap from final convolutional layer

# Explainable neural network

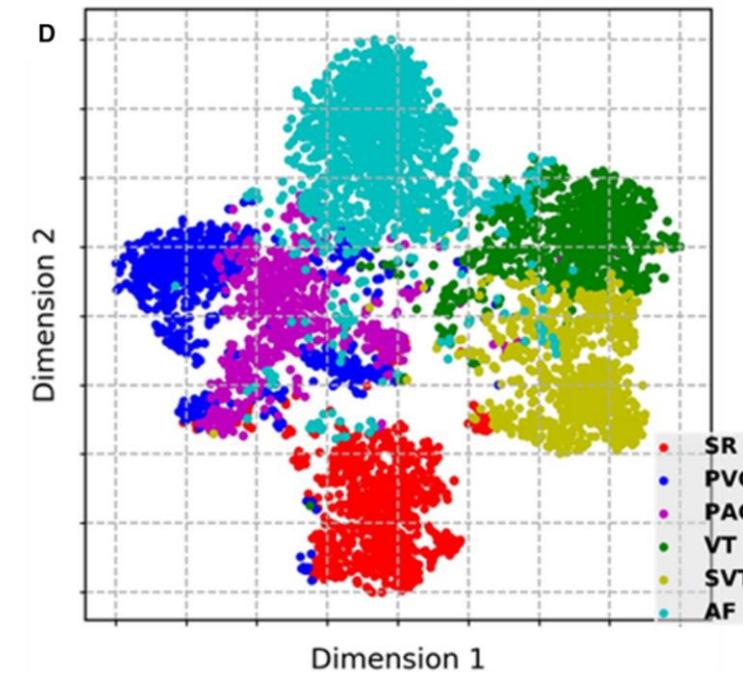
- t-distributed stochastic neighbor embedding



Layer 2



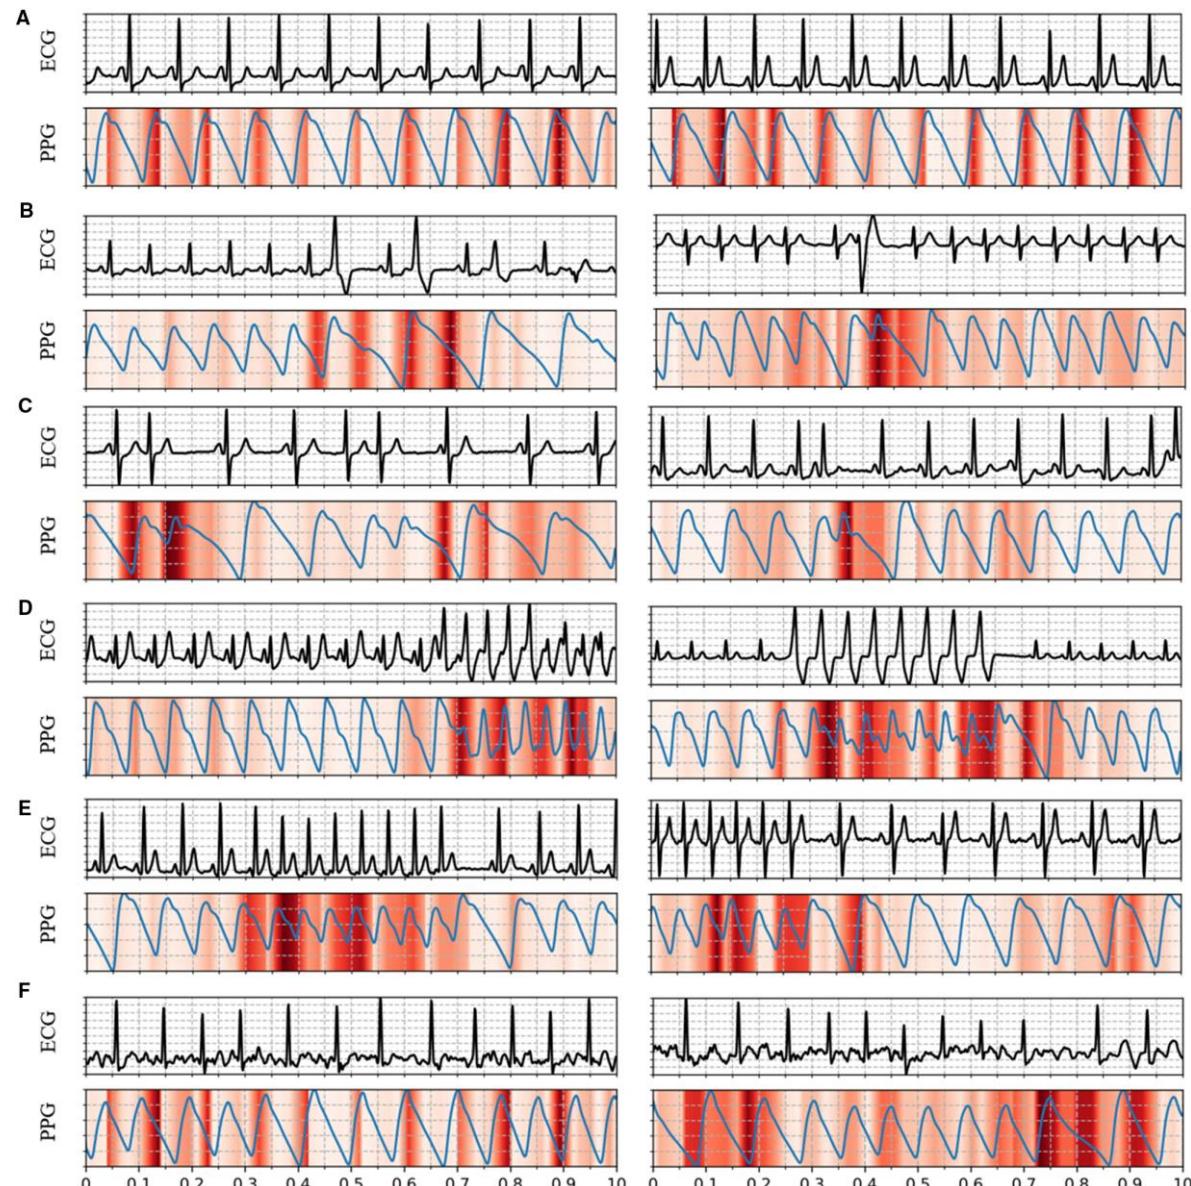
Layer 7



Layer 13

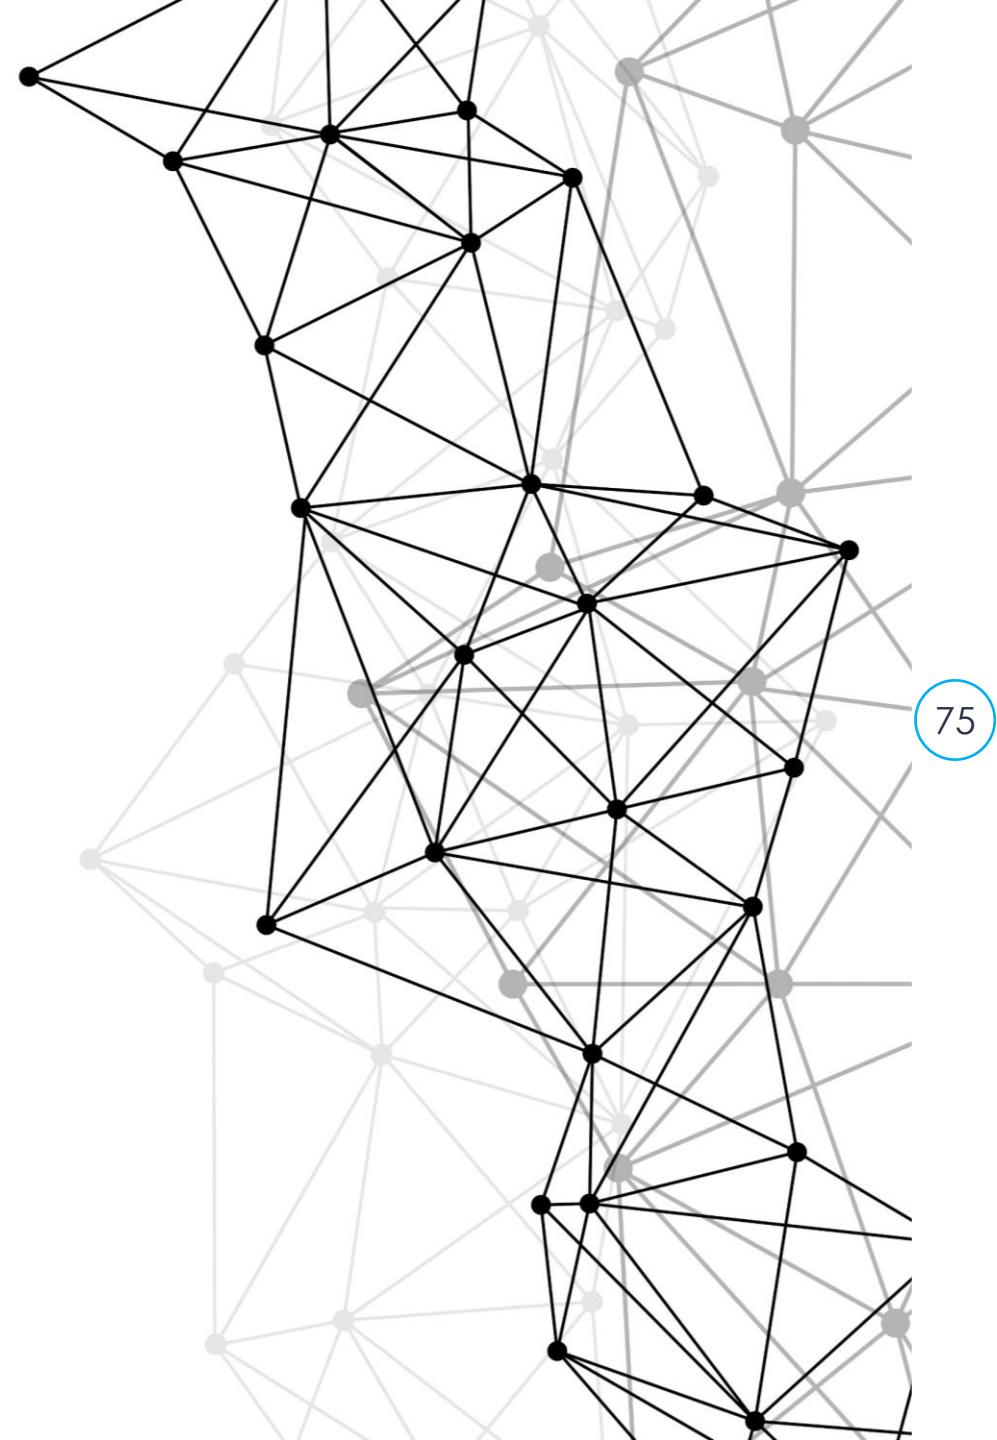
# Explainable neural network

- Sinus rhythm
- PVC  
Premature ventricular contraction
- PAC  
Premature atrial contraction
- VT  
Ventricular tachycardia
- SVT  
Supraventricular tachycardia
- AF  
Atrial fibrillation



# Content

- Neural network recap
- Regularization
- Different architectures
- Labs



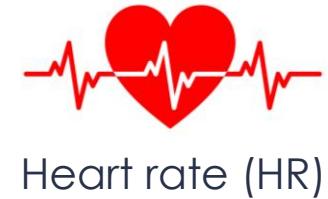
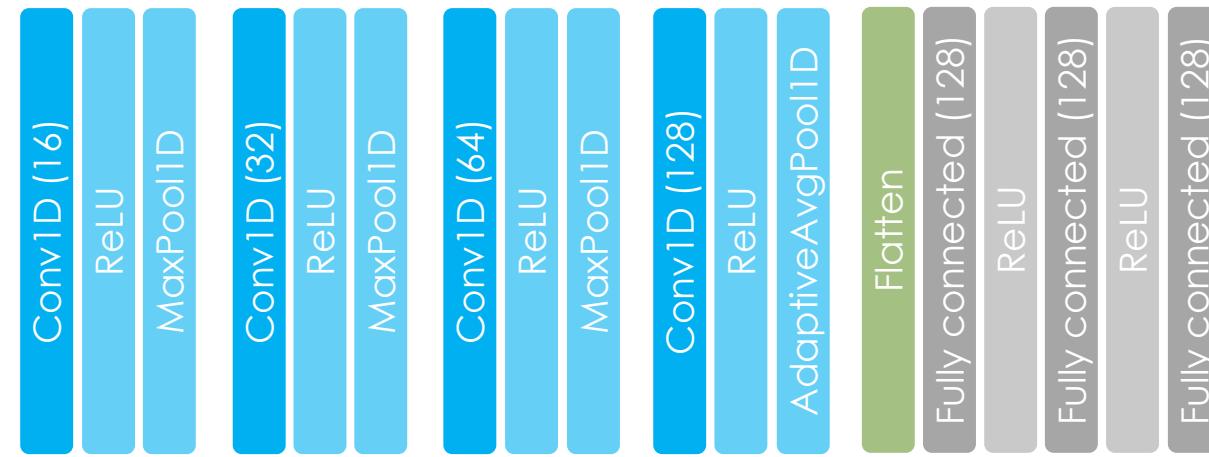
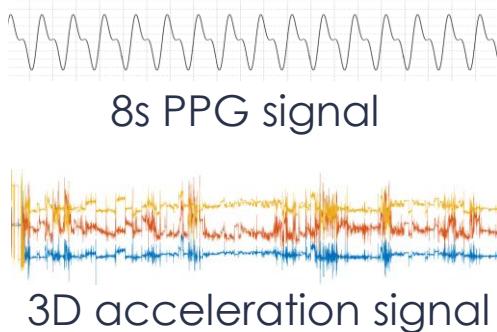
# Lab – Instructions

- Submit report as **single PDF file**
- Recommended to work in groups of **3 students**
- You can prepare one single report for the group  
(name1\_name2\_name3\_lab\_NN2.pdf)
- But every member must upload the file on Moodle
- Python code is given and provided as **Jupyter notebooks**
- This practical session is not focused on coding but on questions testing your understanding and interpretation of the results.
- **2 exercises** in this lab session on real-life biomedical problems.

# Lab – Heart rate classification

- GOAL: Estimate heart rate (HR) from PPG and acceleration signal
  - Focus on sitting and walking
- DATA:
  - PPG-DaLiA dataset <https://archive.ics.uci.edu/ml/datasets/PPG-DaLiA>
  - 15 subjects
  - Photoplethysmography (PPG) signals
  - 3D acceleration signals
  - Reference HR computed from ECG signal

# Lab – Heart rate classification



## Model training

Adam optimizer  
Mean squared error loss  
30 epochs  
Batch size = 100  
Learning rate = 0.0001

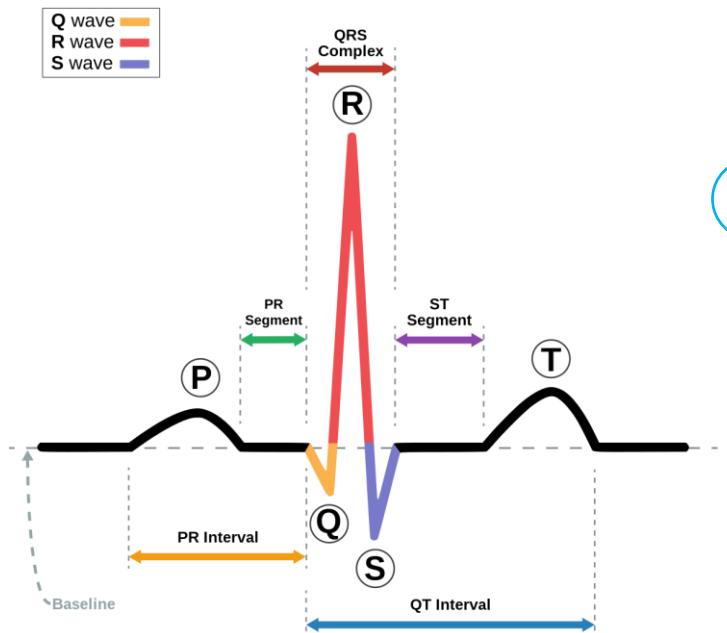
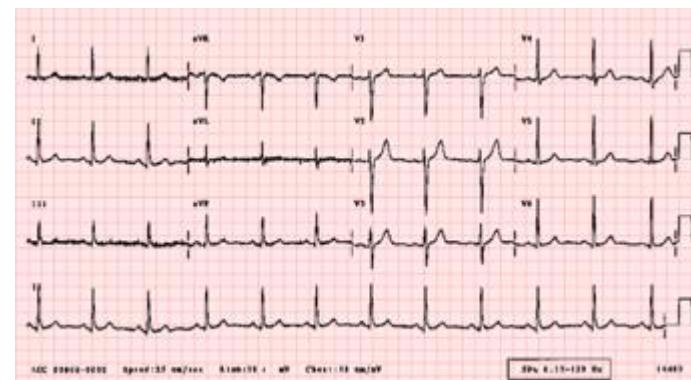
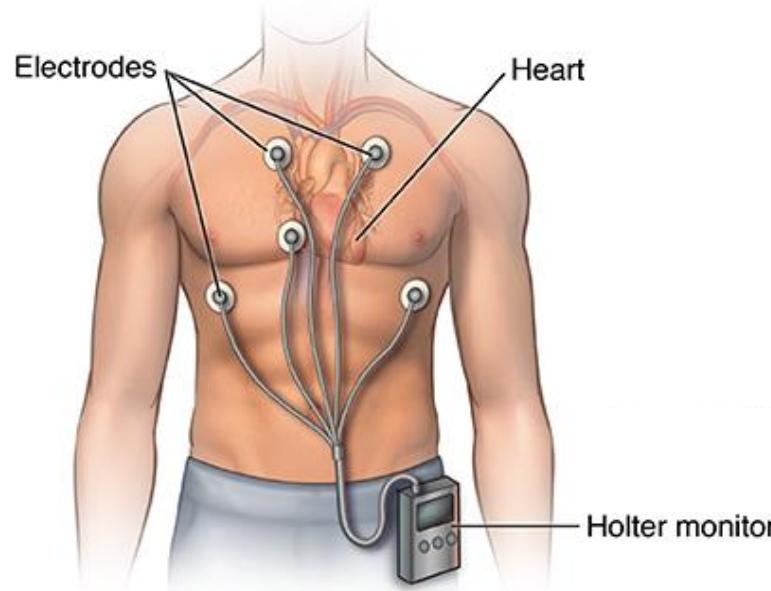
With or without acceleration signal  
Add batch normalization  
Add dropout

# Lab – ECG rhythm classification

- **GOAL** – Train NN to classify different cardiac rhythms from single-lead ECG signals
  - Normal sinus rhythm
  - Sinus bradycardia
  - Sinus tachycardia
  - Atrial fibrillation
  - Atrial Flutter
- **DATA**
  - Subset of large scale 12-lead ECG database for arrhythmia study  
<https://physionet.org/content/ecg-arrhythmia/1.0.0/>
  - 1500 single-lead (lead II) ECG signals of each rhythm

# Lab – ECG rhythm classification

- **ECG** – Recording of the heart's electrical activity



# Lab – ECG rhythm classification

- **Normal sinus rhythm**

- P wave before QRS complex,
- Regular rhythm
- Rate = 60-100 bpm at rest



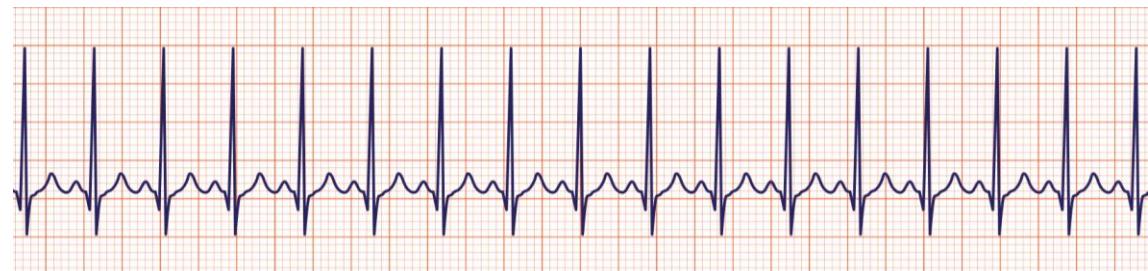
- **Sinus bradycardia**

- P wave before QRS complex
- Regular rhythm
- Rate < 60 bpm at rest



- **Sinus tachycardia**

- P wave before QRS complex
- Regular rhythm
- Rate > 100 bpm at rest



# Lab – ECG rhythm classification

- **Atrial fibrillation**

- No distinct P wave before QRS complex
- Very irregular rhythm



- **Atrial flutter**

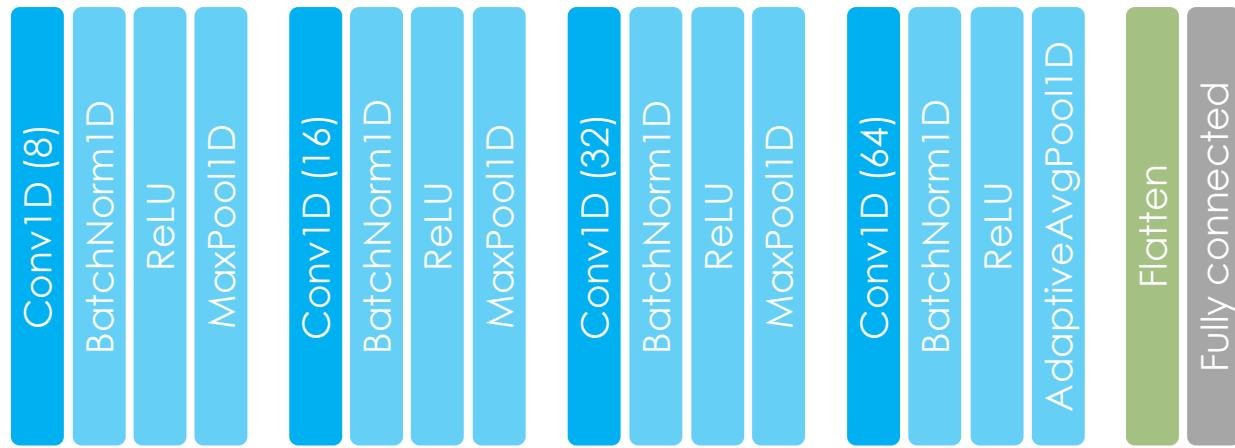
- “Sawtoothed” pattern of P wave
- Irregular rhythm



# Lab – ECG rhythm classification



10s ECG signal



Encoded rhythm

83

## Model training

Adam optimizer  
Cross entropy loss  
50 epochs  
Batch size = 100  
Learning rate = 0.0001

## Custom architectures

Add convolutional layers  
Add fully connected layers  
Add regularization  
Etc.

# Labs

- Process locally with a virtual environment rather than using noto.
- (Recommended) run the experiments using last week notebooks.
  1. Complete last week lab directory with the new notebooks and the new datasets.
  2. Open a terminal in the directory.
  3. Activate virtual environment.
    - Linux: `source venv/bin/activate`
    - Windows: `venv\Scripts\activate`
  4. Start [JupyterLab](#).  
`python -m jupyter lab`