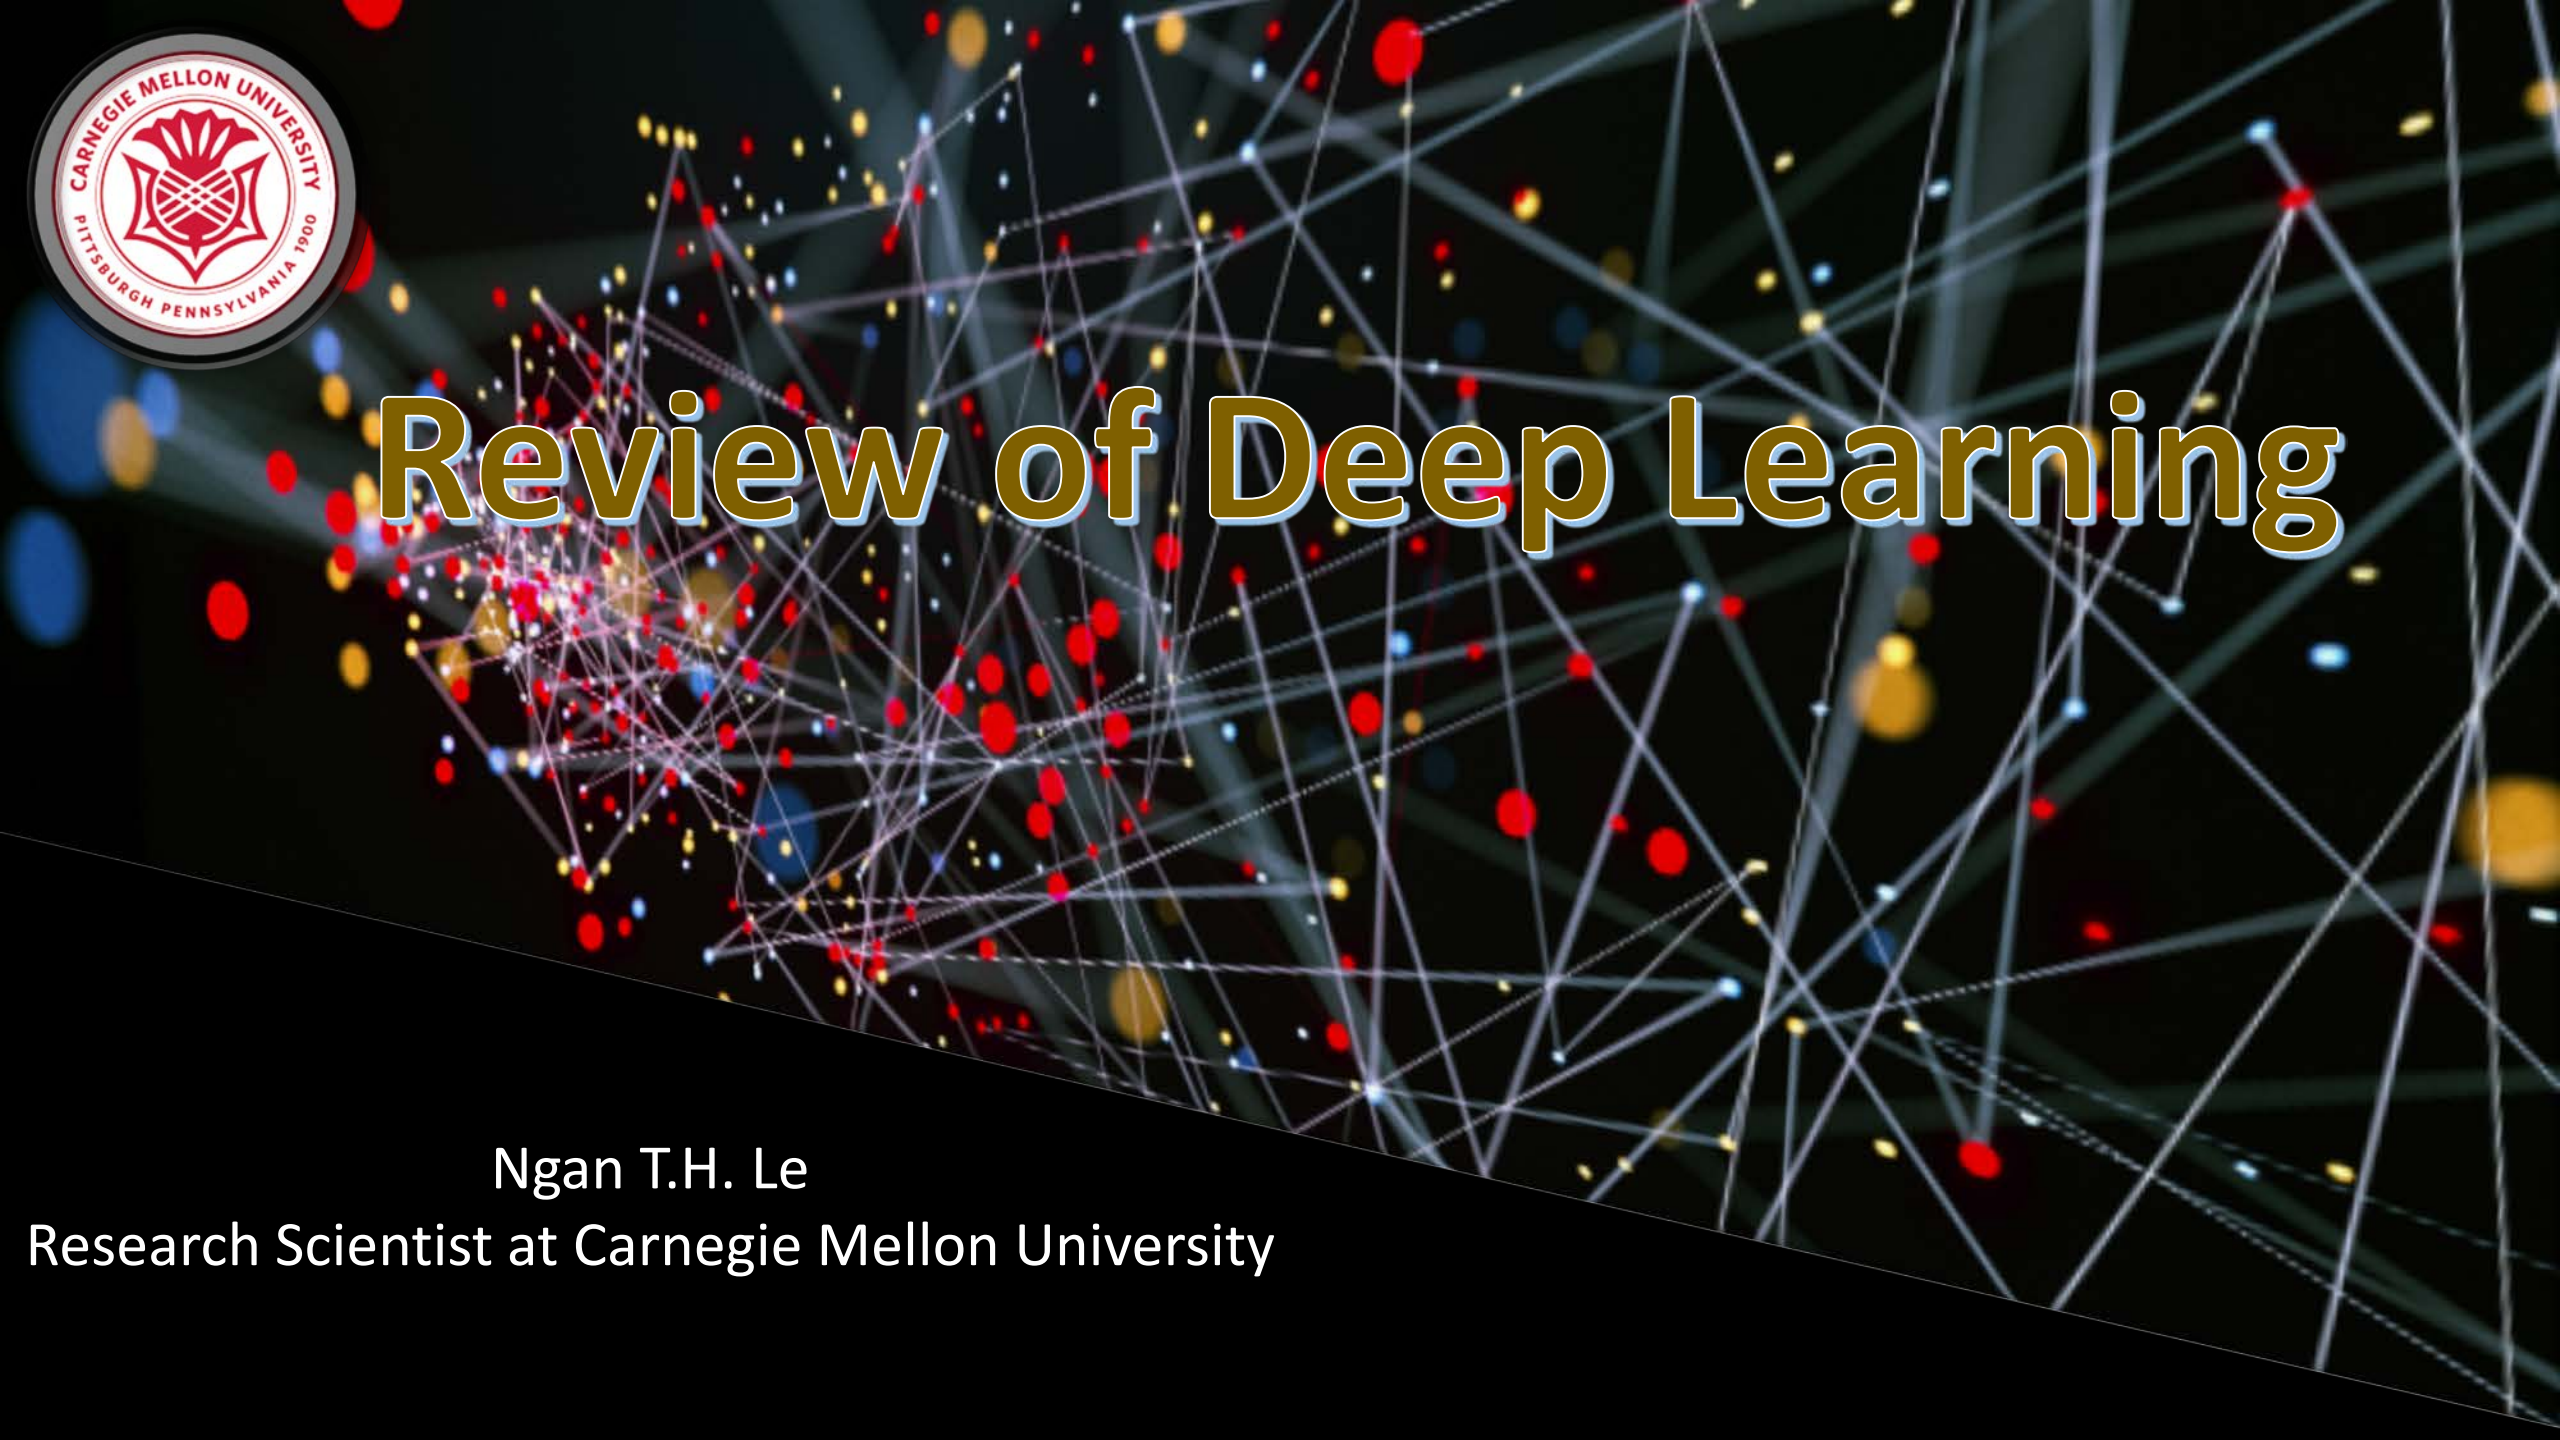

# Review of Deep Learning

Ngan T.H. Le
Research Scientist at Carnegie Mellon University

Deep Neural Network - revisit

Deep Learning in Medical Imaging

Dataset, Contest, Challenging

# Convolutional Neural Networks (CNNs)

Function

Beach?
Sand?
Person?

Function
Function
Function
Function
Function
Function
Function

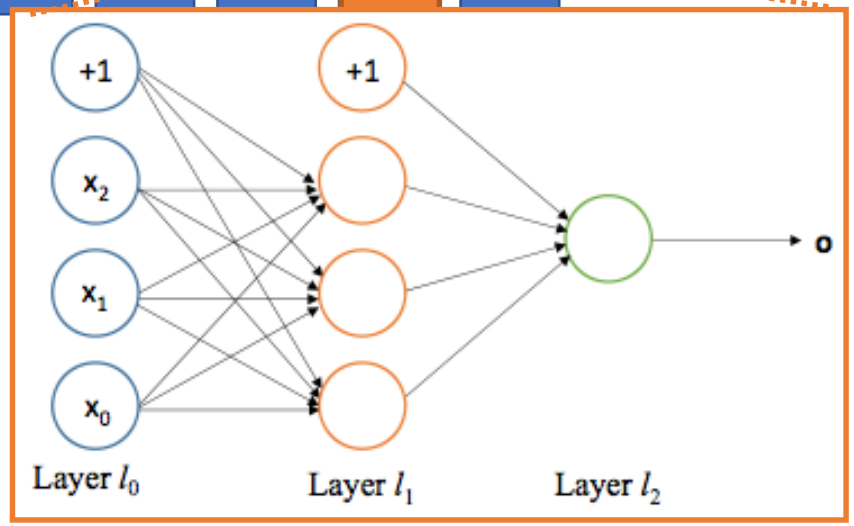


Neuron network

# Convolutional Neural Networks (CNNs)

## VGG16

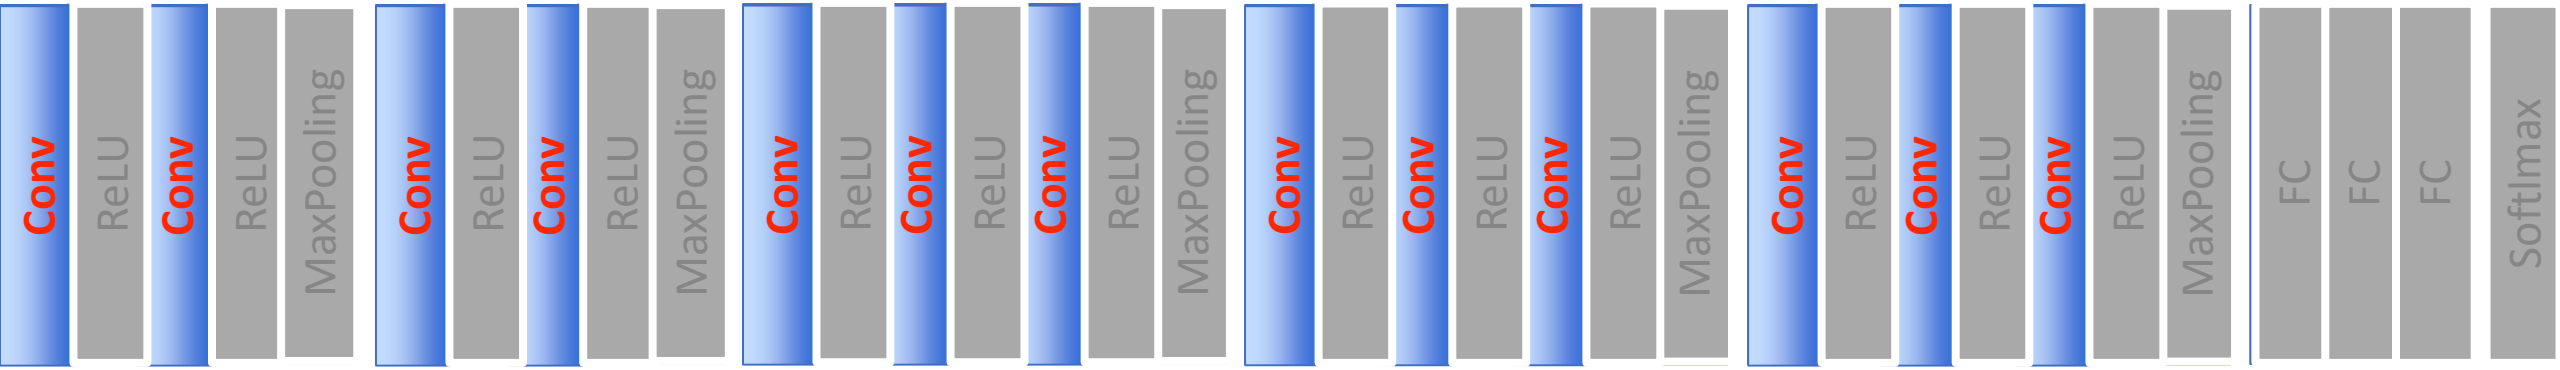


## Convolution layer

To learn feature representations

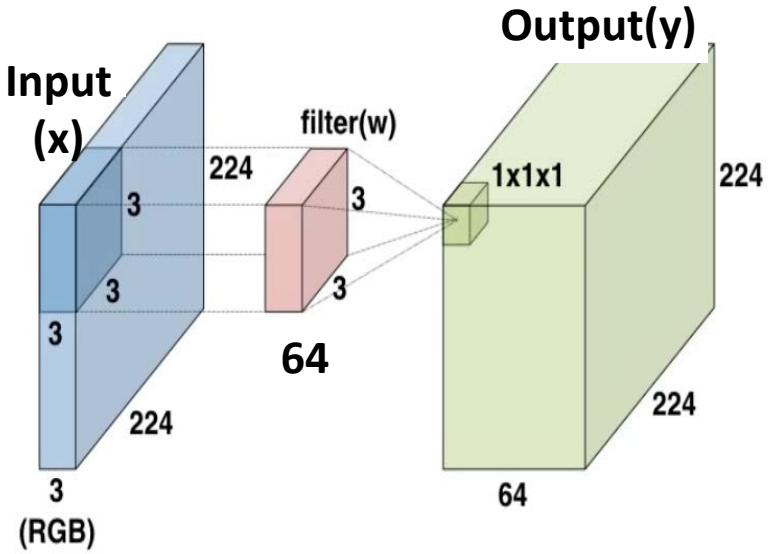


| | | |
|---|---|---|
| a | b | c |
| d | e | f |
| g | h | i |

| | |
|---|---|
| 1 | 2 |
| 3 | 4 |

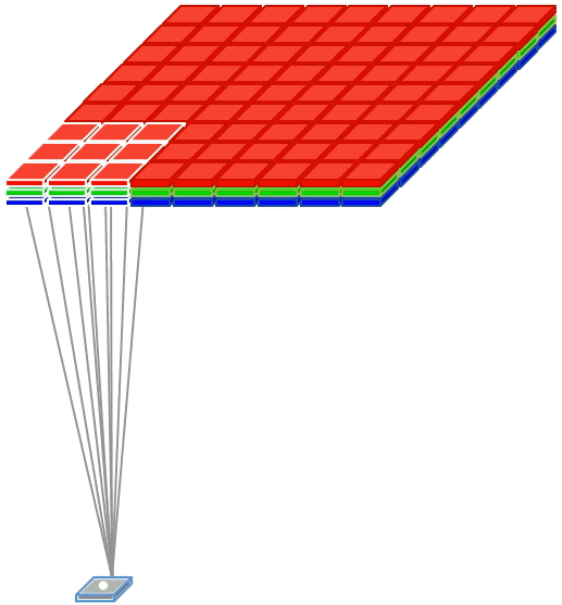

| | |
|---|---|
| 1a + 2b + 3d + 4e | 1b + 2c + 3e + 4f |
| 1d+2e+3g+4h | 1e+2f+3h+4i |

# Convolutional Neural Networks (CNNs)

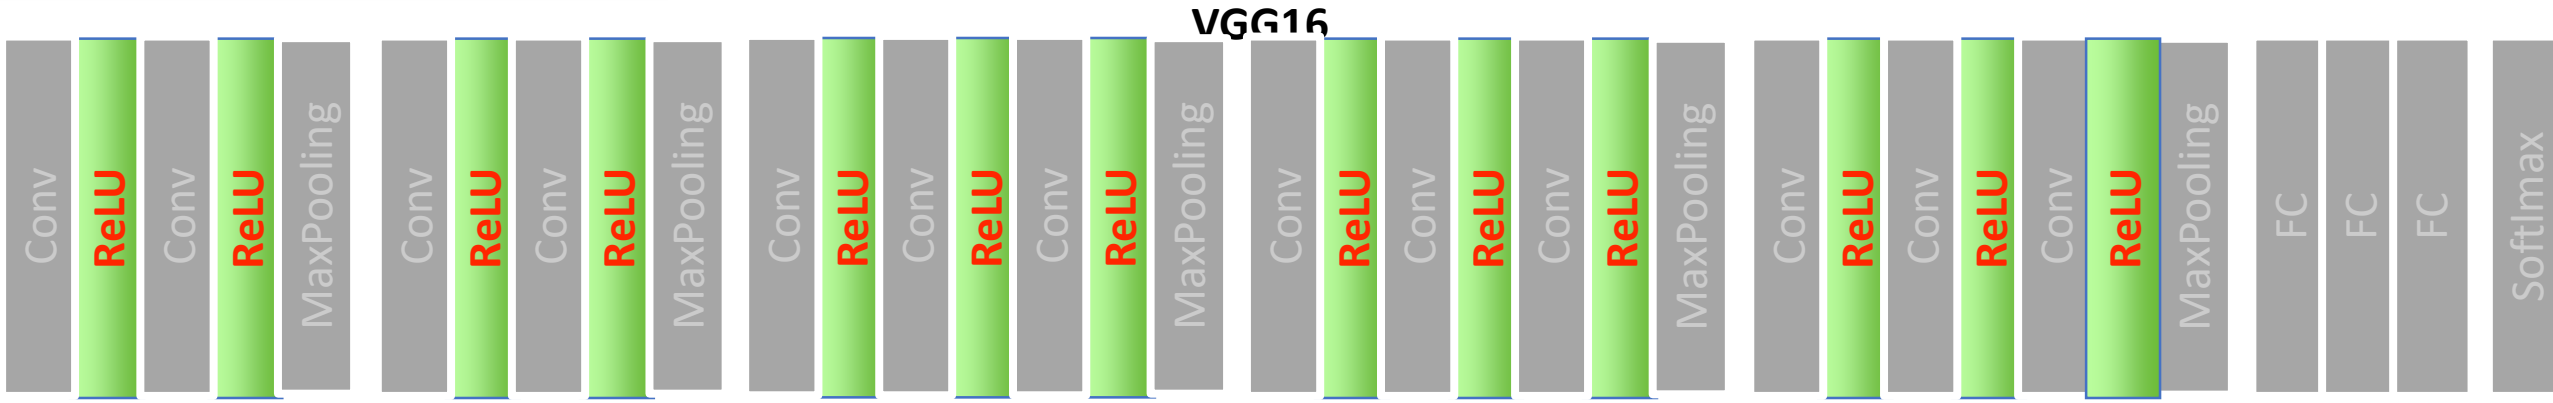


## Activation Layer

- Introduces nonlinearities to CNNs
- Controls how the signal flows from one layer to the next

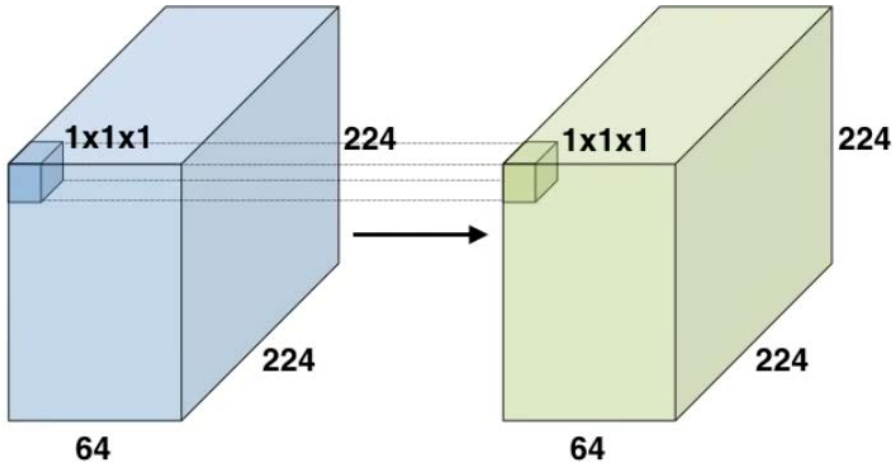


Sigmoid

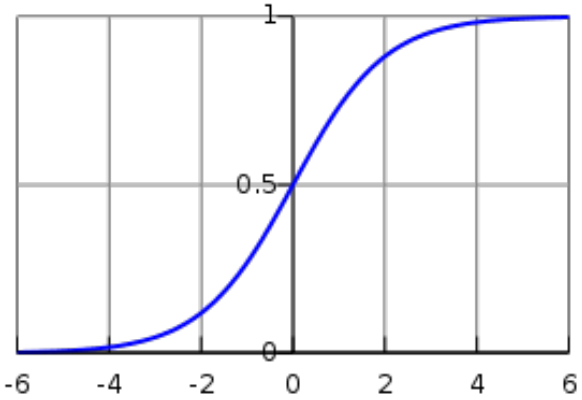


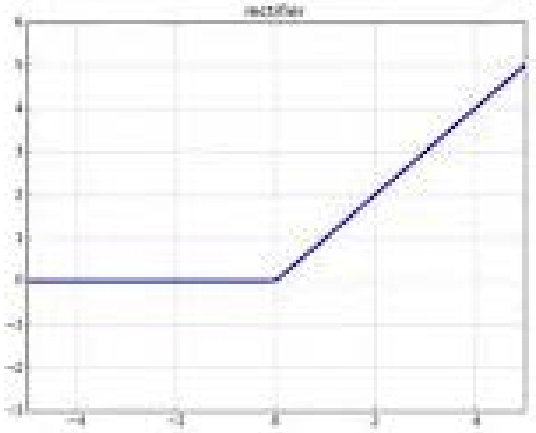

ReLU

Tanh

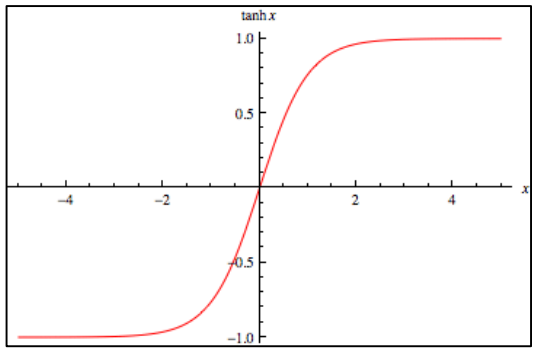

# Convolutional Neural Networks (CNNs)

## VGG16

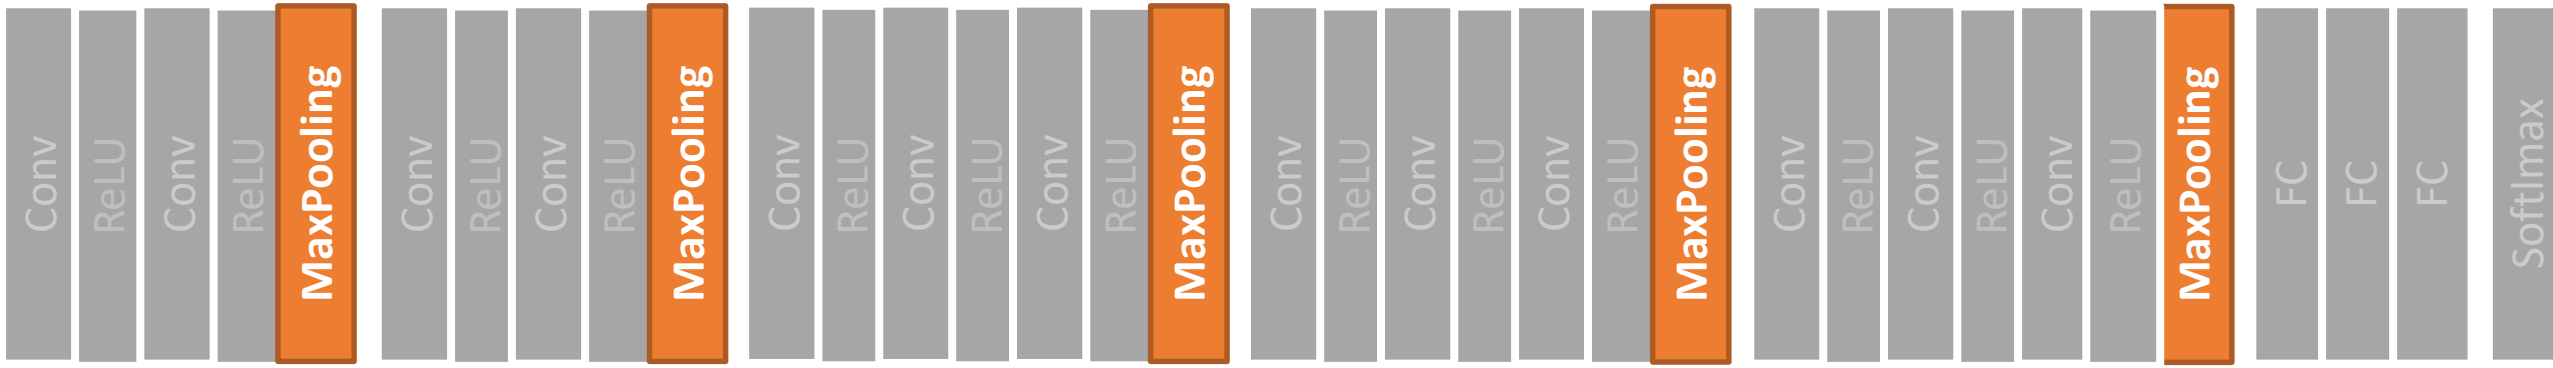


## Pooling Layer

- To smooth the input from the convolutional layer
- Helps to reduce the sensitivity of the filters to noise and variations

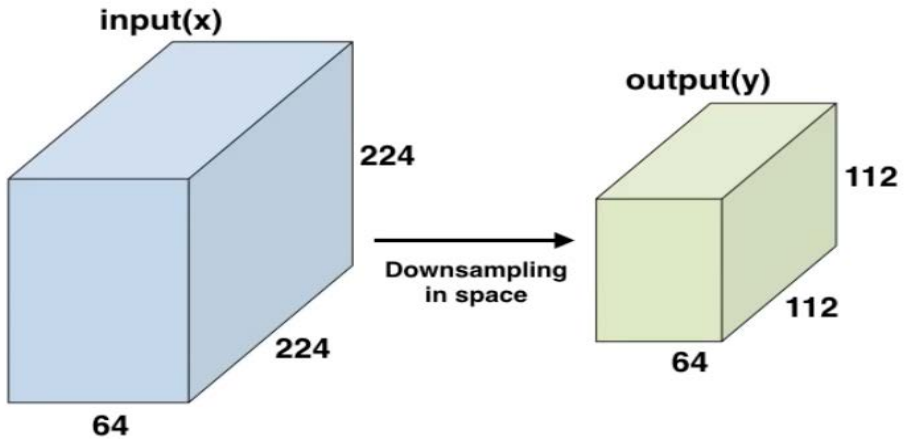


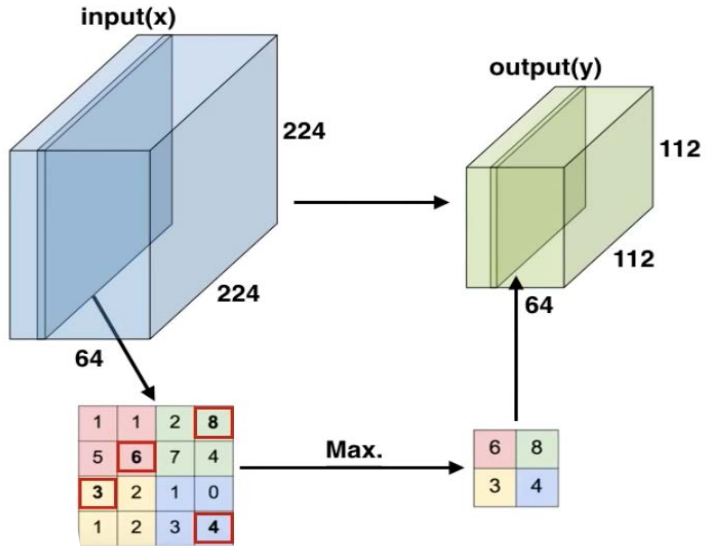

# Convolutional Neural Networks (CNNs)

**VGG16**

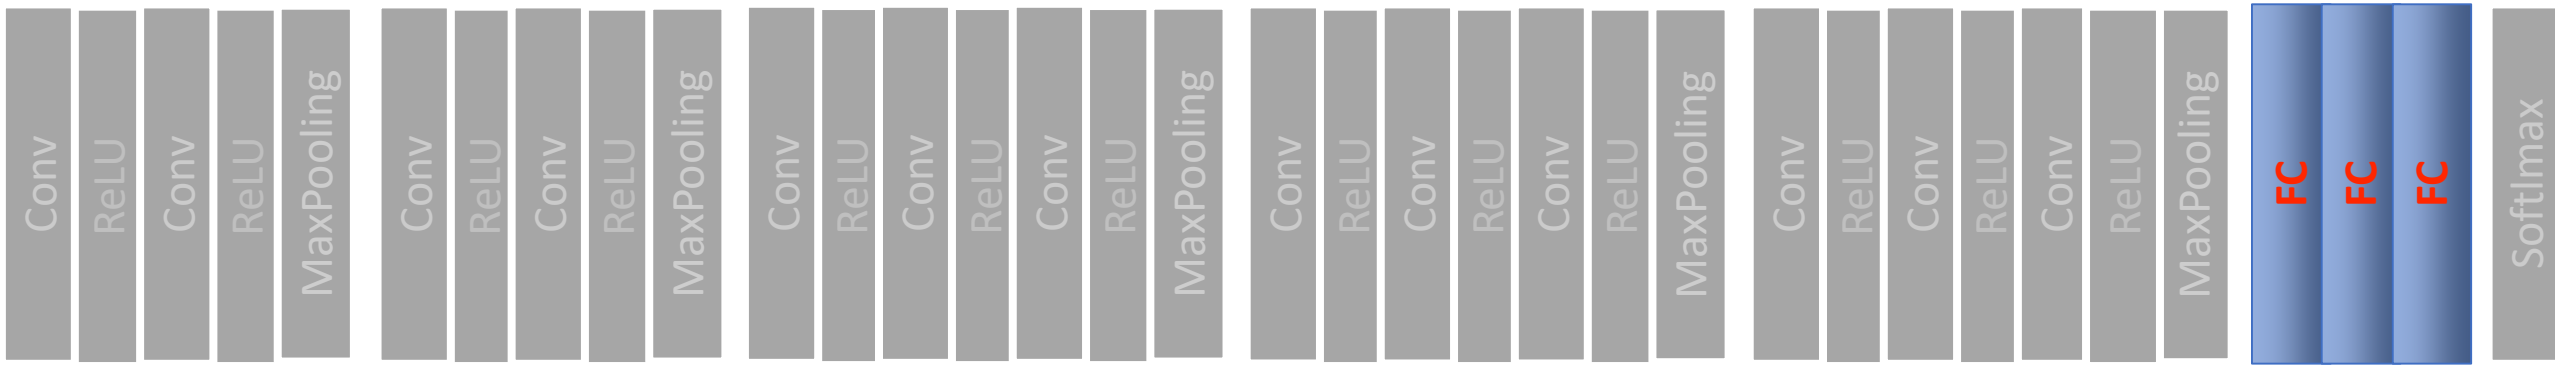


## Fully Connected Layer

- Implies that every neuron in the previous layer is connected to every neuron on the next layer
- Mimics high level reasoning where all possible pathways from the input to output are considered

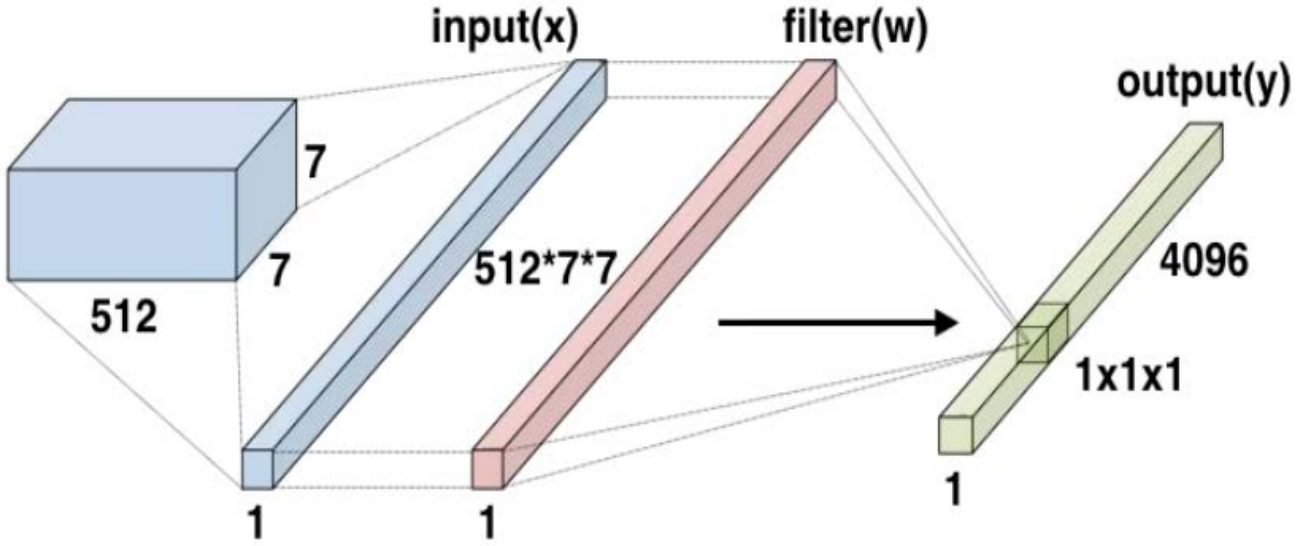

# Convolutional Neural Networks (CNNs)

**VGG16**

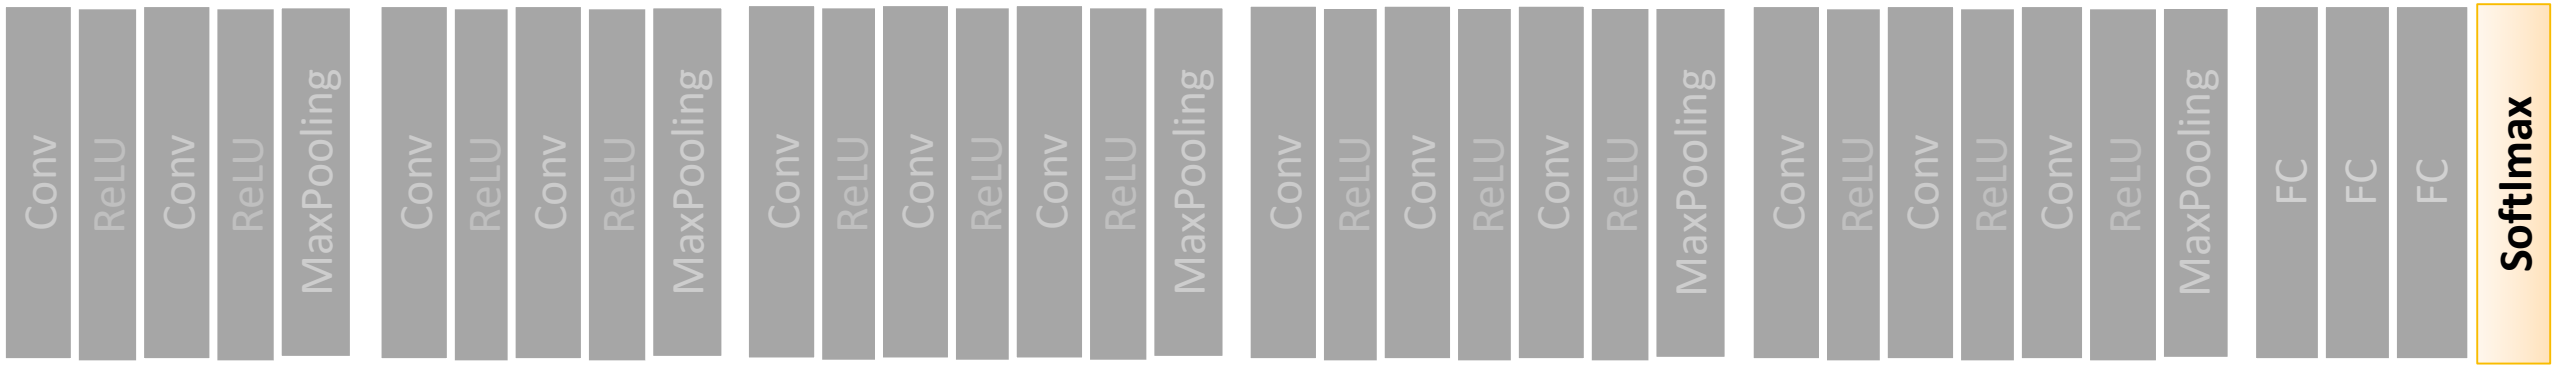


## Loss Layer

Choose an appropriate loss function for a specific task

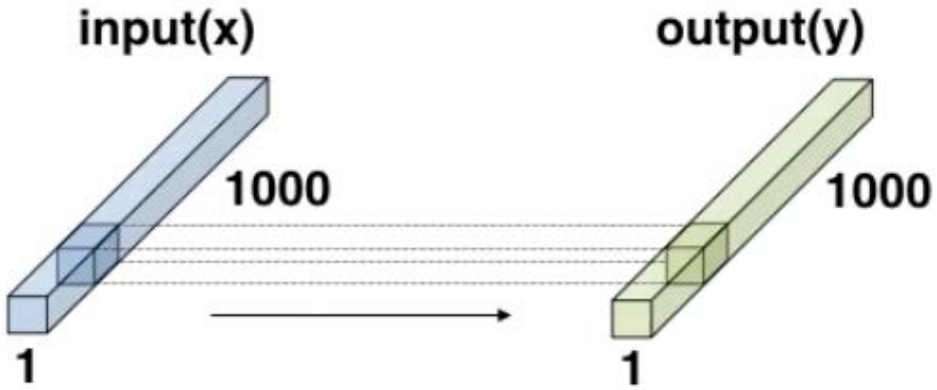


$$y_i = \frac{e^{x_i}}{\sum_j^{1000} e_j^x}$$

# VGG16

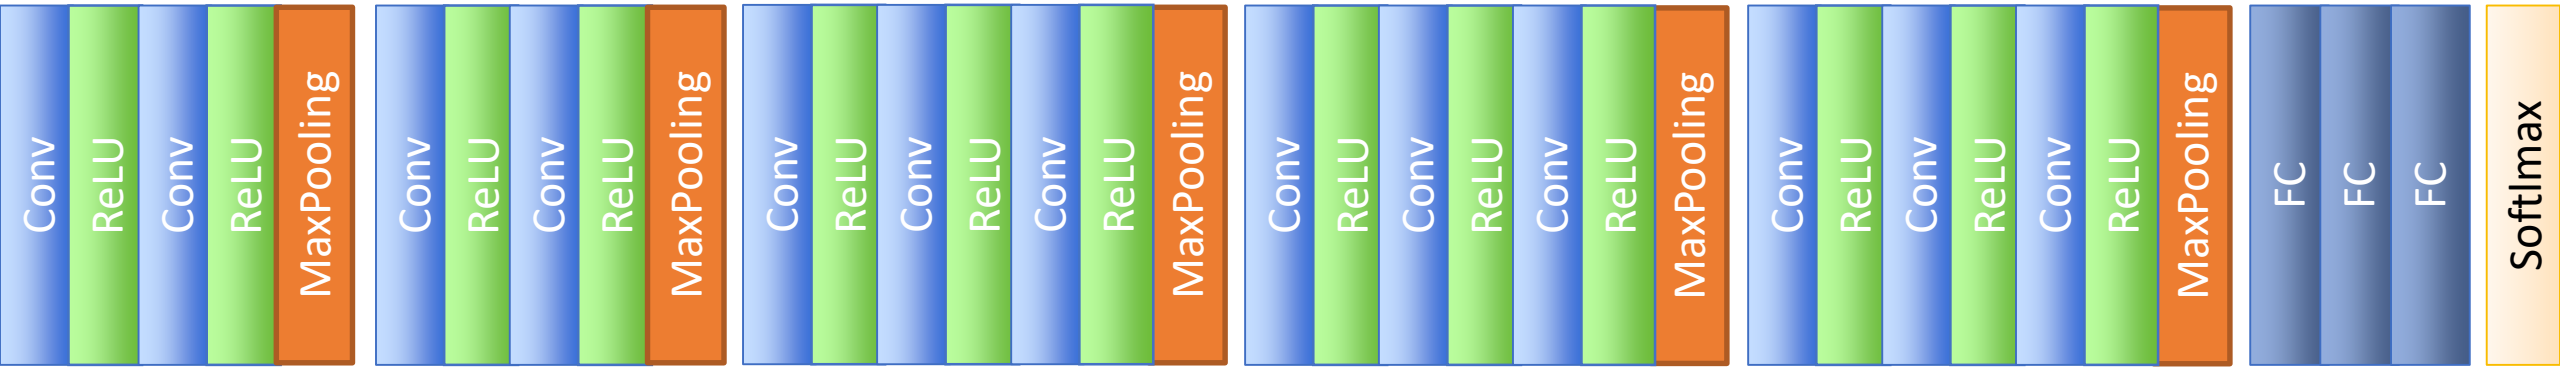


Shallow

Deep

Low level feature

Mid level feature

High level feature

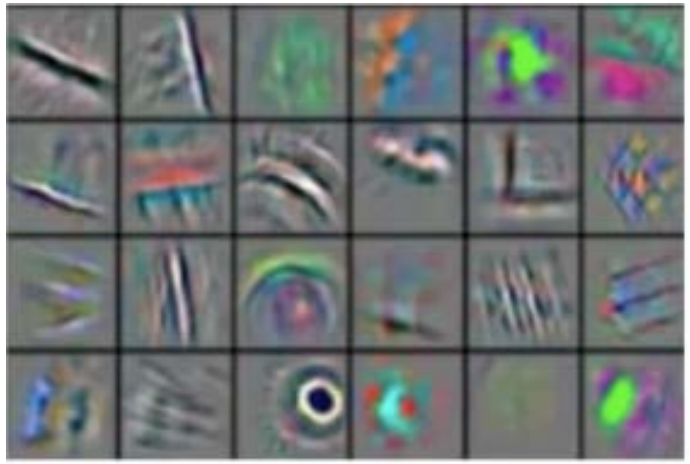

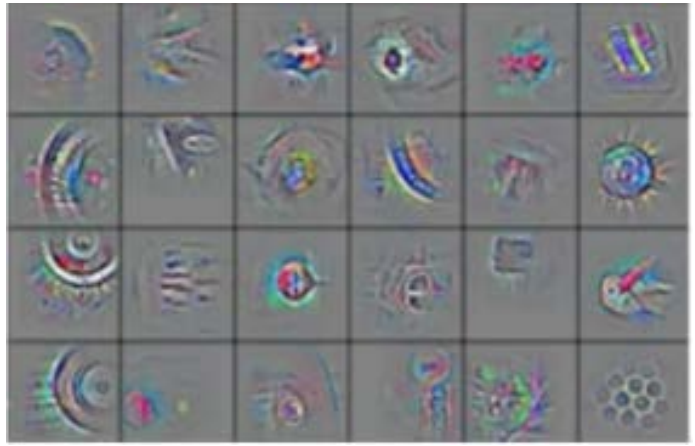

# CNN handles complex nonlinear function?

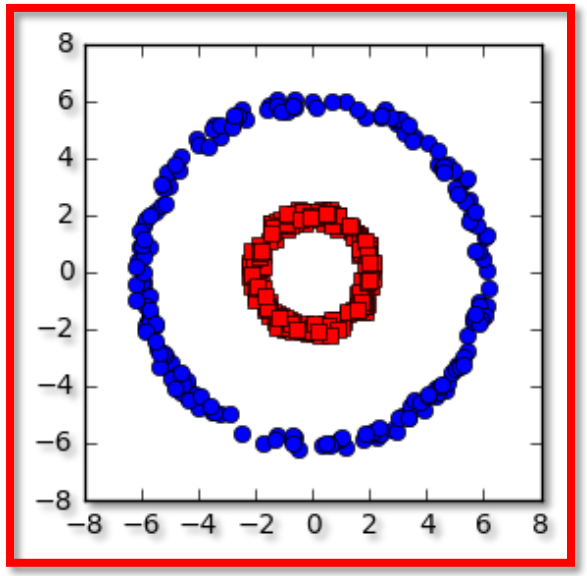

$$Y = \omega X + b$$

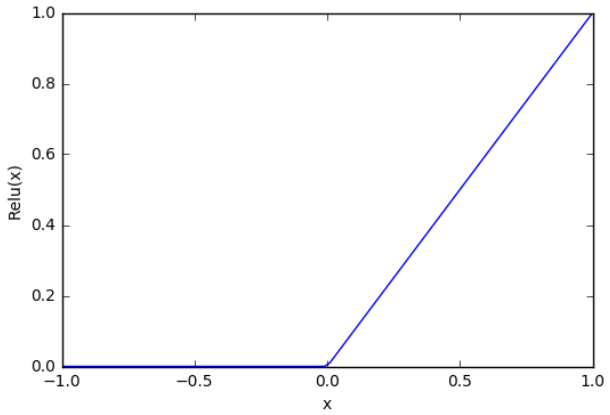


$$\omega_1 x_1 + x_2 + b > 0$$

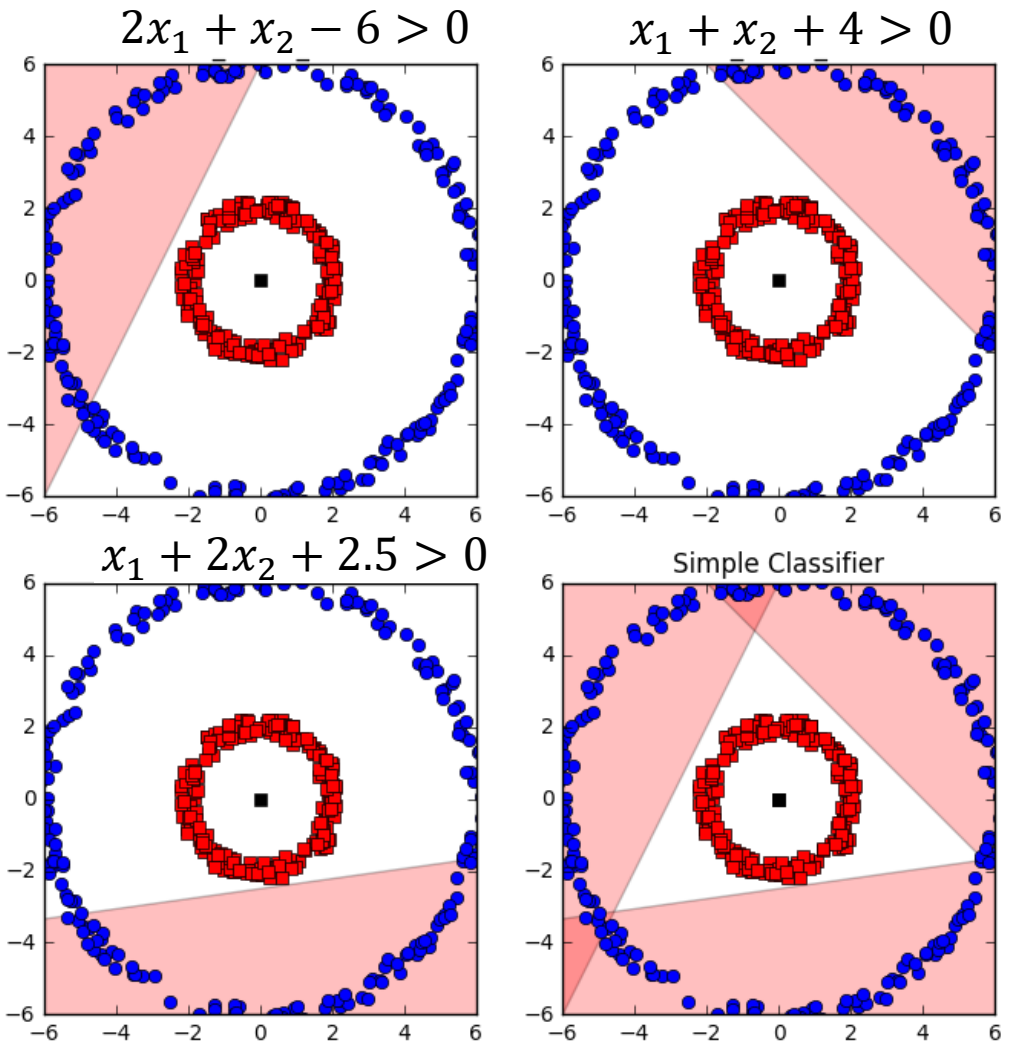


Red > 0

White = 0

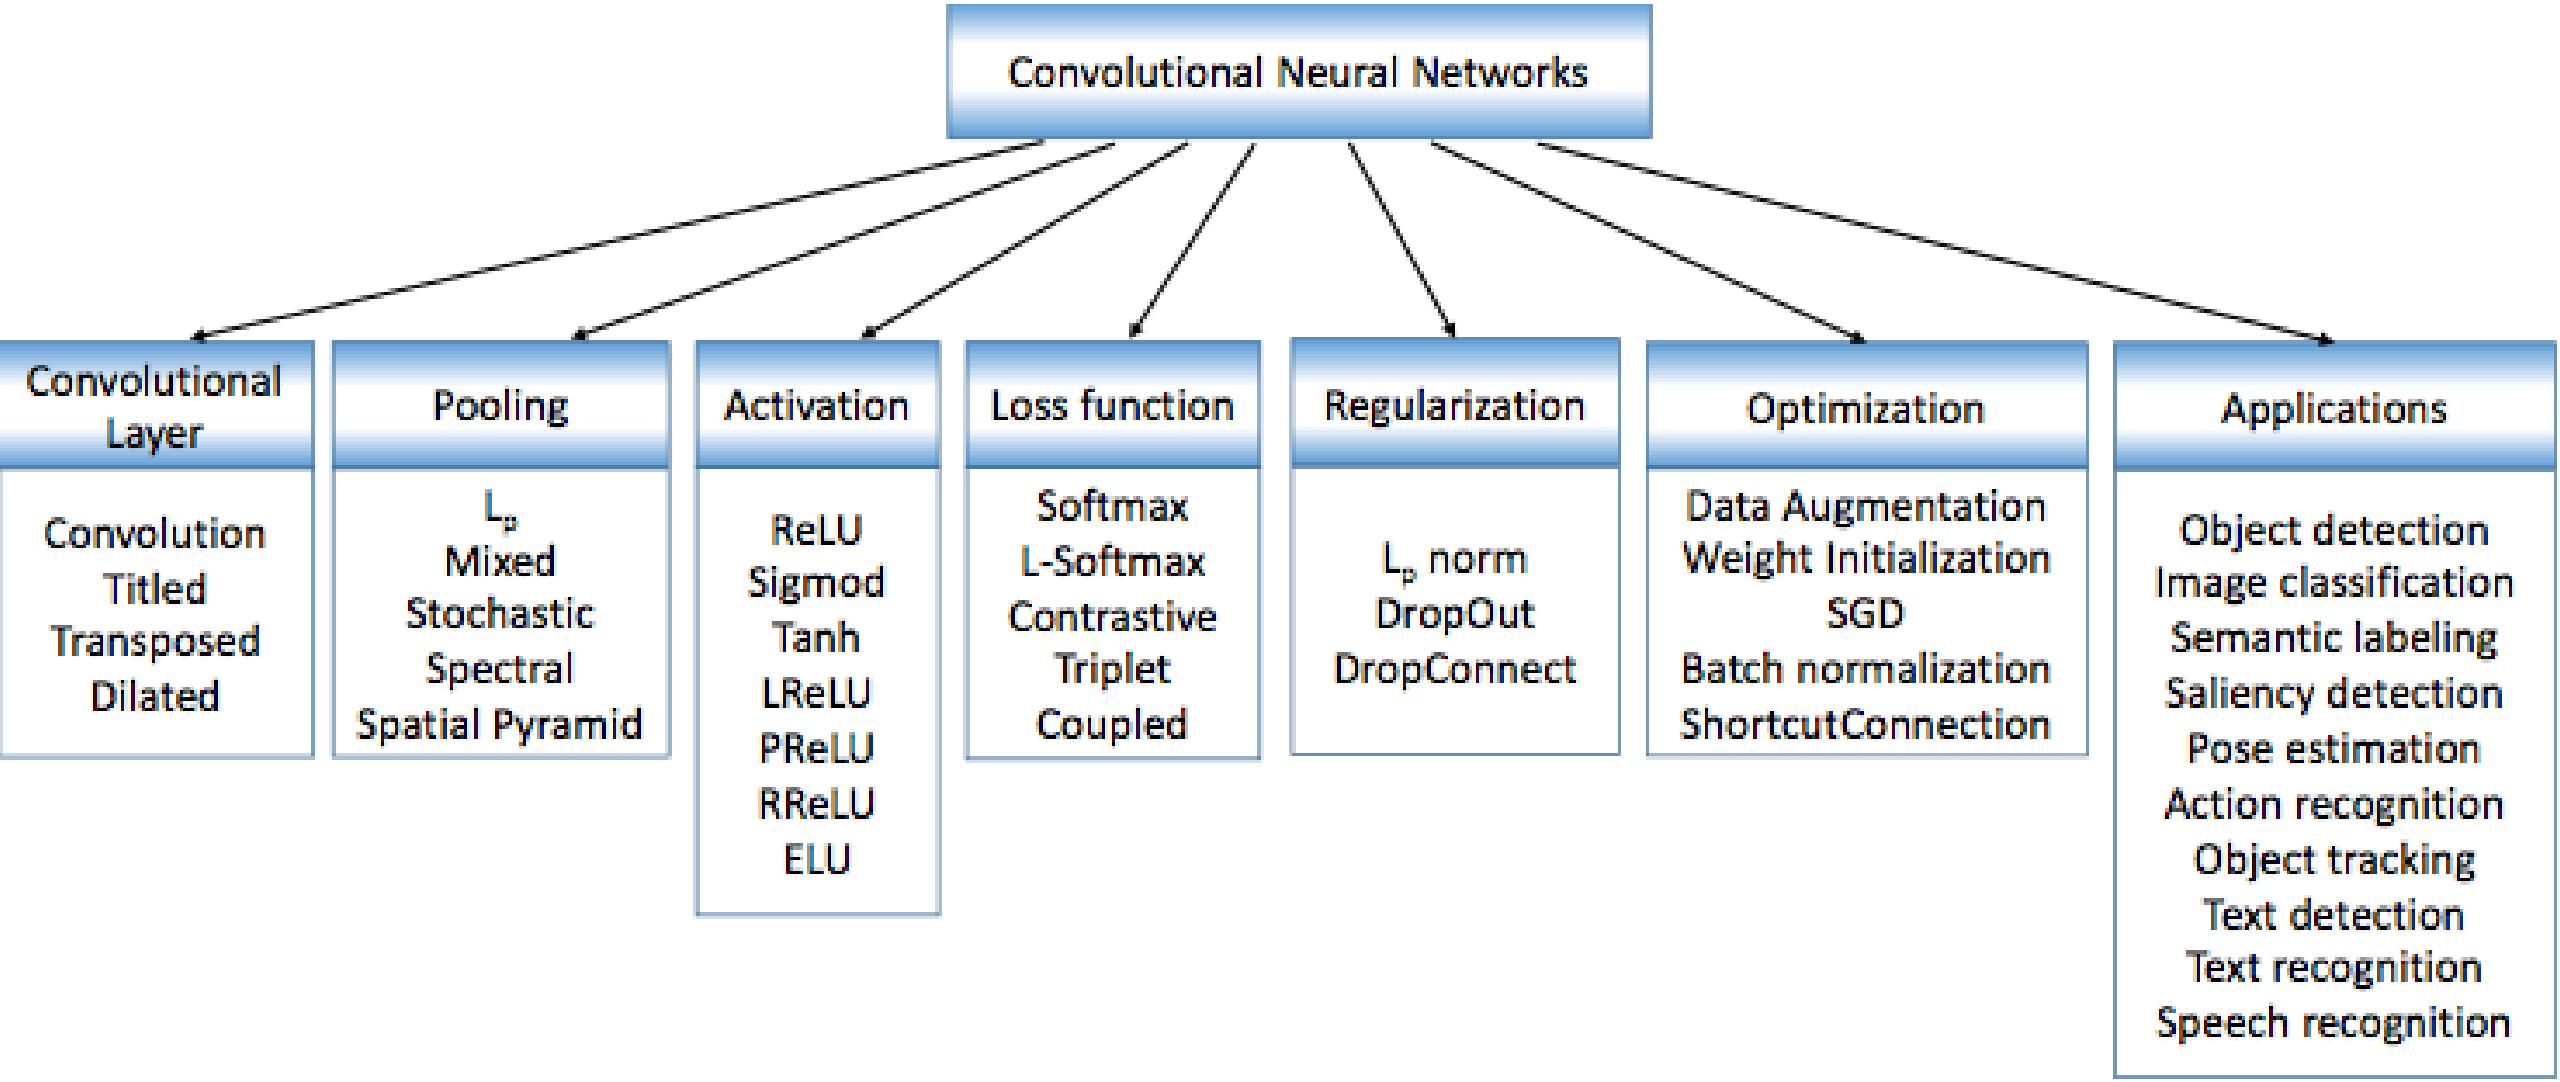
Convolutional Neural Networks
Convolutional Layer
Convolution
Titled
Transposed
Dilated
Pooling
$L_p$
Mixed
Stochastic
Spectral
Spatial Pyramid
Activation
ReLU
Sigmod
Tanh
LReLU
PReLU
RReLU
ELU
Loss function
Softmax
L-Softmax
Contrastive
Triplet
Coupled
Regularization
$L_p$ norm
DropOut
DropConnect
Optimization
Data Augmentation
Weight Initialization
SGD
Batch normalization
ShortcutConnection
Applications
Object detection
Image classification
Semantic labeling
Saliency detection
Pose estimation
Action recognition
Object tracking
Text detection
Text recognition
Speech recognition

## Convolutional Layer

Convolution
Strided
Transposed
Dilated

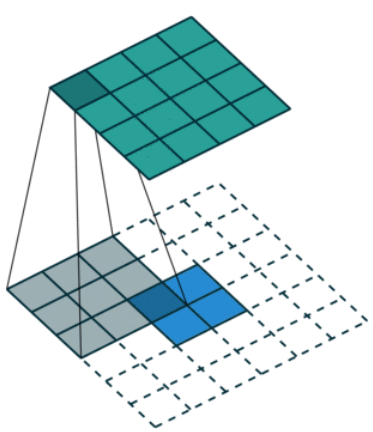

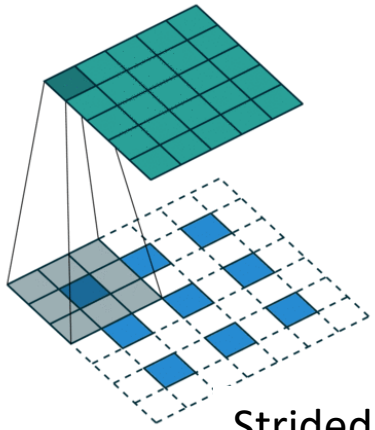

Strided

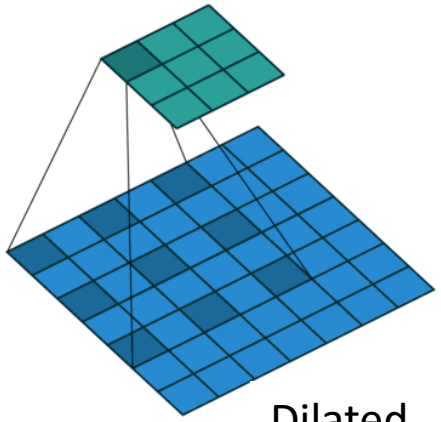

Dilated

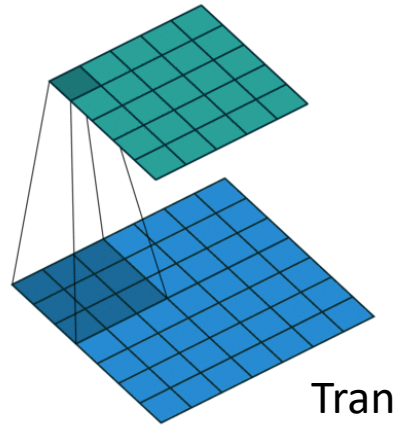

Transposed

## Pooling

$L_p$
Mixed
Stochastic
Spectral
Spatial Pyramid

Lp: average pooling (p=1) or max pooling (p=inf)

Mixed pooling has ability to address the over-fitting

Stochastic pooling is a dropout-inspired pooling

Spectral pooling: performs dimensionality reduction in frequency domain

Spatial pyramid pooling (SPP): generates a fixed-length representation

## Activation

ReLU
Sigmod
Tanh
LReLU
PReLU
RReLU
ELU

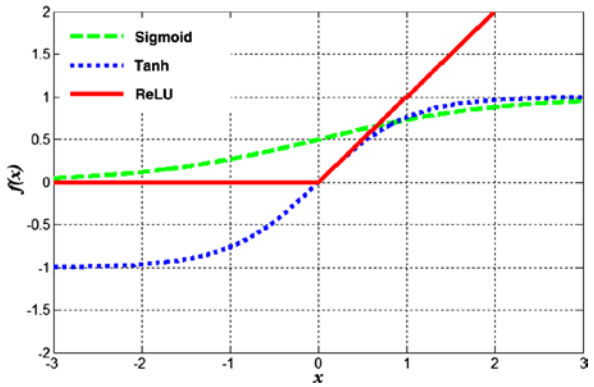


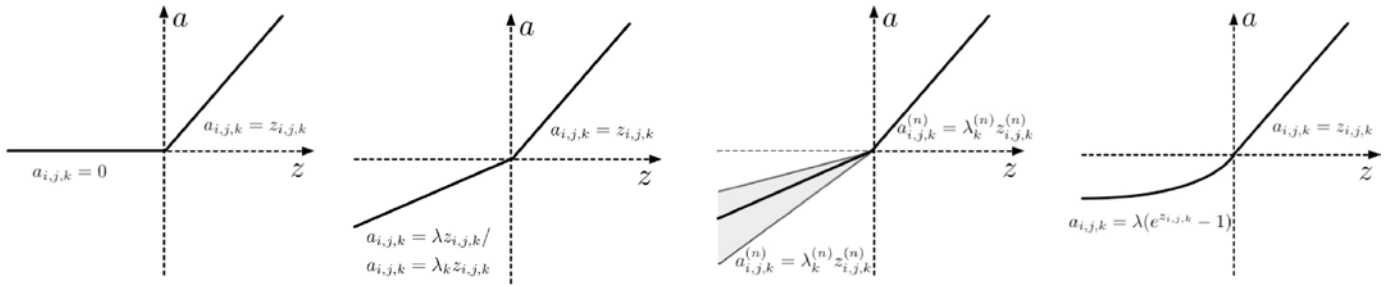


(a) ReLU (b) LReLU/PReLU (c) RReLU (d) ELU

| Flatform | Language | Pros/Cons | By |
|---|---|---|---|
| Tensorflow | C, Python | • Good amount of documentation & has the ability to do partial subgraph computation<br>• The most commonly used deep learning framework<br>• Available on both desktop and mobile and supports languages Python, C++ and R<br>• Adopted by many big companies eBay, Coca Cola, Twitter, Airbus, IBM, and Uber | Google 11 09, 2015 |
| Caffe/Caffe2 | C, Python | • Have a large repository of pre-trained neural network models<br>• priority to expression, speed, and modularity<br>• Caffe2 offers users to use pre-trained models to build demo applications without extra hassle<br>• Need to write C++ / CUDA for new GPU layers | BVLC 2014 Facebook, 2017 |
| Torch / Pytorch | Lua/Python | • Easy to set up & large amount of sample code and tutorials<br>• Can import trained NN models from Caffe's Model Zoo<br>• Difficult to set up in CentOS<br>• Pytorch: competitor to TensorFlow. | Facebook |
| CNTK | Python, C | • Computational Network Toolkit<br>• Easy training and combination of popular model types across servers<br>• RL or GAN can done easily using the toolkit | Microsoft |
| Mxnet | Python, R, C++ and Julia | • Gives the user the ability to code in a variety of programming languages<br>• MXNet supports LSTM along RNN, CNN<br>• Adopted by Microsoft, Intel, and Amazon Web Services. | Apache |
| Other: Chainer, Deeplearning4j, Spark, Hadoop, SystemML<br>RIP: Theano | | | |

Lenet, 1990

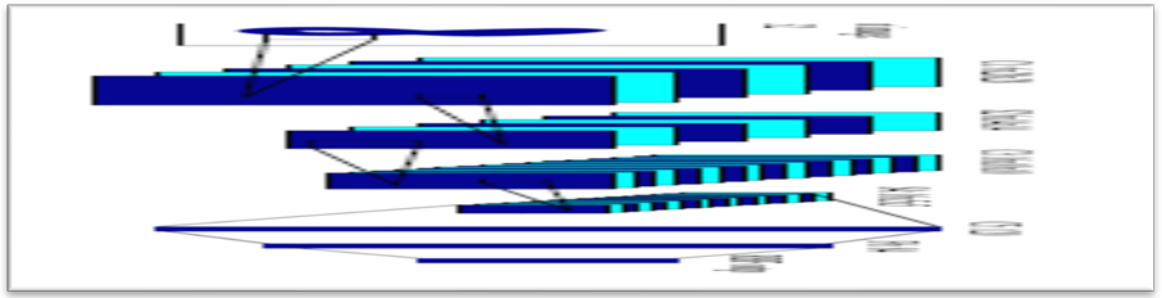

Alexnet,
ILSVRC 12: 15.3% top 5 error

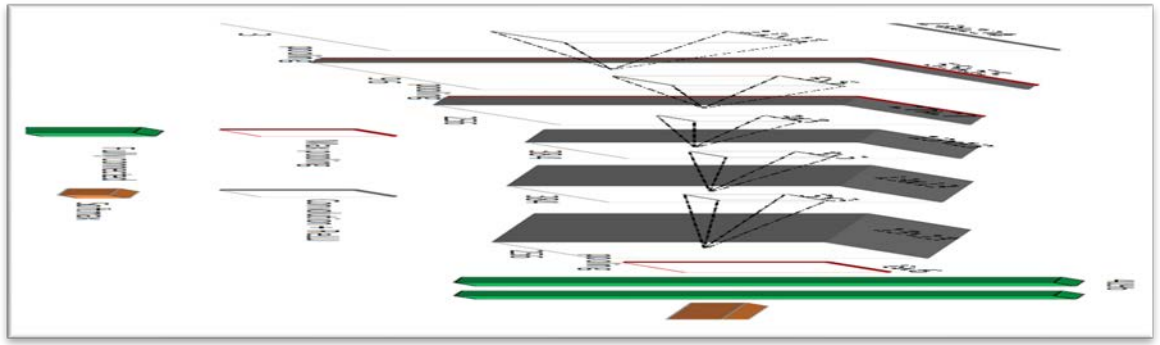

ZFNet
ILSVRC 13: 11.2% top 5 error

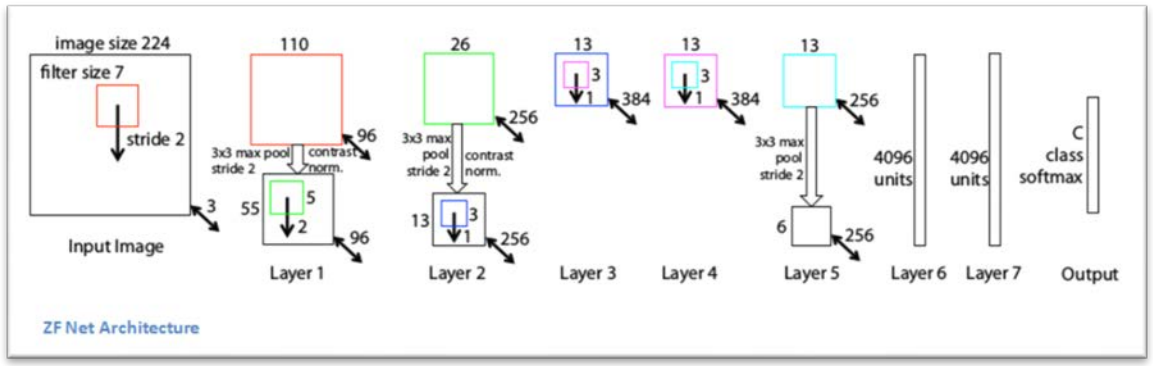


GoogLenet
ILSVRC 14: 6.7% top 5 error

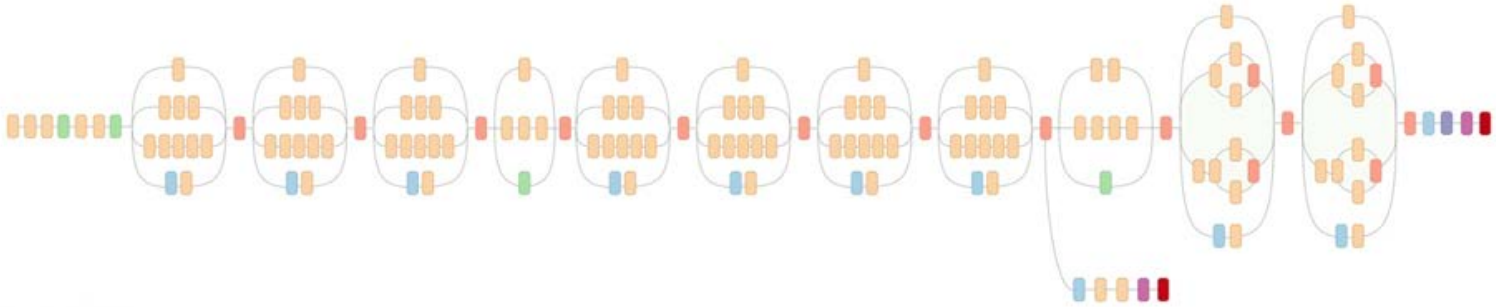

VGG-16
ILSVRC 14: 11.2% top 5 error

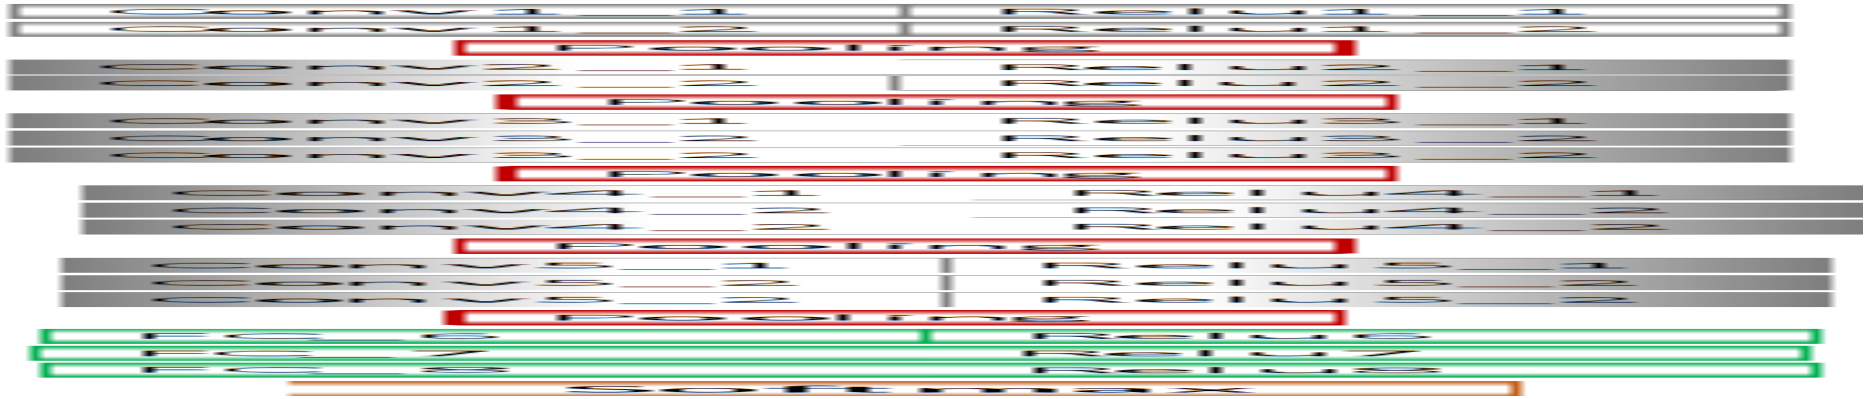

Resnet
ILSVRC 15: 3.6% top 5 error

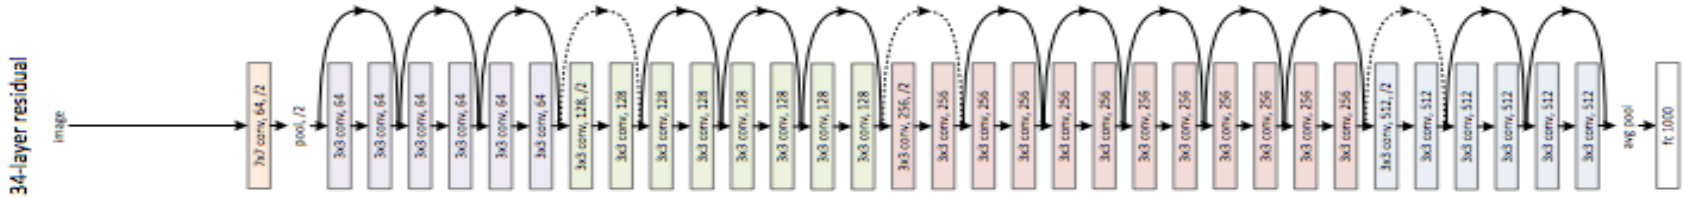


DenseNet

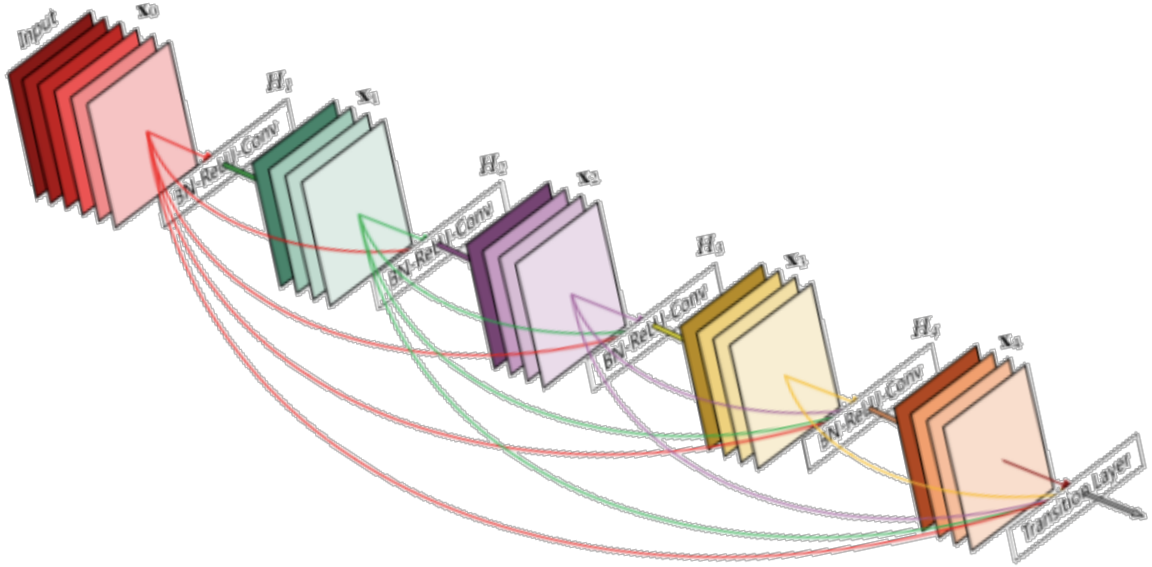


Condensenet

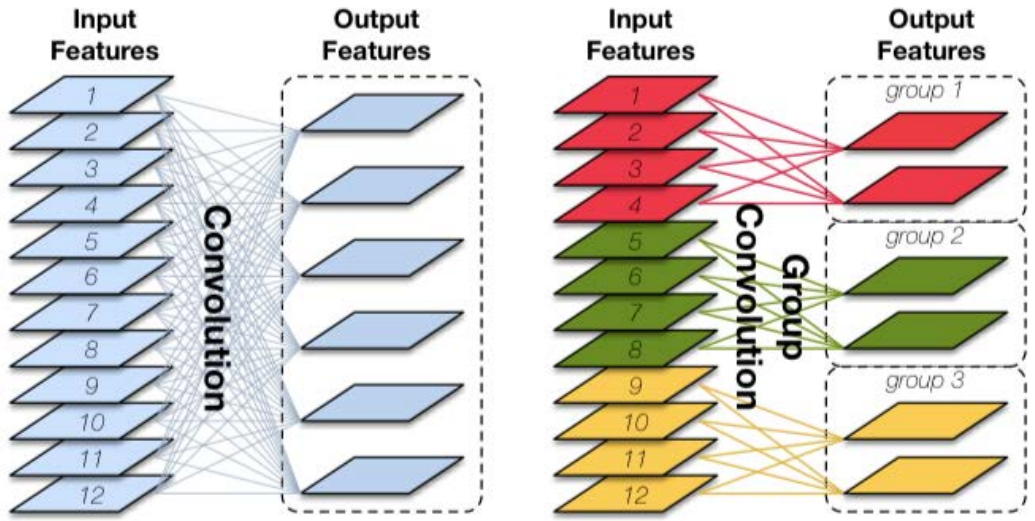

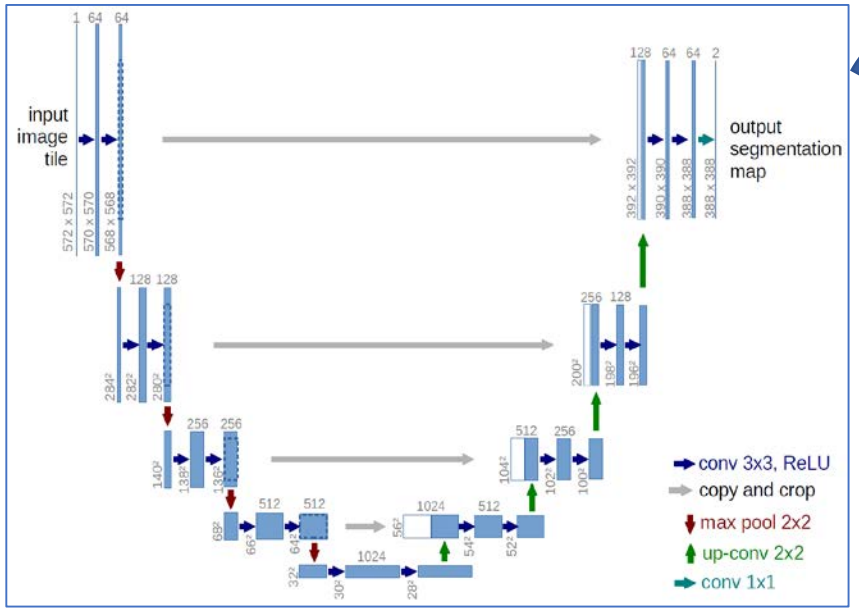


Unet, 2015

3D-Unet, 2016

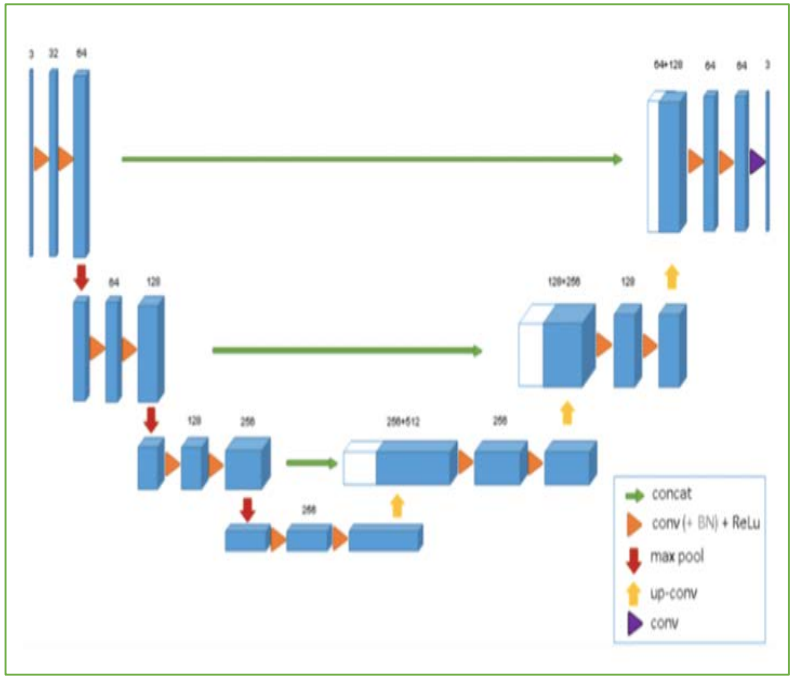


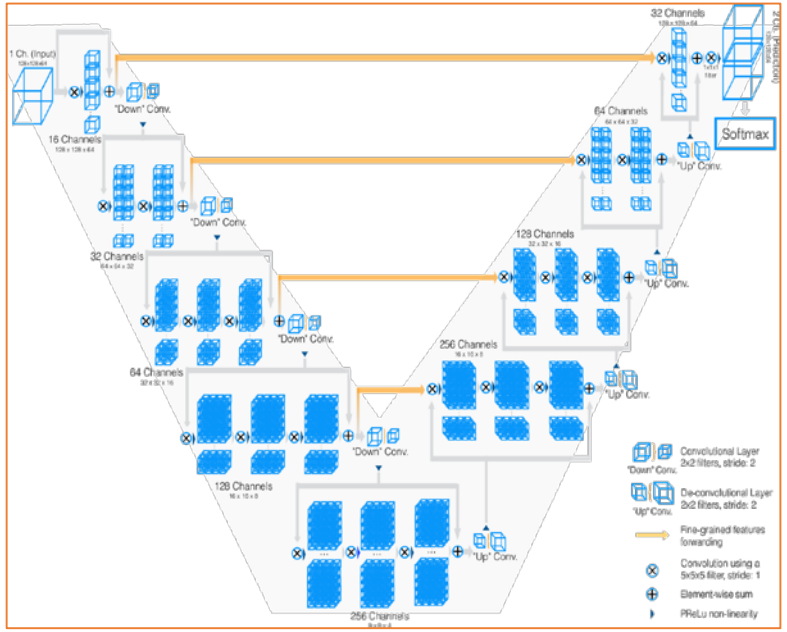


Vnet, 2016

DeepMedic, 2016

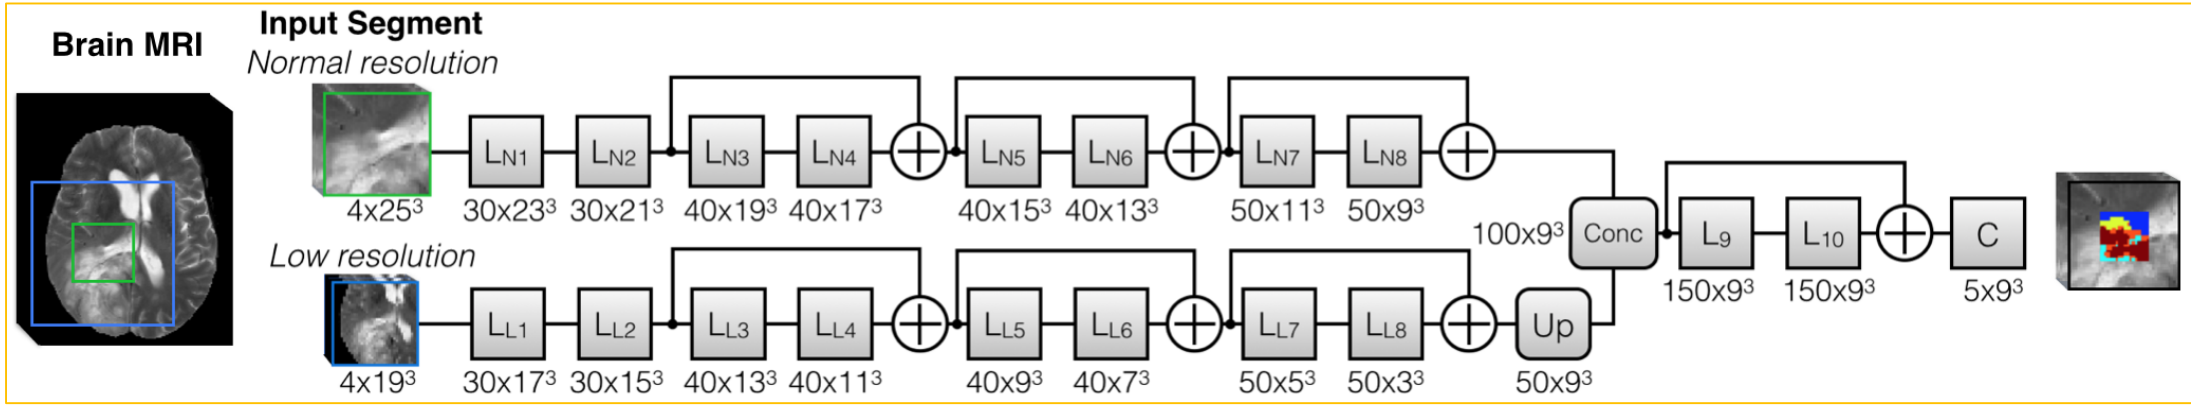

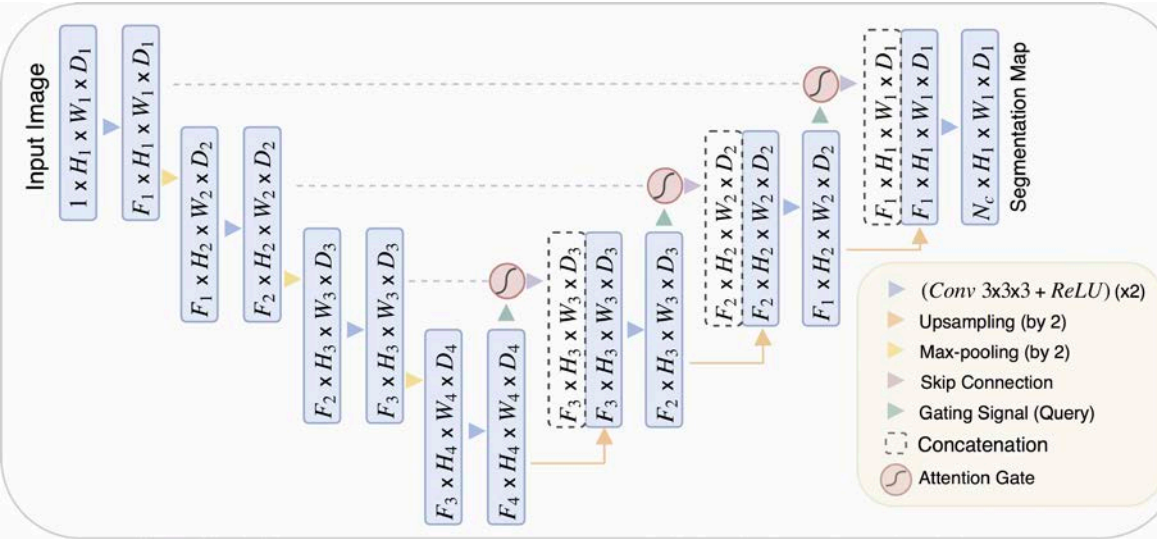


Attention U-net, 2017

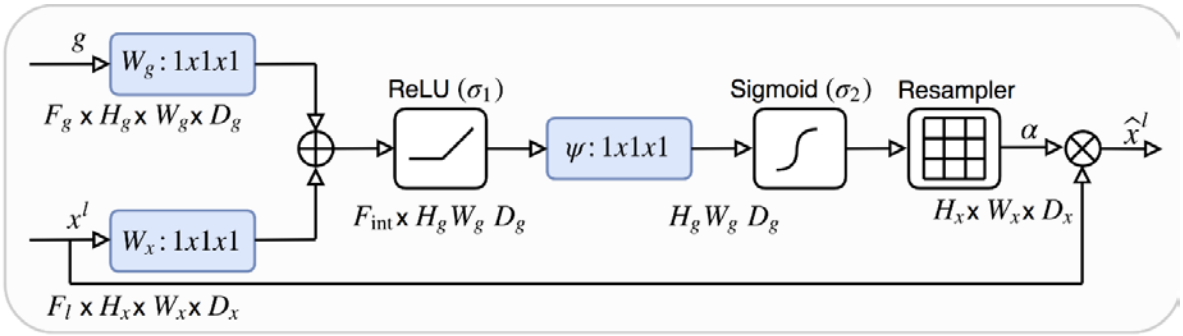


Attention gate

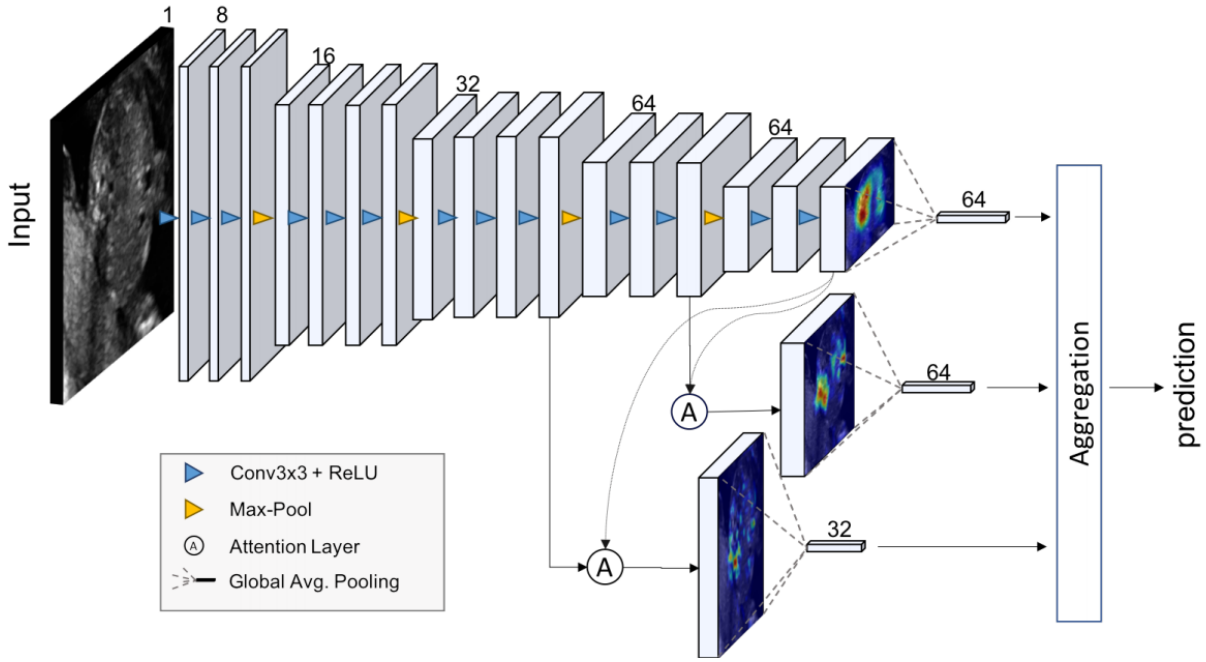


Attention Gated Network, 2017

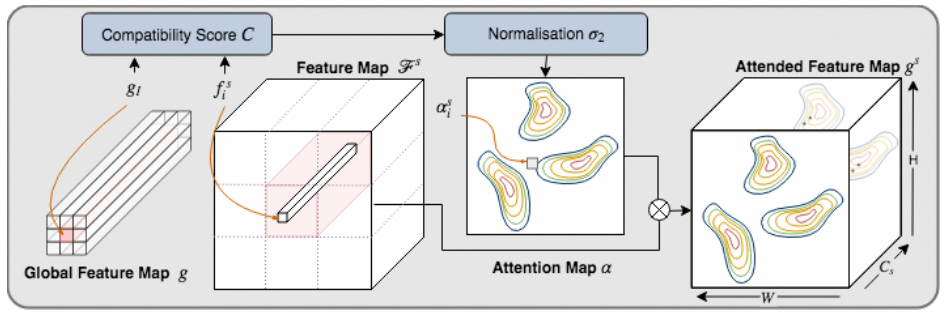


Attention unit

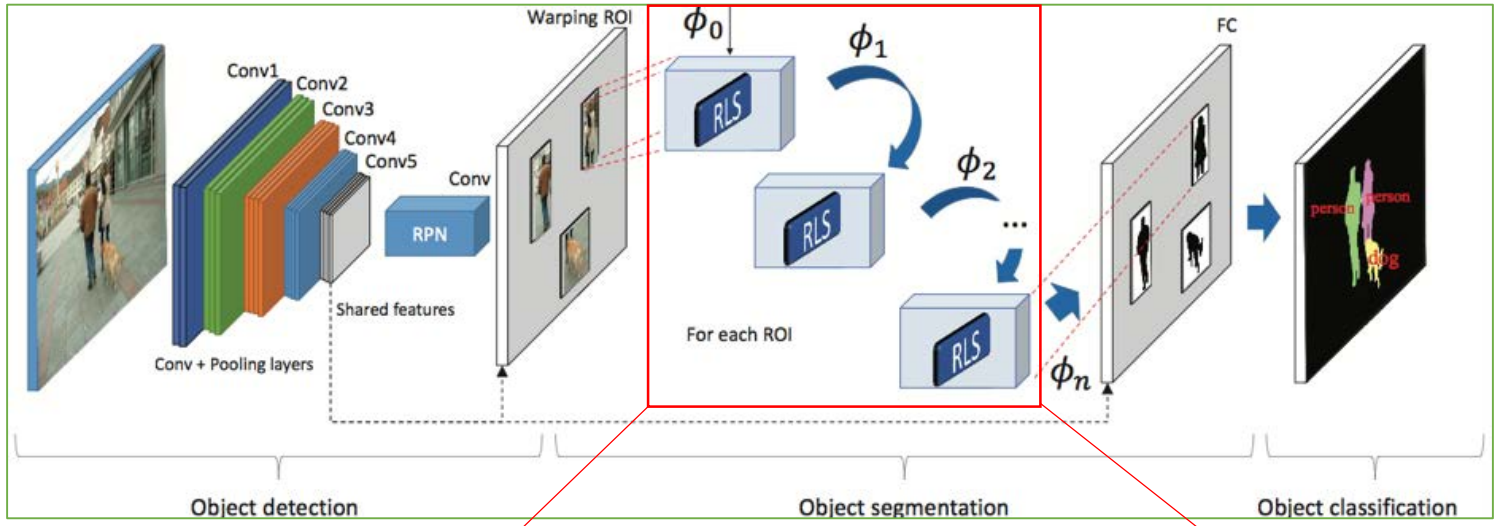


CRLS, 2018

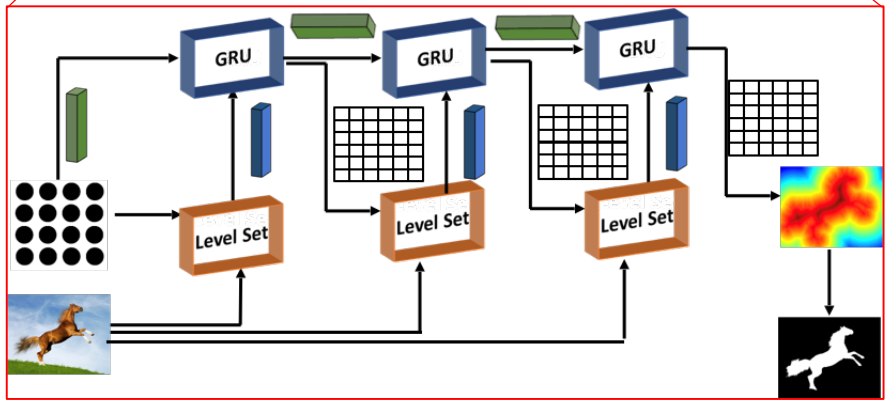


**T.H.N Le**, et al. "Reformulating Level Sets as Deep Recurrent Neural Network Approach to Semantic Segmentation" , **TIP 2018**

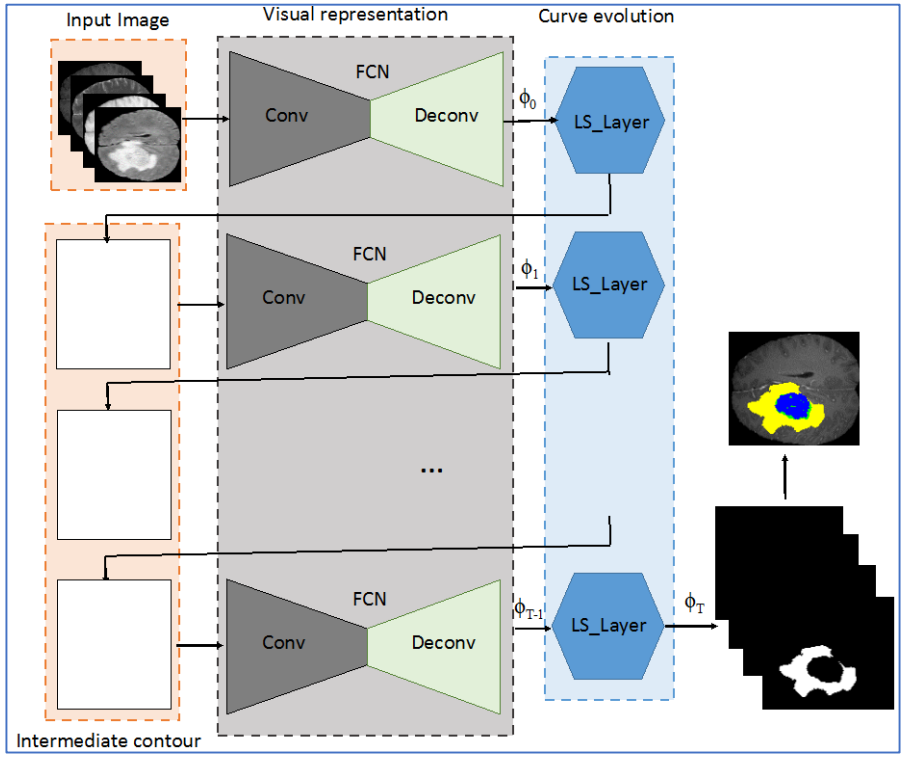


DRLS, 2018

**T.H.N. Le** , et al. “Deep Recurrent Level Set for Segmenting Brain Tumors”, MICCAI 2018

Poster T-144

## *Deep Learning Techniques*

- Neural Networks
- Multilayer Perceptron
- Restricted Boltzmann Machine
- Stacked Auto-Encoders
- Convolutional Auto-Encoder
- Fully Convolutional Networks
- Recurrent Neural Networks
- Generative Adversarial Networks

## *Imaging Modality*

- Ultrasound
- Magnetic Resonance Imaging
- Positron Emission Tomography
- Computed Tompgraphy
- Optical Images

## Application

- Annotation
- Classification
- Detection
- Segmentation
- Registration
- Regression

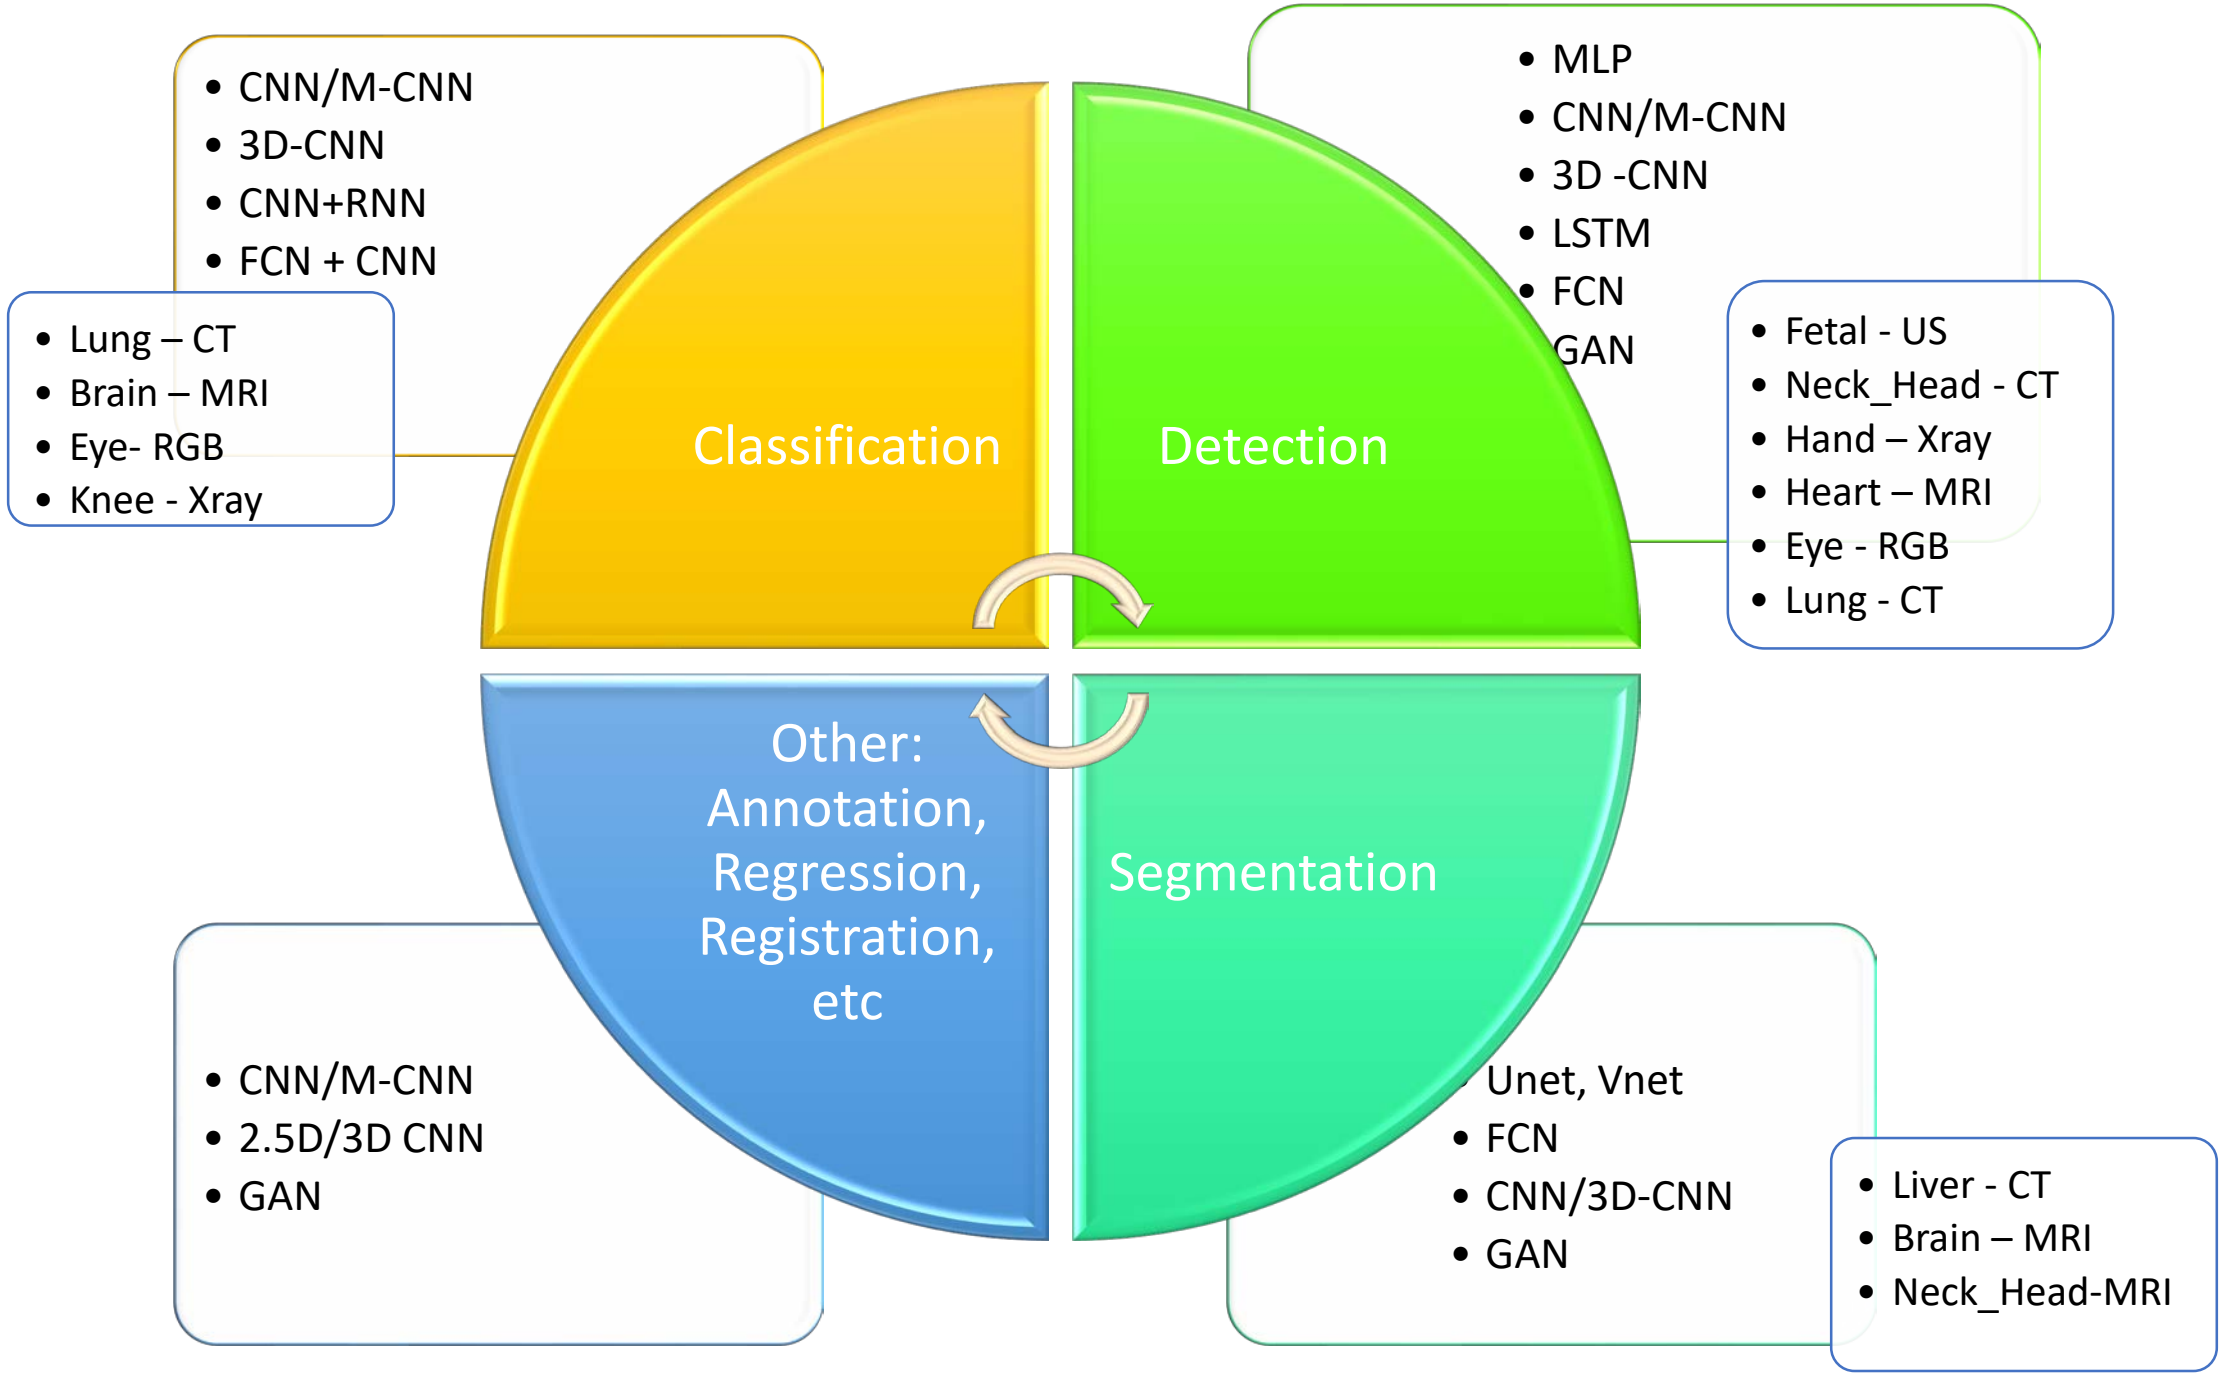
Classification
Detection
Other: Annotation, Regression, Registration, etc
Segmentation
• CNN/M-CNN
• 3D-CNN
• CNN+RNN
• FCN + CNN
• Lung – CT
• Brain – MRI
• Eye- RGB
• Knee - Xray
• MLP
• CNN/M-CNN
• 3D -CNN
• LSTM
• FCN
• GAN
• Fetal - US
• Neck_Head - CT
• Hand – Xray
• Heart – MRI
• Eye - RGB
• Lung - CT
• CNN/M-CNN
• 2.5D/3D CNN
• GAN
• Unet, Vnet
• FCN
• CNN/3D-CNN
• GAN
• Liver - CT
• Brain – MRI
• Neck_Head-MRI

# Book


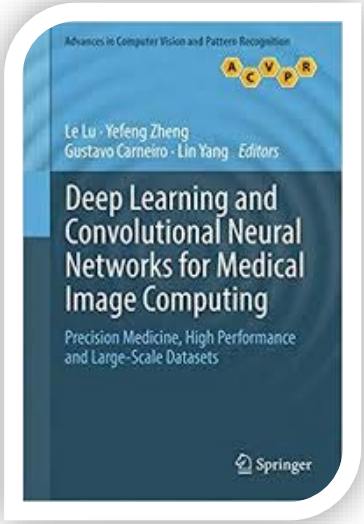


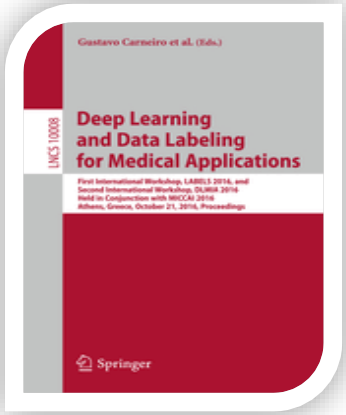


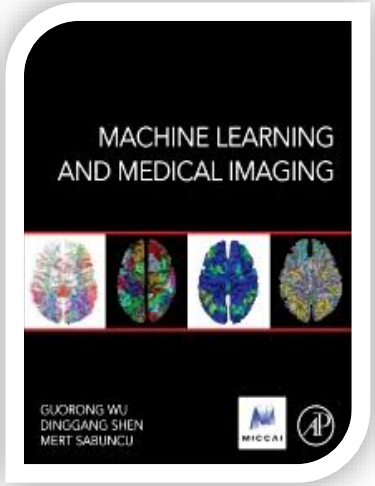


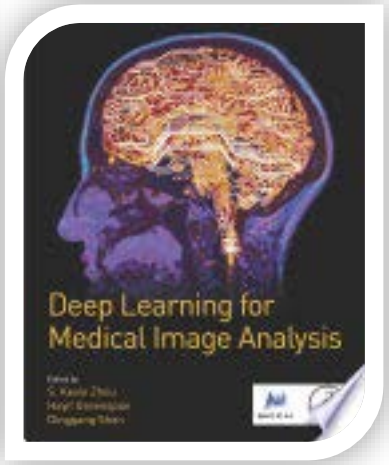


# Article


**Dermatologist-level classification of skin cancer**

An artificial intelligence trained to classify images of skin lesions as benign lesions or malignant skin cancers achieves the accuracy of board-certified dermatologists.

In this work, we pretrain a deep neural network at general object recognition, then fine-tune it on a dataset of ~130,000 skin lesion images comprised of over 2000 diseases.


Brain tumor segmentation with Deep Neural Networks

Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation


Guest Editorial
Deep Learning in Medical Imaging: Overview and Future Promise of an Exciting New Technique


Deep Learning in Medical Image Analysis


Contest – Challenge - Dataset

# Medical database

Acquiring, annotating and distributing medical image data sets are costly

Requires high levels of expertise from clinicians with limited time

Due to privacy concerns, sharing data sets between institutions, let alone internationally, is logistically and legally challenging

Typical data sets remain small

Capturing high-resolution data in multiple dimensions
Data dimensionality: 2D – 5D

Stored in different formats than in many computer vision tasks, e.g DICOM, NIfTI, Analyze

DeepIGeoS: semi-automated annotation
GIFT-Cloud: data sharing

# Grand Challenges in Biomedical Image Analysis

https://grand-challenge.org/challenges/

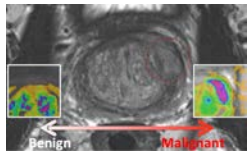


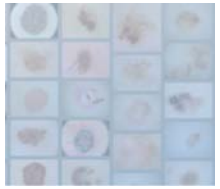

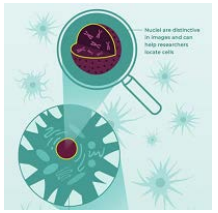



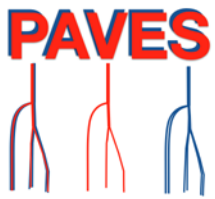


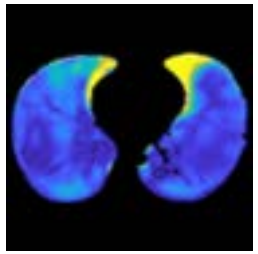

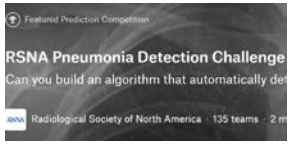


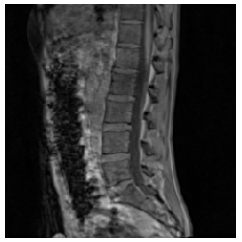

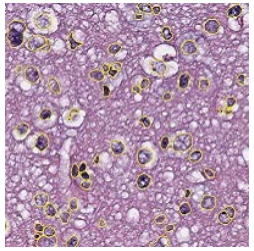

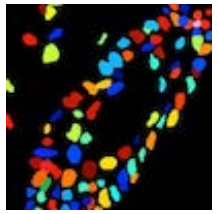


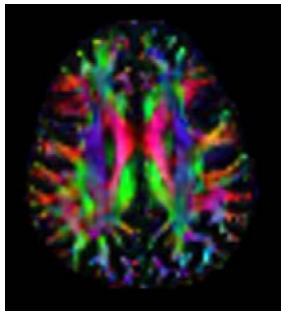

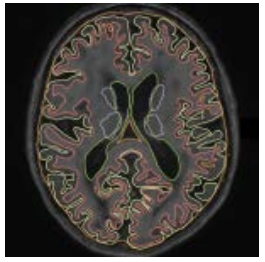





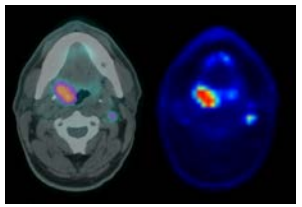

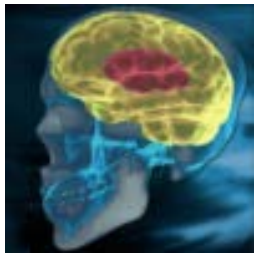

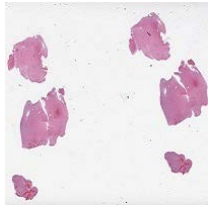


**Liver Tumor Segmentation Challenge**

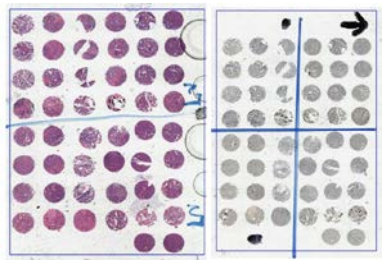

**Tissue Microarray Analysis in Thyroid Cancer Diagnosis**


**Skin Lesion Analysis Towards Melanoma Detection**

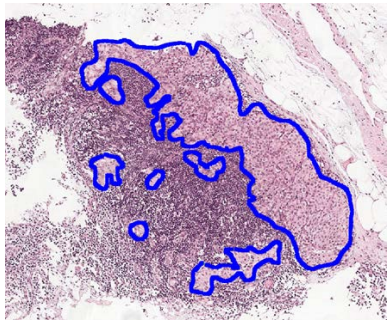

**CAMELYON17**



**3-D Validation of Tractography with Experimental MRI (3D VoTEM)**


**Diabetic Retinopathy – Segmentation and Grading Challenge**

Lung Nodule Malignancy Prediction, Based on Sequential CT Scans

DREAM CHALLENGES

*powered by Sage Bionetworks*

http://dreamchallenges.org

**Tumor Deconvolution DREAM Challenge**

**DREAM Single Cell Transcriptomics Challenge**

**IDG-DREAM Drug-Kinase Binding Prediction Challenge**

CDC Centers for Disease Control and Prevention

CDC 24/7: Saving Lives. Protecting People.™

https://data.cdc.gov/browse


https://stanfordmlgroup.github.io/competitions/mura/

# Toolkits


https://github.com/DLTK/DLTK

Neural networks toolkit written in python, on top of TensorFlow

Provides easy to use baselines for deep learning on medical images.

It enables fast prototyping and is simply installed via pypi


https://github.com/MITK/MITK

Free open-source software system for development of interactive medical image processing software

Combines the Insight Toolkit (ITK) and the Visualization Toolkit (VTK) with an application framework.

MITK is a cross-platform C++ toolkit and officially supports: Windows, macOS, Linux

# Deep Neural Network

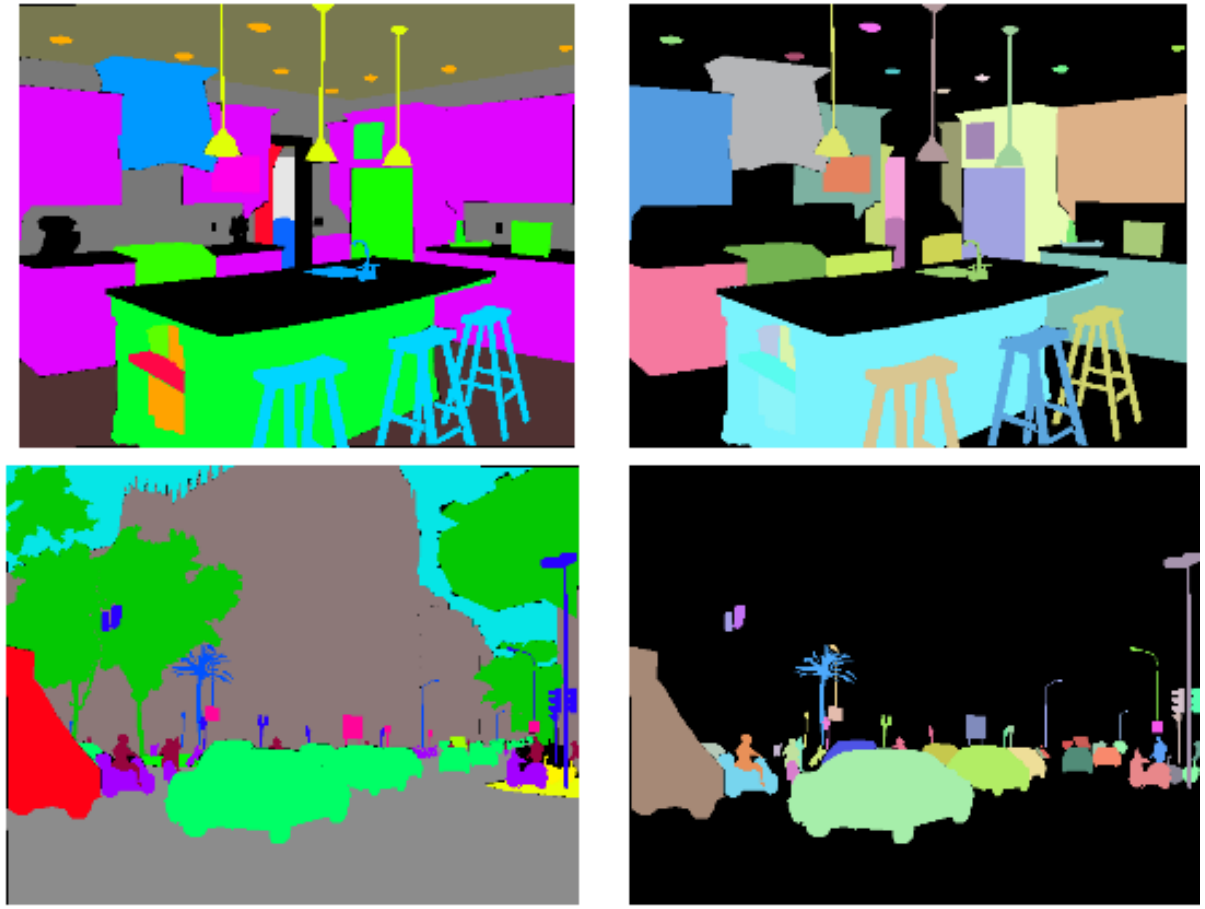

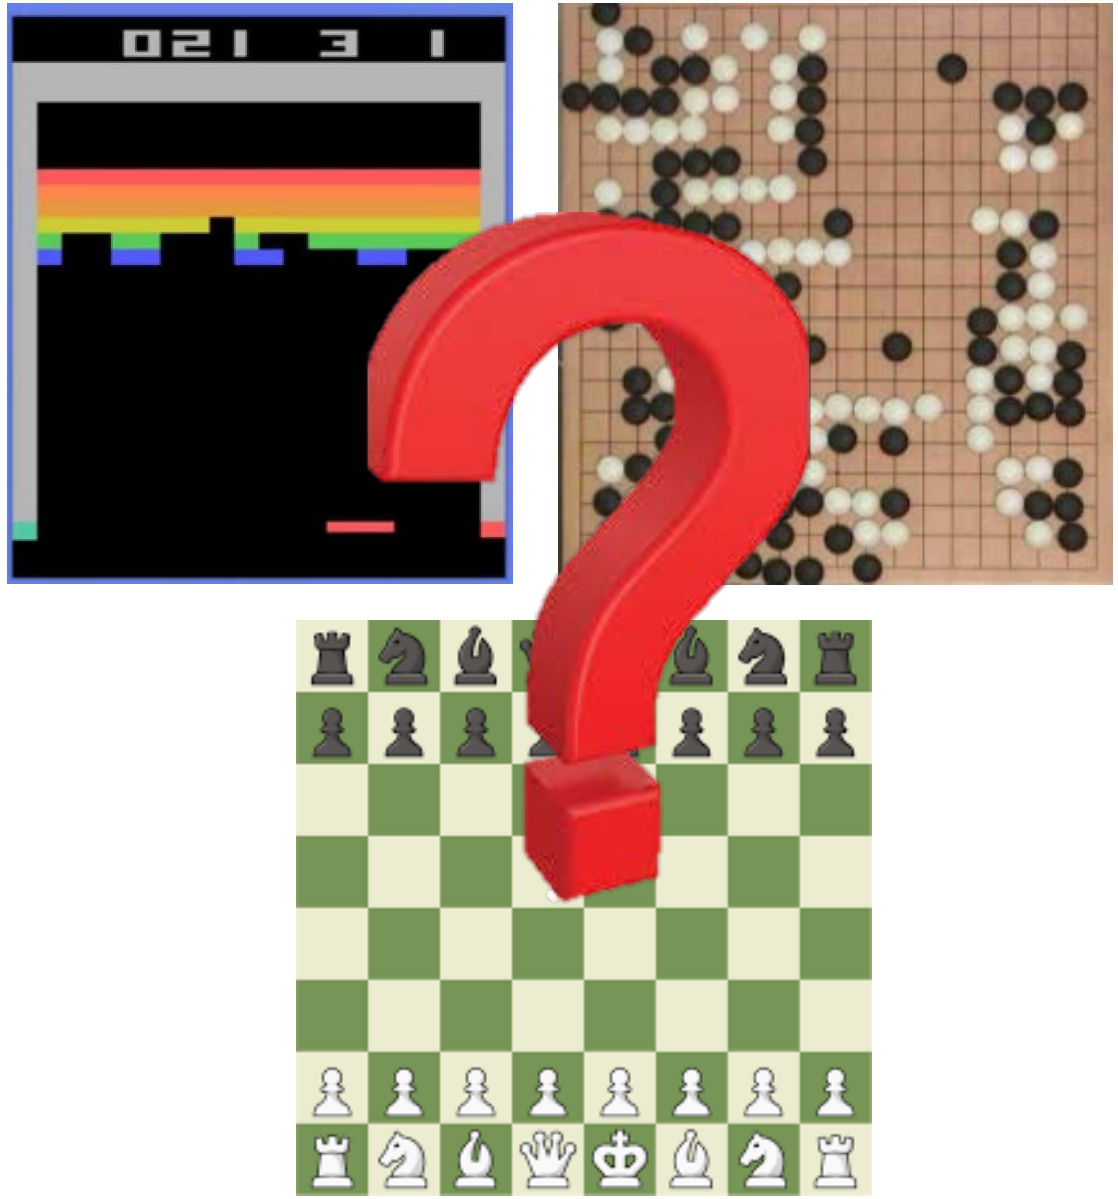


THANK YOU!
QUESTIONS?

# Backup

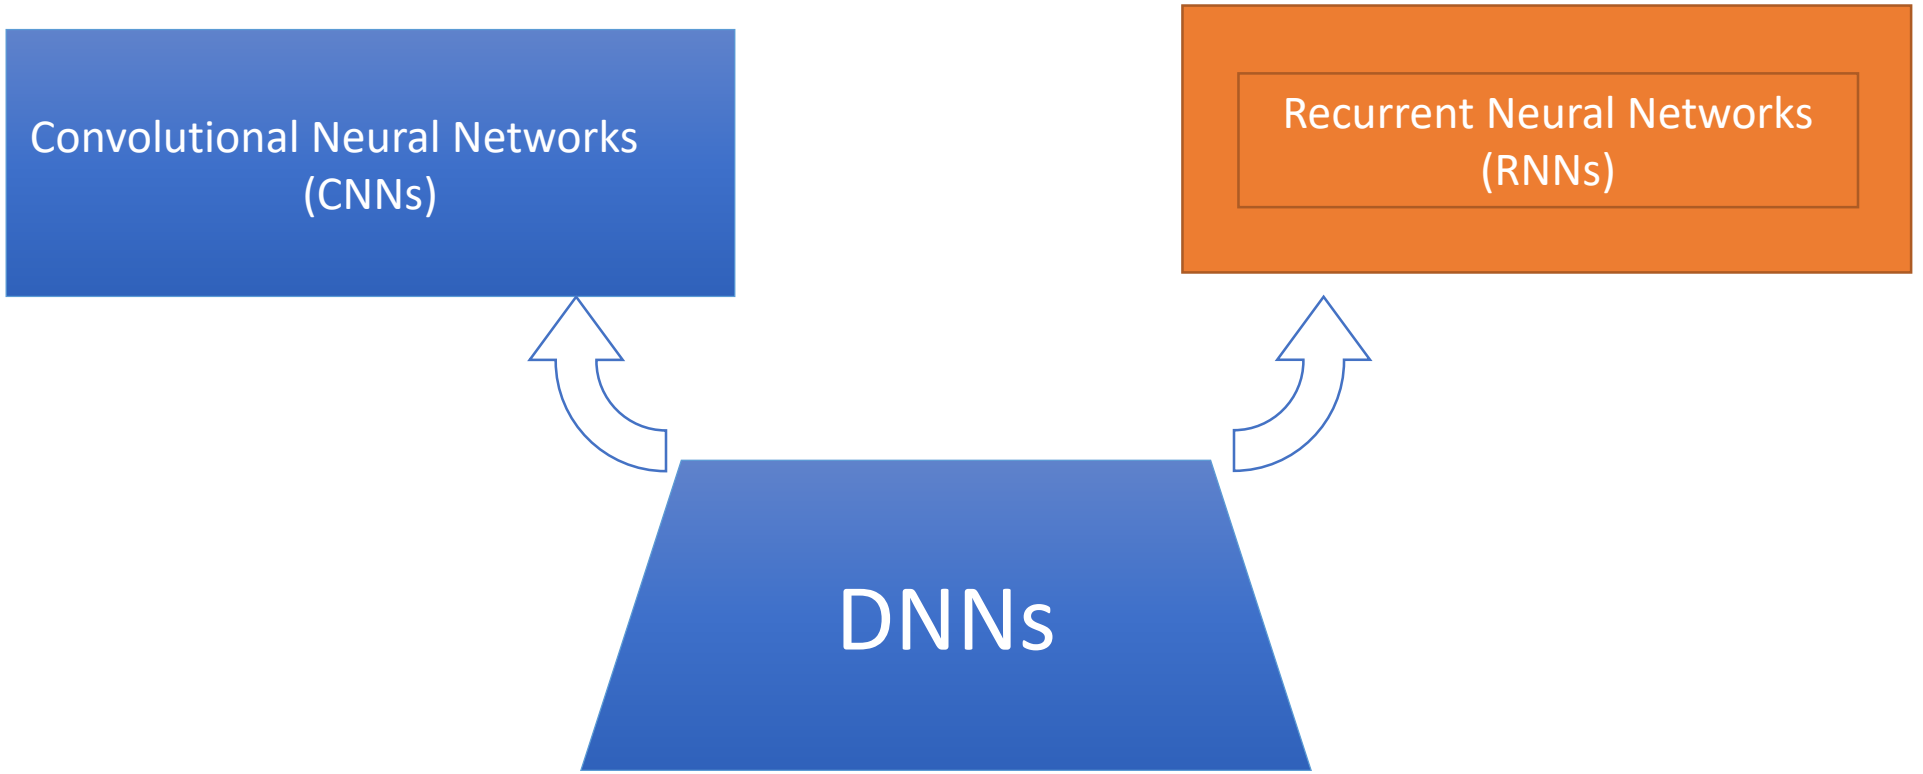
Convolutional Neural Networks
(CNNs)
Recurrent Neural Networks
(RNNs)
DNNs

$$\mathbf{h}_t = tanh(\mathbf{U}\mathbf{x}_t + \mathbf{W}\mathbf{h}_{t-1})$$

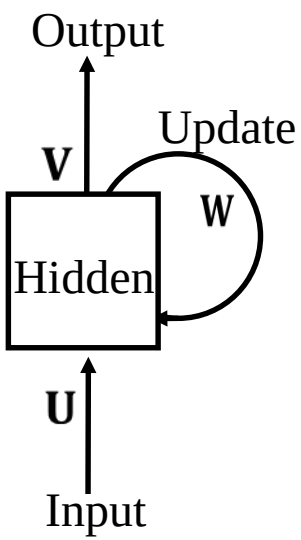


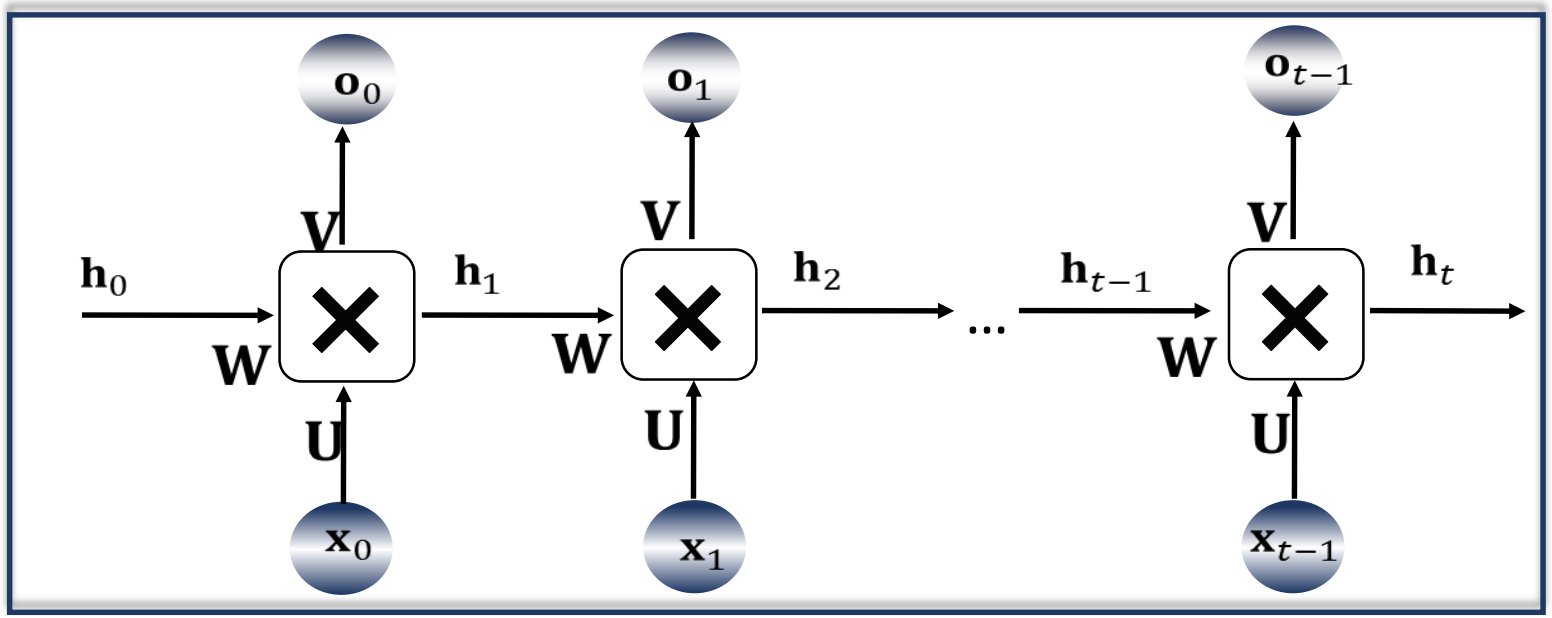

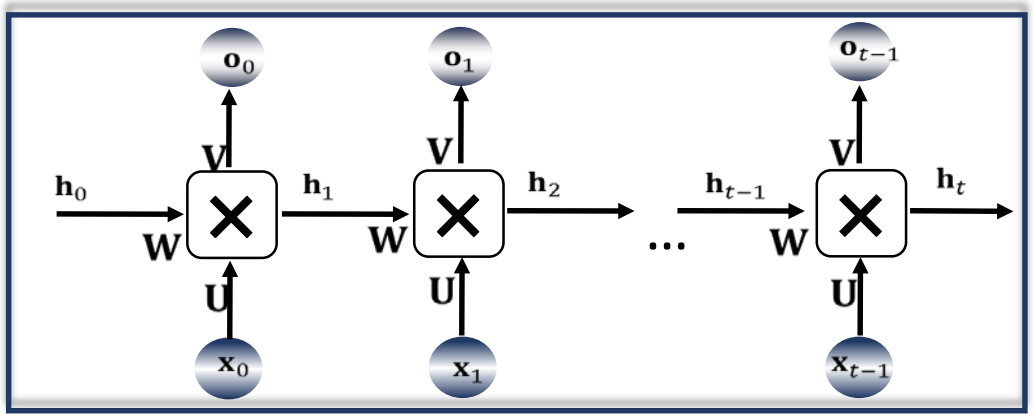


# Sentiment Analysis

Score: 5

FC

RNN cell

RNN cell

RNN cell

RNN cell

RNN cell

$h_0$ $h_1$ $h_2$ $h_3$ $h_4$ $h_5$

$x_0$ $x_1$ $x_2$ $x_3$ $x_4$

The Movie Is Very Intersting

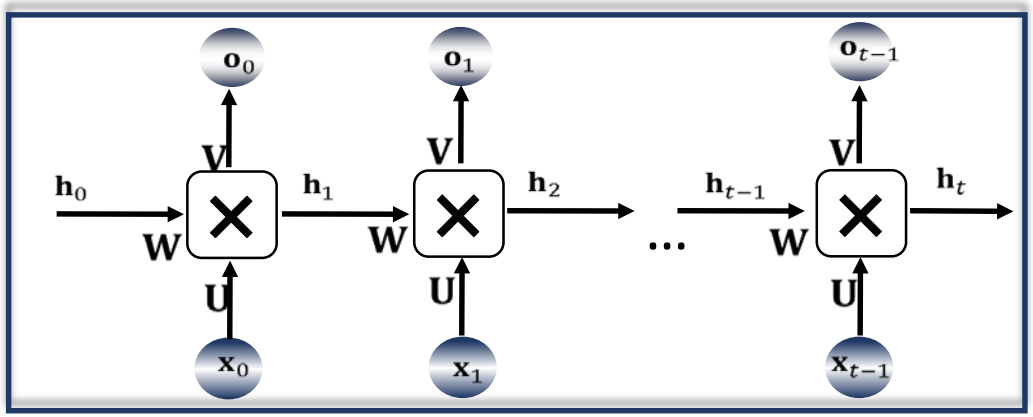


# Machine Translation

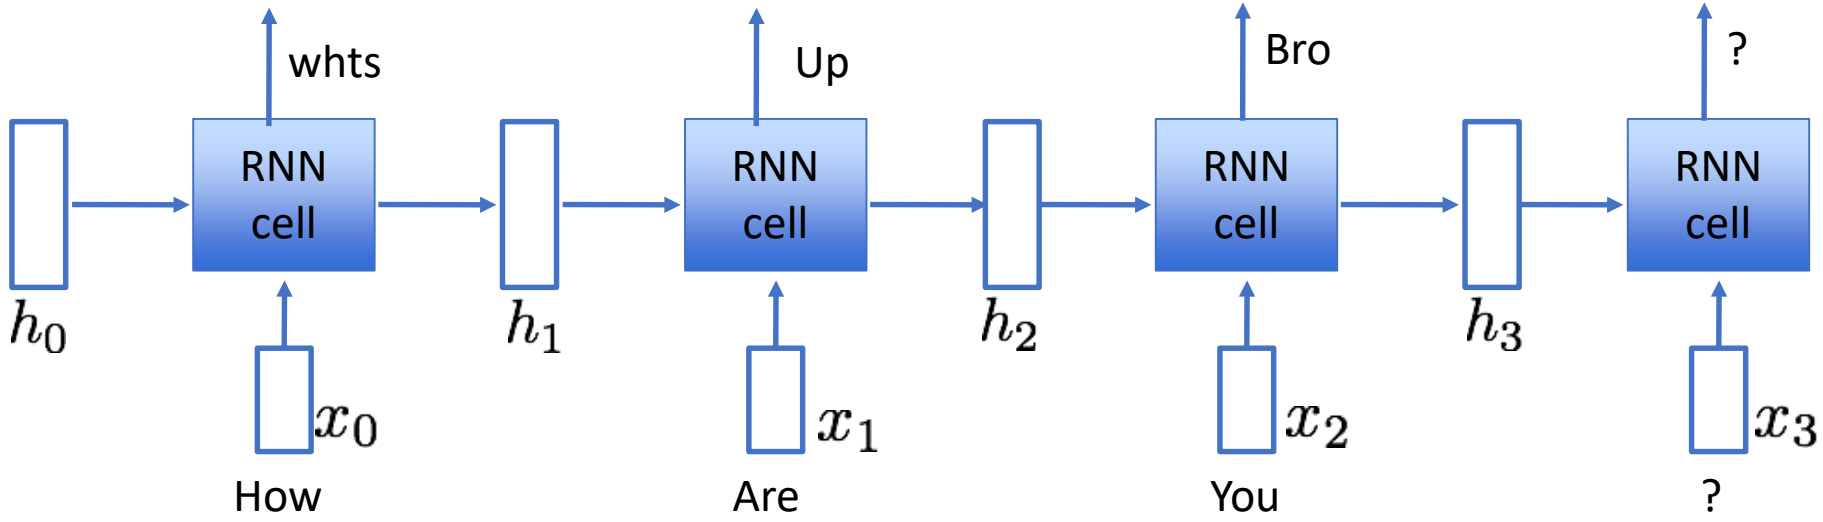

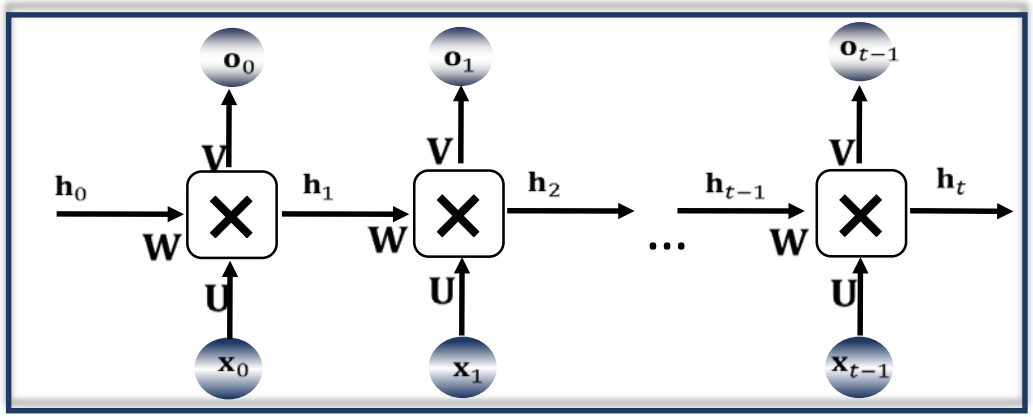


# Machine Translation

Decoder

Encoder

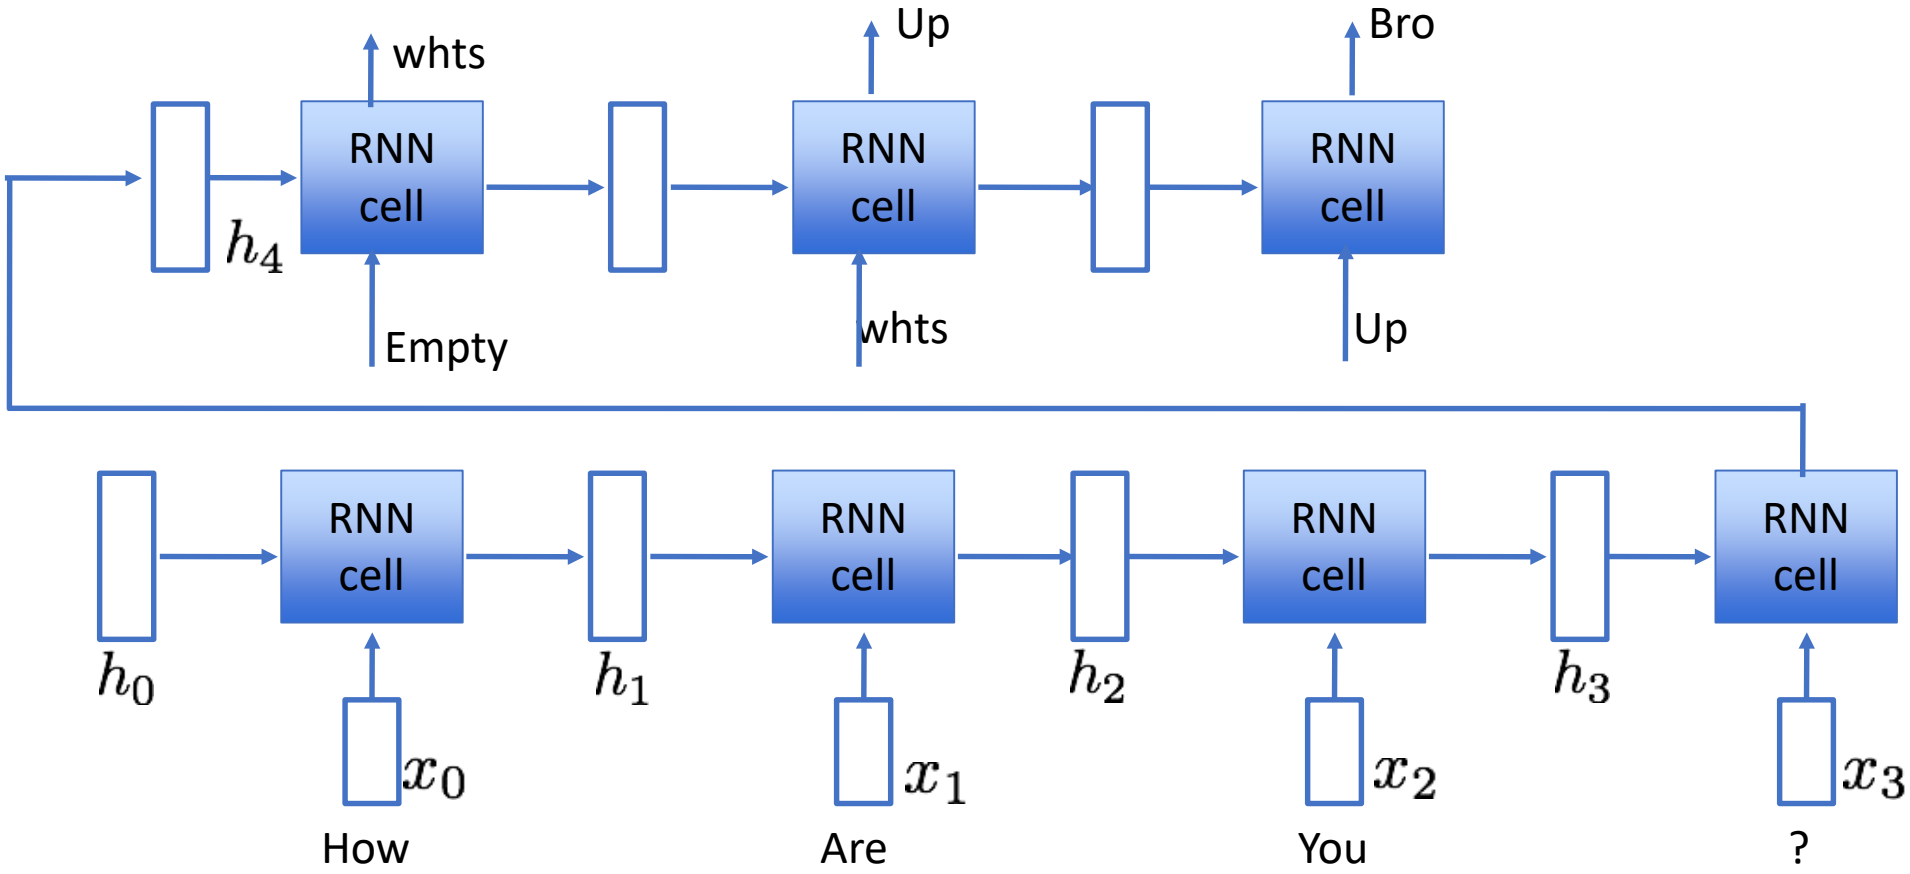

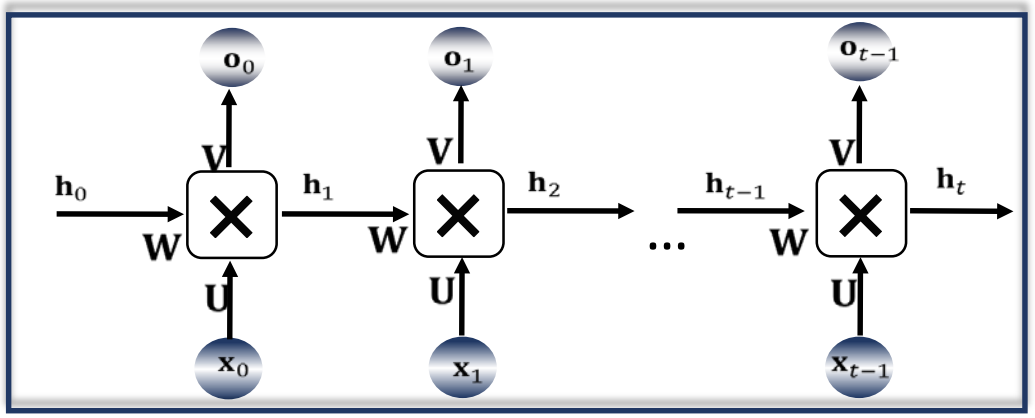


# Vanishing/Exploding Gradient

## Backpropagation

$$w = w + \triangle w$$

$$\triangle w = \frac{de}{dw}$$

$$e = (GT - predicted)^2$$

**Vanishing**

No update for weight

if $\frac{de}{dw} <<<<<< 1$

$\triangle w <<<<<<< 1$

**Exploding**

if $\frac{de}{dw} >>>>> 1$

$\triangle w >>>>>> 1$

RNN

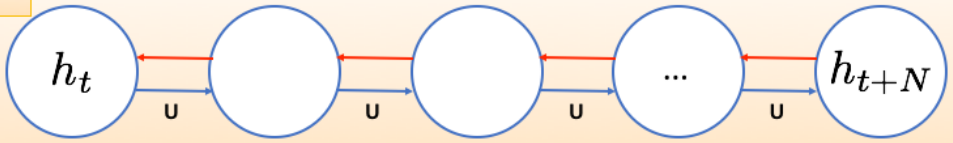


The error must back propagate through all the intermediate nodes

Create shortcut

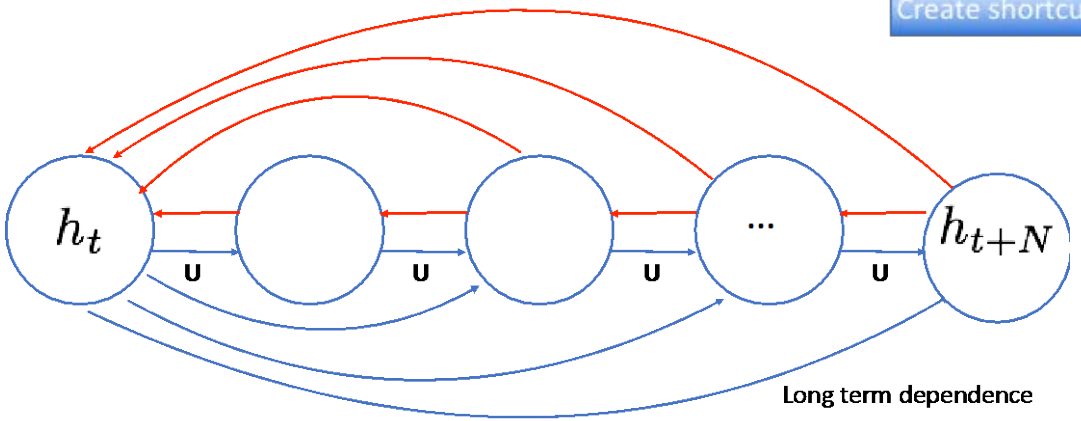


Long term dependence

Adaptive shortcut connection to decide how much to pay attention to the past as well as to the present

# LSTM

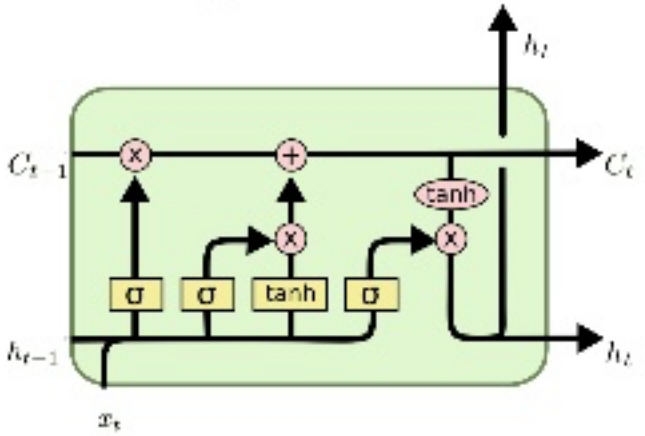


# GRU

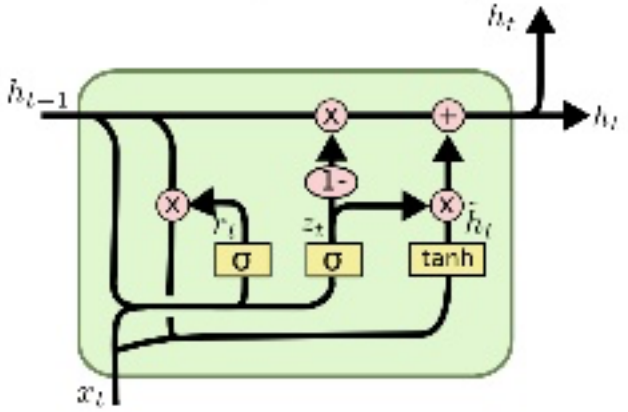


Deep Learning in Medical Imaging