



DEEP
LEARNING
INSTITUTE



PRACE Workshop: Deep Learning and GPU programming workshop

7 – 10 September 2020



VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER



VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER



Preliminaries on Convolutional and Recurrent Neural Networks

07.09.2020 | PD Dr. Juan J. Durillo

Agenda



DEEP
LEARNING
INSTITUTE



- Part 1
 - Introduction
 - Introduction to (Deep) Neural Networks for Machine Learning
 - Computer Vision as working example
 - Introduction to Convolutional Neural Networks
 - Deep Neural Network Architecture
- Part 2
 - More than Images
 - Representing Languages
 - Recurrent Neural Networks



VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

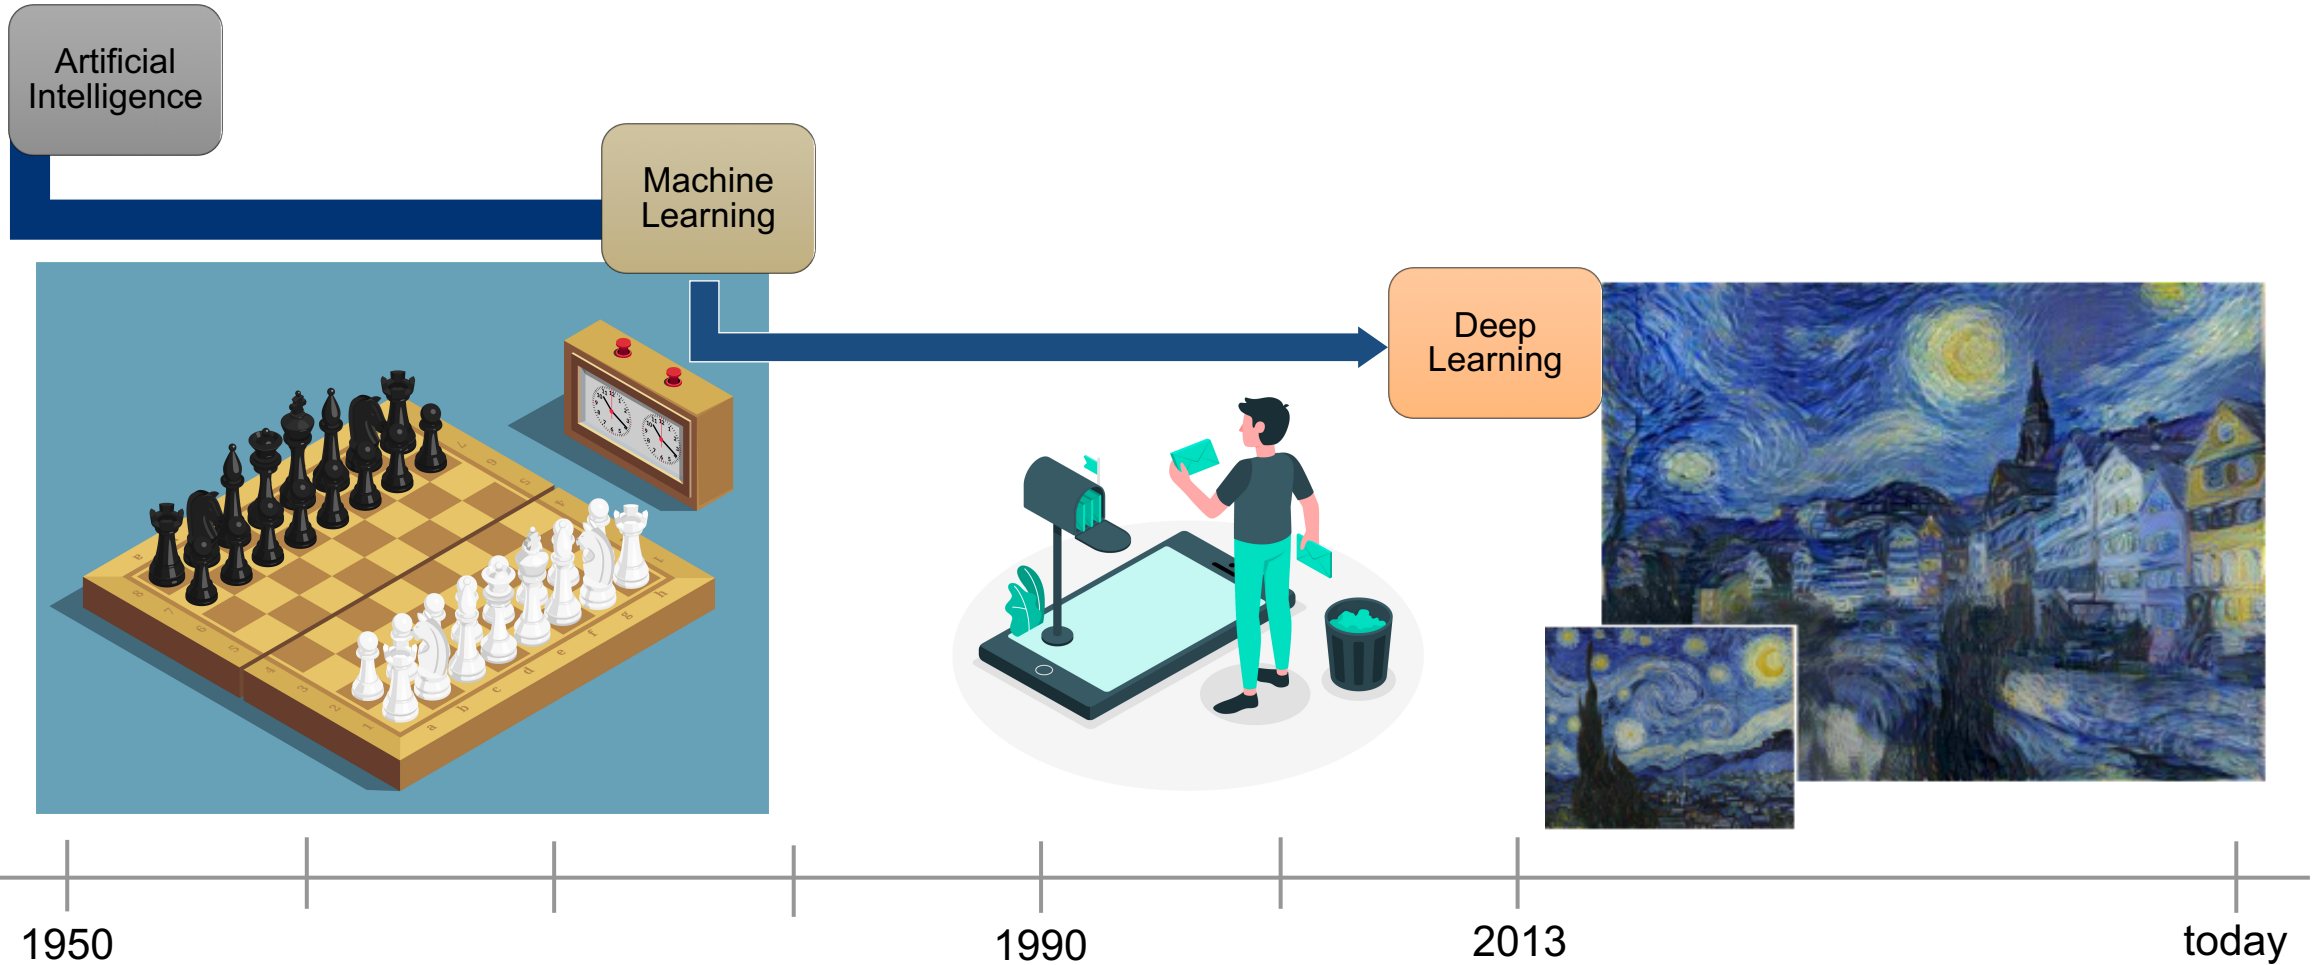


Part 1

Artificial + Intelligence



DEEP
LEARNING
INSTITUTE




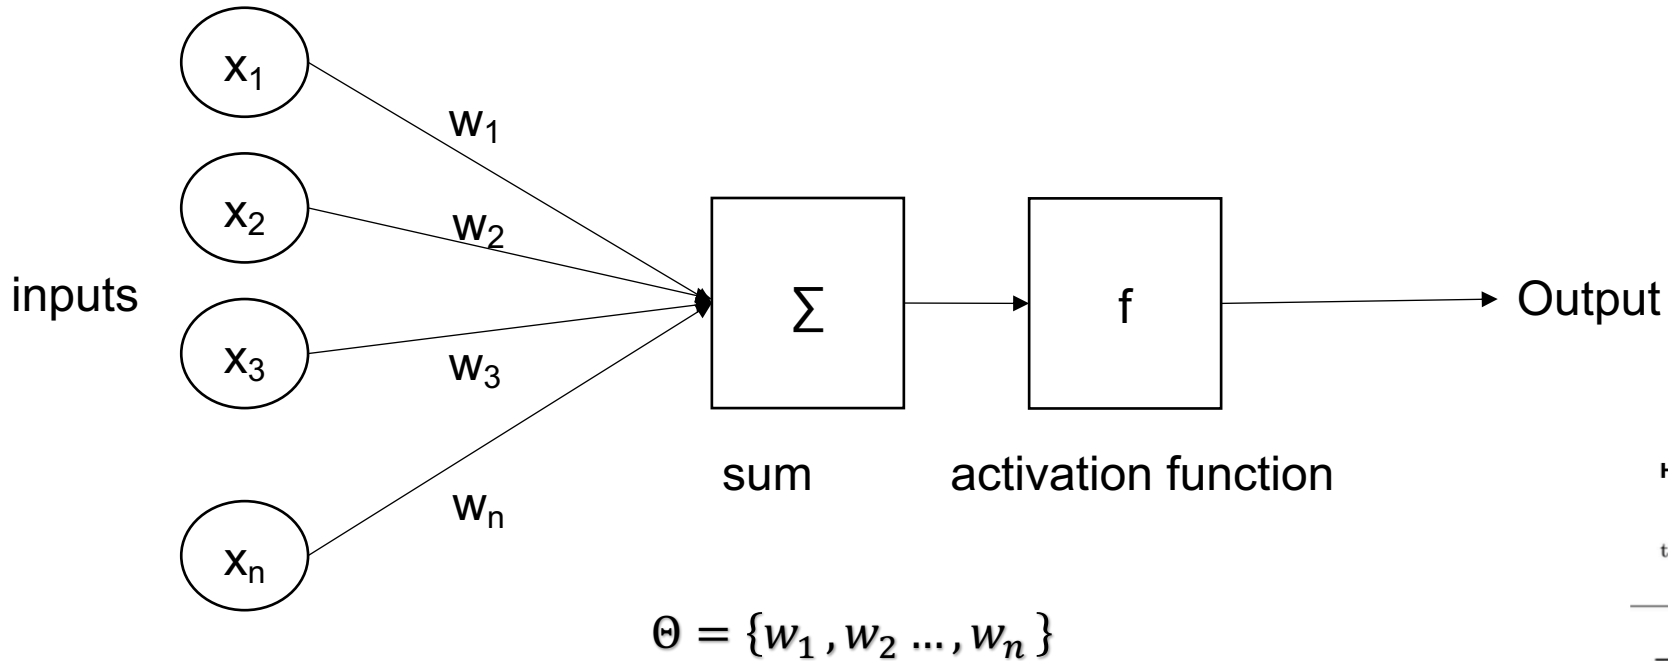
VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

Perceptron – Artificial Neuron




DEEP
LEARNING
INSTITUTE




most popular activation functions


Hyper Tangent Function




ReLU Function



Sigmoid Function



Identity Function



Single artificial neurons work well for linearly separable datasets (indeed output is the activation effect on a linear combination of the input)



VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

Neural Network




DEEP
LEARNING
INSTITUTE



Input Layer


Intermediate Layer

Output



$$\Theta = \{w_{1,1}^1, w_{1,2}^1, w_{1,3}^1, w_{2,1}^1, w_{2,2}^1, w_{2,3}^1, w_{1,1}^2, w_{2,1}^2, w_{2,3}^2\}$$

- Even when the data is not linearly separable



VSb TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

(Supervised) Learning

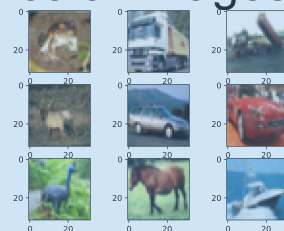


- Data domain $Z: X \times Y$

$X \rightarrow$ domain of the input data

$Y \rightarrow$ set of labels (knowledge)

$X: 32 \times 32$
color images



$Y: \text{labels}$

{ truck, car, horse, bird, boat }

Example (CIFAR10 dataset)

- Data Distribution is a probability distribution over a data domain
- Training set z_1, \dots, z_n from Z assumed to be drawn from the Data Distribution D
- Validation set v_1, \dots, v_m from Z also assumed to be drawn from D
- A machine learning model is a function that given a set of parameters Θ and z from Z produces a prediction
- The prediction quality is measured by a differentiable non-negative scalar-valued loss function, that we denote $\ell(\Theta; z)$



(Supervised) Learning



- Given Θ we can define the expected loss as: $L(\Theta) = \mathbb{E}_{z \sim D}[\ell(\Theta; z)]$
- Given D , ℓ , and a model with parameter set Θ , we can define learning as:
“The task of finding parameters Θ that achieve low values of the expected loss, while we are given access to only n training examples”
- The mentioned task before is commonly referred to as *training*

- Empirical average loss given a subset of the training data set $S(z_1, \dots, z_n)$ as:

$$\hat{L}(\Theta) = \frac{1}{n} \sum_{t=1}^n [\ell(\Theta; z_t)]$$

- Usually a proxy function, easier to understand by humans, is used for describing how well the training is performed (e.g., accuracy)



(Supervised) Learning



- The dominant algorithms for training neural networks are based on mini-batch stochastic gradient descent (SGD)
- Given an initial point Θ_0 SGD attempt to decrease \hat{L} via the sequence of iterates

$$\Theta_t \leftarrow \Theta_{t-1} - n_t g(\Theta_{t-1}; B_t)$$

$$g(\Theta; B) = \frac{1}{|B|} \sum_{z \in B} \nabla \ell(\Theta; z)$$

Definitions

B_t : random subset of training examples

n_t : positive scalar (learning rate)

epoch: update the weights after going over all training set



Computer Vision



DEEP
LEARNING
INSTITUTE



- Why? Focus on a kind of Deep Neural Network called Convolutional Neural Network (CNN)
- CNNs ability to extract multi-scale localized spatial features and compose them to construct highly expressive representations led to breakthroughs in almost all machine learning areas



VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

COMPUTER VISION TASKS




Image Classification




predicting the type or class of an object in an image

Image Classification + Localization




predicting the type or class on an object in an image and draw a bounding box around it

Object Detection



predicting the location of objects in an image via bounding boxes and the classes of the located objects

Image Segmentation





predicting the class to which each pixel in the image belongs to



Training Neural Networks



DEEP
LEARNING
INSTITUTE



Stochastic Gradient Descent

$$\theta_t \leftarrow \theta_{t-1} - \eta_t g(\theta_{t-1}; B_t)$$

$$g(\theta; B) = \frac{1}{|B|} \sum_{z \in B} \nabla \ell(\theta; z)$$



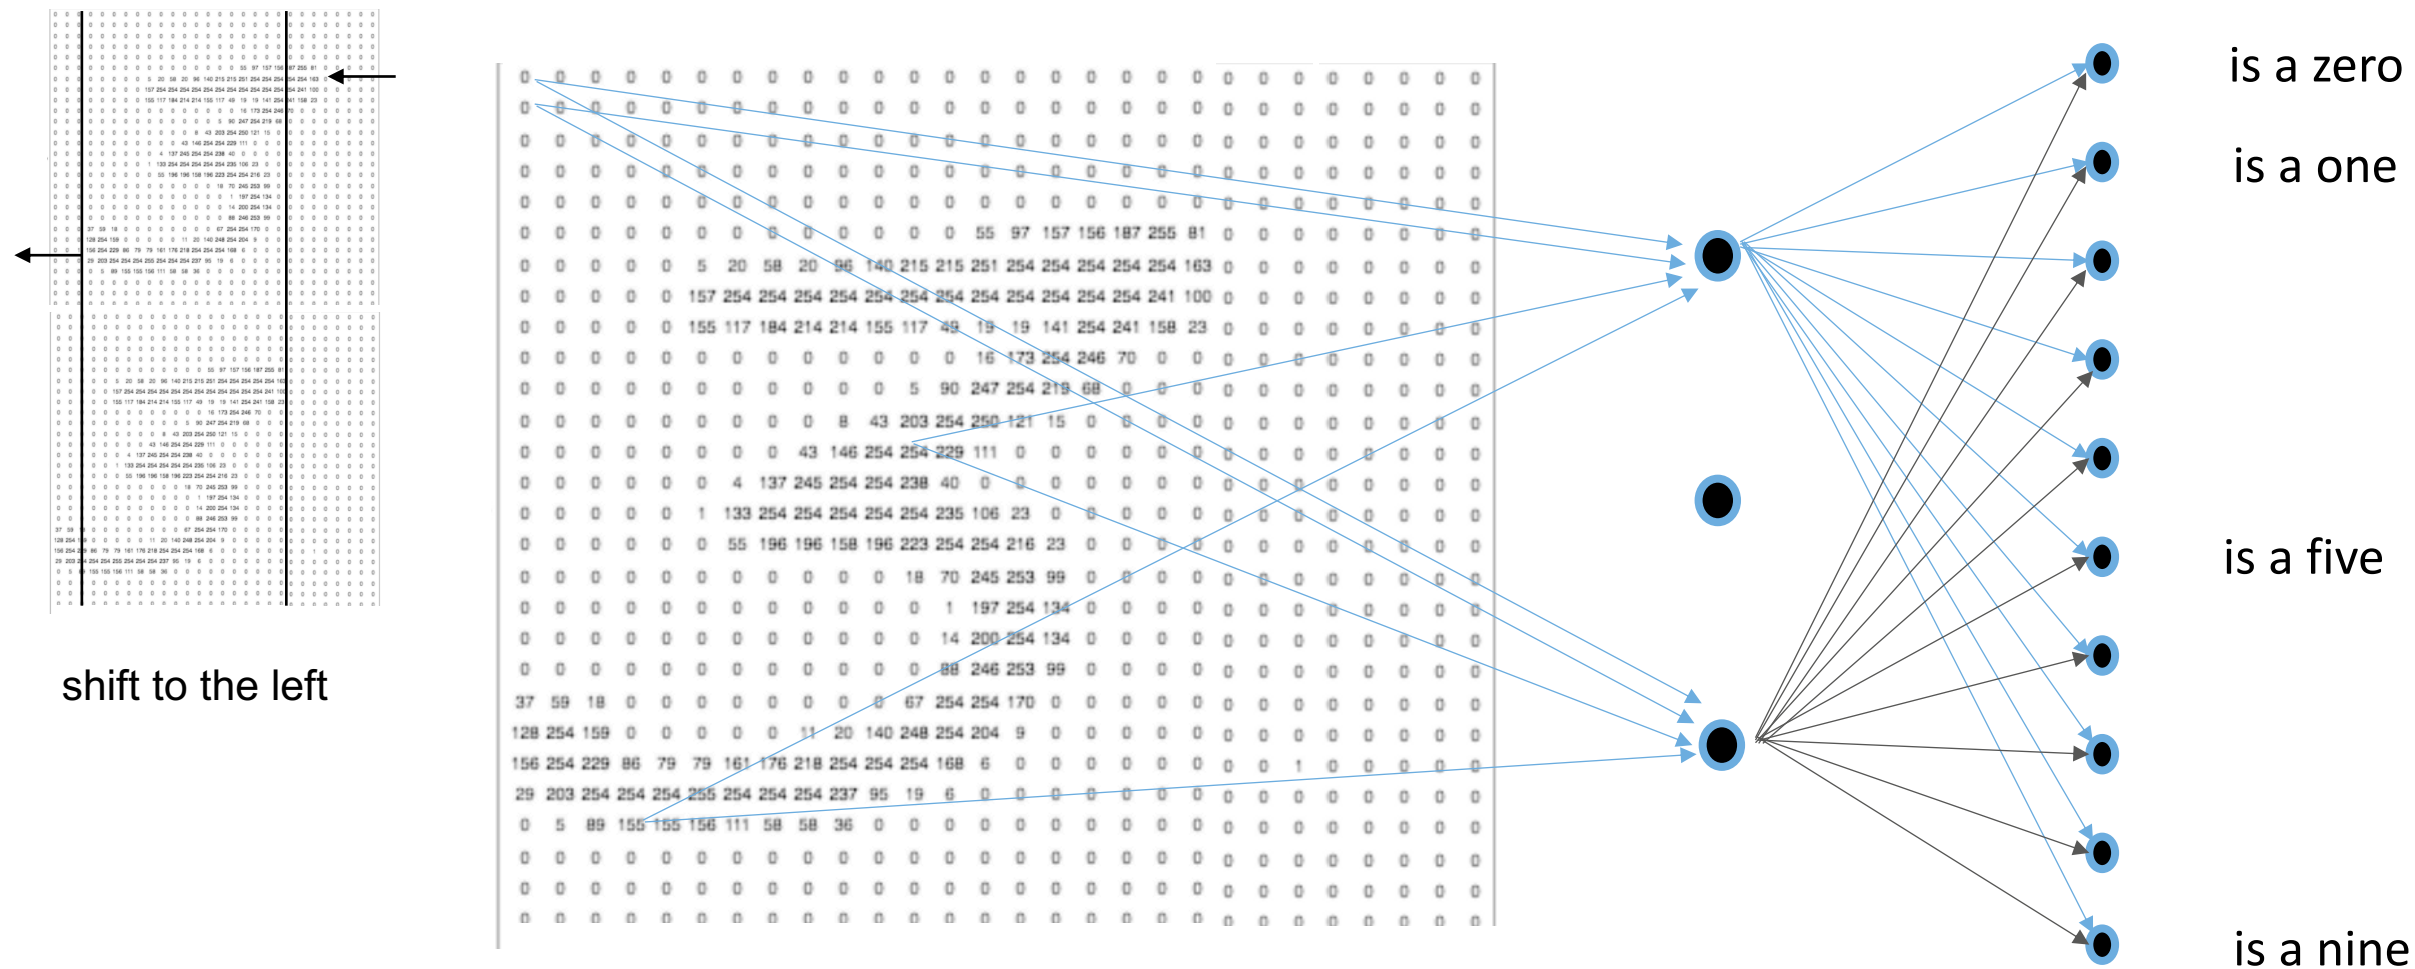
VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

Neural Networks for Image Classification



DEEP
LEARNING
INSTITUTE





VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

No More Feature Engineering



DEEP
LEARNING
INSTITUTE



VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

Learning features from data: Convolutions



DEEP
LEARNING
INSTITUTE



Input Image

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

Filter

-1	0	1
-2	1	2
-3	0	3

Convolved Image

	4						

receptive field

$$1 \times (-1) + 0 \times 0 + 1 \times 1 + 0 \times (-2) + 1 \times 1 + 0 \times 2 + 0 \times (-3) + 0 \times 0 + 1 \times 3 = 4$$

Filter is convolved with all the pixels of the image

How many units the filter moves horizontally or vertically is called **stride** and can be different in both dimensions

The stride defines the size of the convolved image

1	-1	0	1	0	1	0	1
0	-2	1	2	1	0	1	0
0	-3	0	3	0	1	0	1
1	0	1	0	0	1	0	0
0	0	0	0	1	0	1	0
0	0	1	0	0	1	1	1
0	0	0	0	0	0	1	0
0	0	1	0	0	1	0	1

1	0	1	0	0	1	0	1
0	1	0	0	1	0	1	0
0	-1	0	1	0	1	0	1
1	-2	1	2	0	1	0	0
0	-3	0	3	1	0	1	0
0	0	1	0	0	1	1	1
0	0	0	0	0	0	1	0
0	0	1	0	0	1	0	1

1	0	1	0	0	1	0	1
0	1	0	0	1	0	1	0
0	0	1	0	0	1	0	1
1	0	1	0	0	1	0	0
0	0	0	0	1	0	1	0
0	0	1	0	0	-1	0	1
0	0	0	0	0	-2	1	2
0	0	1	0	0	-3	0	3



VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

Filters



DEEP
LEARNING
INSTITUTE



Input Image:



Can we get only vertical lines out of this picture?

1	0	-1
---	---	----

filter 1



1	0	-1
1	0	-1
1	0	-1

filter 2



1	0	0	0	-1
1	0	0	0	-1
1	0	0	0	-1
1	0	0	0	-1
1	0	0	0	-1

filter 3



try the code yourself (in octave)!

```
I=imread(<path-to-image>);
GRAY=rgb2gray(I)
FILTER=[ 1 0 -1; 1 0 -1; 1 0 -1]; % filter 2
CONVOLUTED=conv2(GRAY,FILTER);
Imwrite(CONVOLUTED, <path-to-result>);
```

*The London skyline image is designed by Freepik



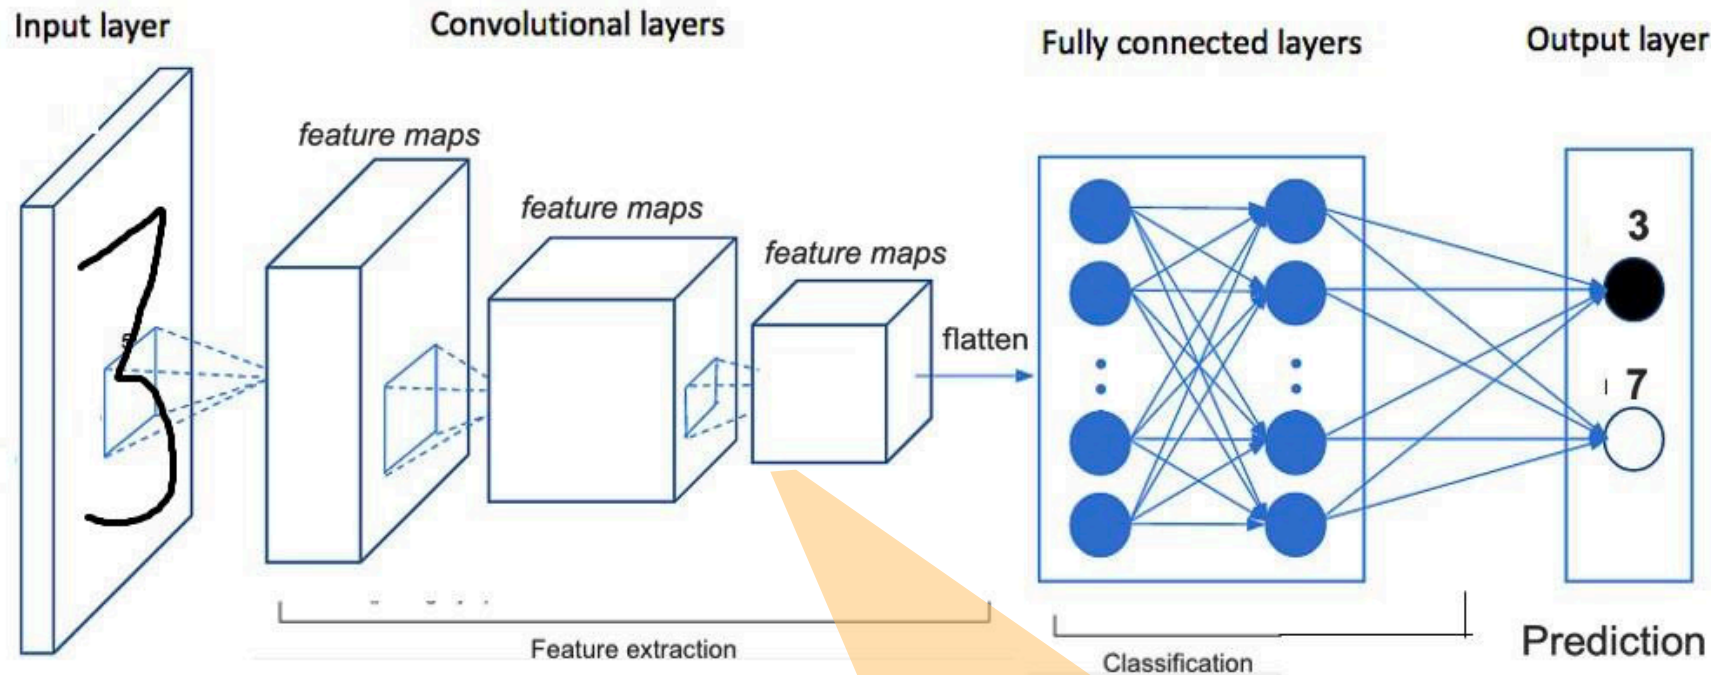
VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER


Convolutional Neural Networks (CNN)



DEEP
LEARNING
INSTITUTE



C
O
N
V
P
O
O
L
C
O
N
V
P
O
O
L



A pooling layer down sample the feature maps produced by a convolution into smaller number of parameters to reduce the computational complexity.


It is a common practice to add pooling layers after each one or two convolutions layers in the CNN architecture.



VS^B TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

CNN Architecture: A Common Pattern and its Influence




The execution time required during a forward pass through a neural network is bounded from below by the number of floating point operations (FLOPs).

This FLOP count depends on the deep neural network architecture and the amount of data.



LeNet Architecture



Architecture summary :

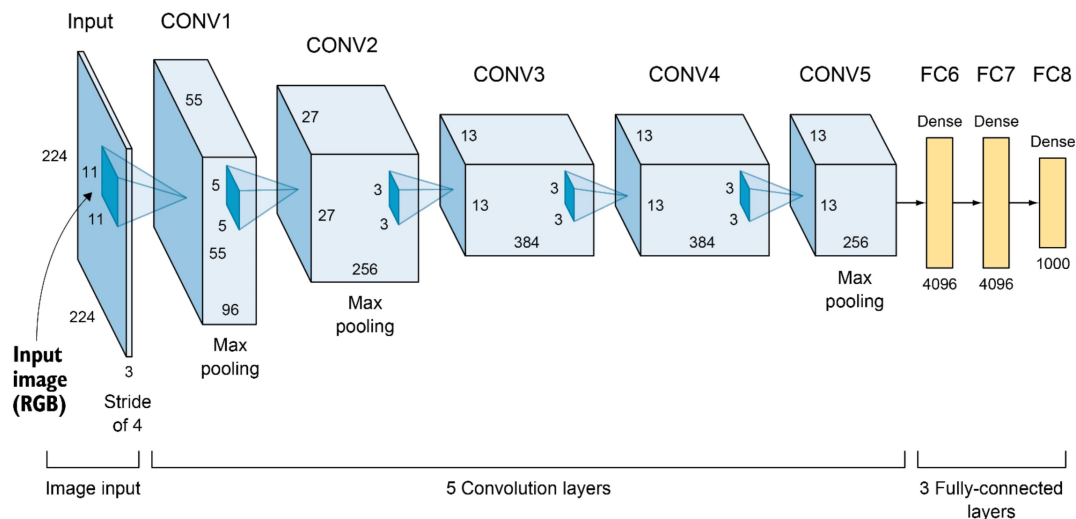
- 3 convolutional layers filters in all the layers equal to 5x5
(layer 1 depth = 6, layer 2 depth = 16, layer 3 depth = 120)
- As activation function the tanh function is used



AlexNet and VGG Architectures

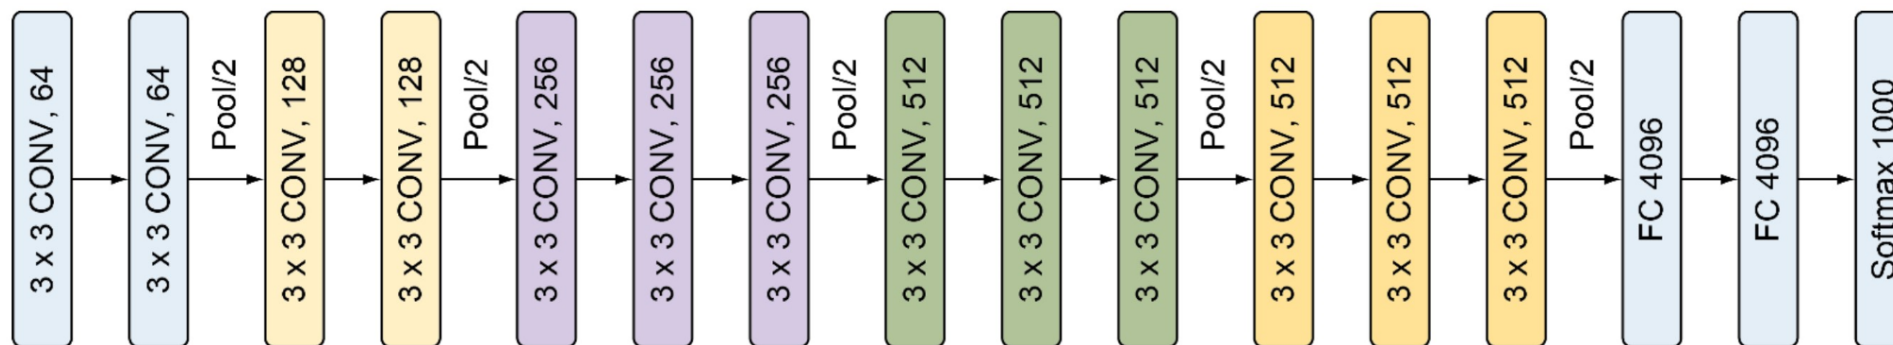


DEEP
LEARNING
INSTITUTE



~60,000,000
parameters

AlexNet



~138,000,000
parameters


VGG16



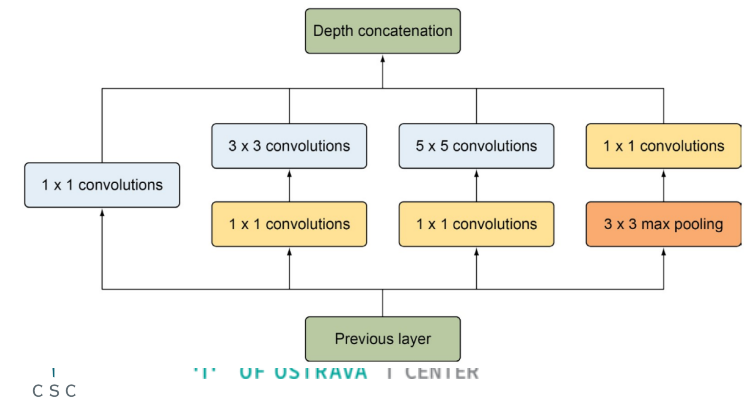
VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

GoogleNet





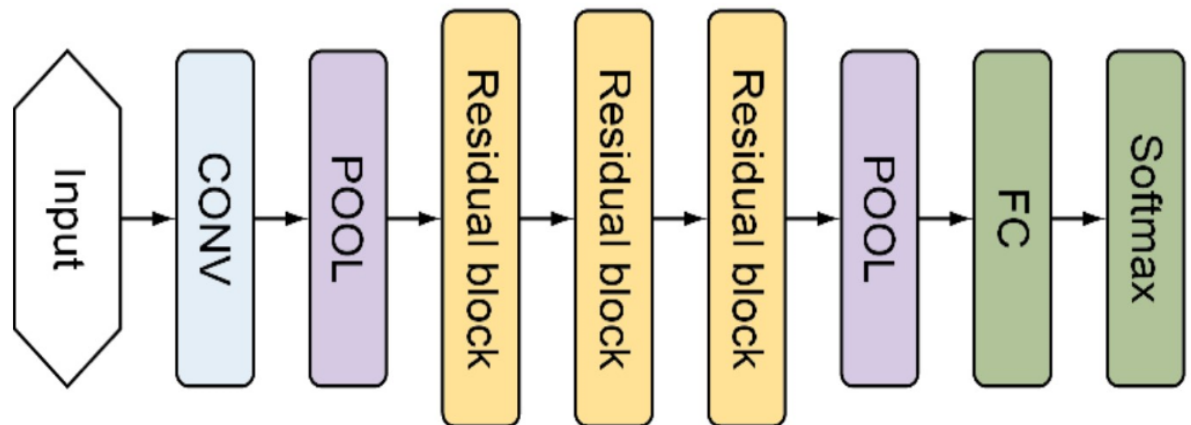
- What is the best kernel size for each layer?
- Concatenating filters instead of stacking them for reducing computational expenses



ResNet



DEEP
LEARNING
INSTITUTE



VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER


Increasing complexity



DEEP
LEARNING
INSTITUTE




7 Exaflops
60 Million Parameters




2015 - Microsoft ResNet
Superhuman Image Recognition

20 Exaflops
300 Million Parameters



2016 - Baidu Deep Speech 2
Superhuman Voice Recognition

100 Exaflops
8700 Million Parameters



2017 - Google Neural Machine Translation
Near Human Language Translation



VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER



Part 2


Images – Input and Output



DEEP
LEARNING
INSTITUTE



100	37	59	87	55	29	13	44
62	79	54	62	23	93	93	26
50	57	93	17	67	53	60	75
3	54	70	37	17	20	69	7
86	42	2	55	90	45	74	77
59	39	100	52	10	8	20	37
61	2	62	92	83	18	12	82
11	7	87	20	5	13	4	34



Deep Neural Network

0.04	0	0.02	0.01	0.92	0.01
Kites	Harrier	Vulture	Hawk	Eagle	Buzzards



Eagle



VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

One-Hot: Turning words into Numbers



- Numerical vector representation for each word
- Dictionary of N words
- Each word is a vector with $N-1$ zeros and one 1, at the position of the word in the dictionary
- A document can be represented as a sequence of these one-hot vectors
- One interesting property of this representation is that no information is lost



ONE-HOT ENCODING



DEEP
LEARNING
INSTITUTE



```
small_dict=['EOS','a','my','sleeps','on','dog','cat','the','bed','floor'] # 'EOS' means end of sentence.
```

```
import numpy as np #numpy is "numerical python" and is used in deep learning mostly for its n-dimensional array
X=np.array([[2,6,3,4,2,8,0],[1,5,3,4,7,9,0]],dtype=np.int32)
print([small_dict[ind] for ind in X[1,:]]) #Feel free to change 1 to 0 to see the other sentence.
```

```
['a', 'dog', 'sleeps', 'on', 'the', 'floor', 'EOS']
```

one-hot encoded inputs

```
[[[ 0.  0.  1.  0.  0.  0.  0.  0.  0.  0.]
   [ 0.  0.  0.  0.  0.  0.  1.  0.  0.  0.]
   [ 0.  0.  0.  1.  0.  0.  0.  0.  0.  0.]
   [ 0.  0.  0.  0.  1.  0.  0.  0.  0.  0.]
   [ 0.  0.  1.  0.  0.  0.  0.  0.  0.  0.]
   [ 0.  0.  0.  0.  0.  0.  0.  0.  1.  0.]
   [ 1.  0.  0.  0.  0.  0.  0.  0.  0.  0.]
```

```
[[ 0.  1.  0.  0.  0.  0.  0.  0.  0.  0.]
 [ 0.  0.  0.  0.  0.  1.  0.  0.  0.  0.]
 [ 0.  0.  0.  1.  0.  0.  0.  0.  0.  0.]
 [ 0.  0.  0.  0.  1.  0.  0.  0.  0.  0.]
 [ 0.  0.  0.  0.  0.  0.  0.  1.  0.  0.]
 [ 0.  0.  0.  0.  0.  0.  0.  0.  0.  1.]
 [ 1.  0.  0.  0.  0.  0.  0.  0.  0.  0.]
```

shape of the input
(2, 7, 10)



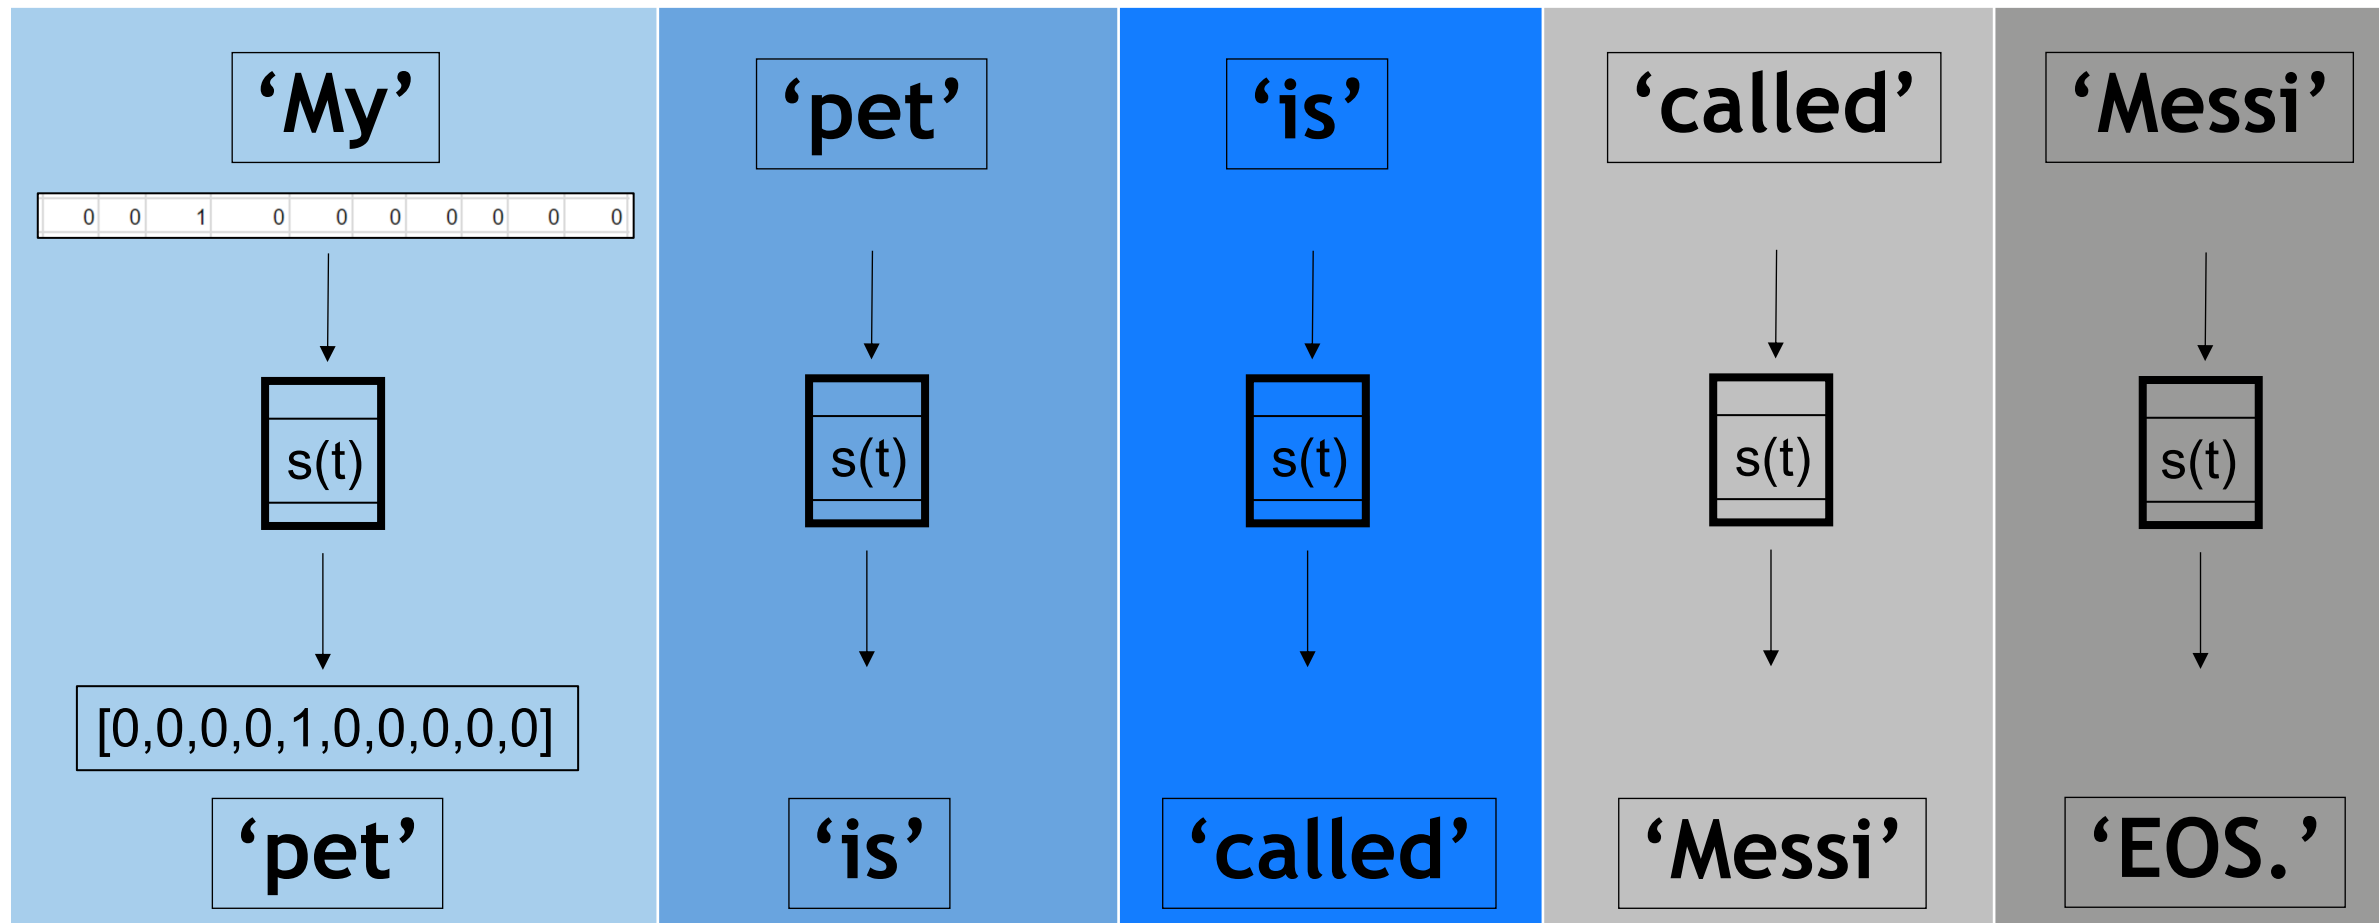
VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

Generating Language



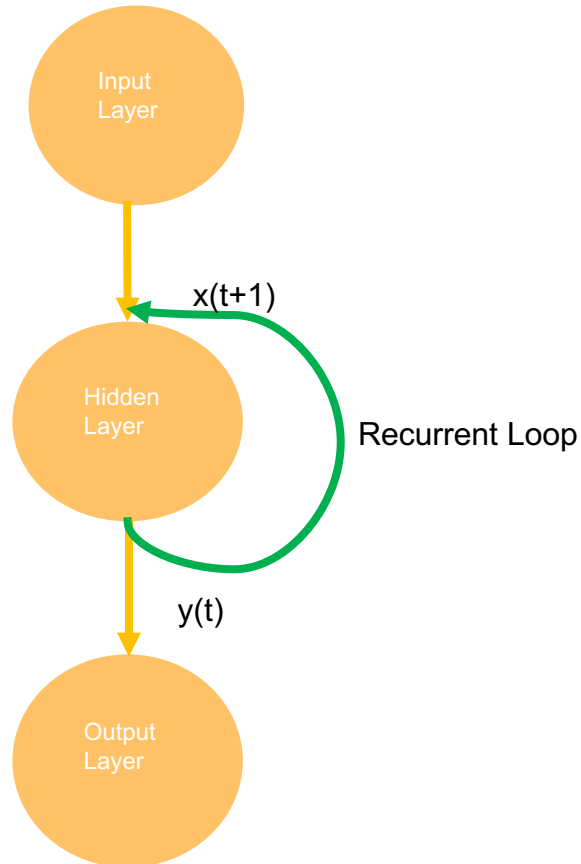
DEEP
LEARNING
INSTITUTE



VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

Recurrent Neural Network (RNN)



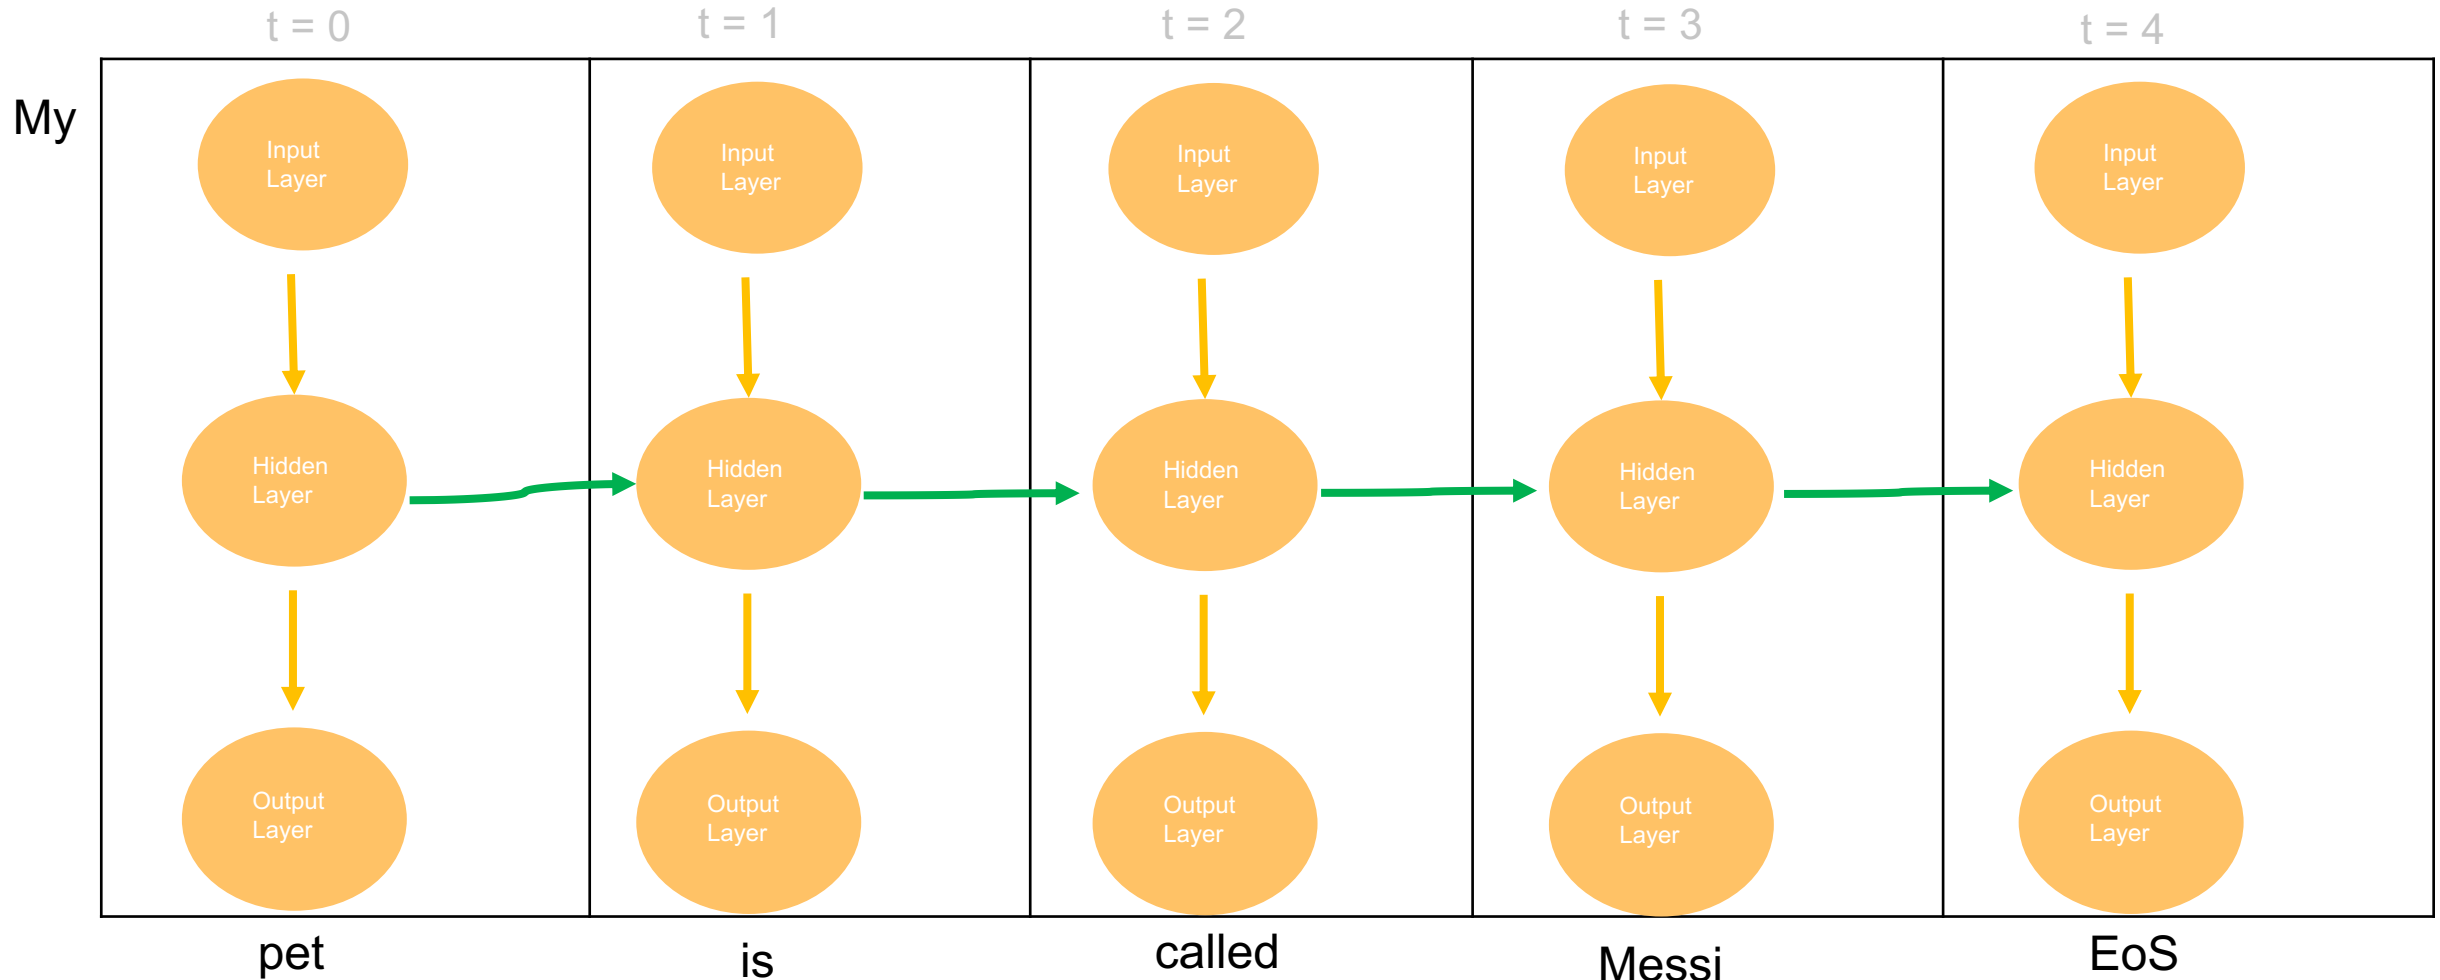
- Enable neural nets to remember past words within a sentence
- Recycle the output of the hidden layer at time t by adding as next input at time $t+1$
- Easier way of understanding its working behavior unrolling the net



RNN Unrolled



DEEP
LEARNING
INSTITUTE



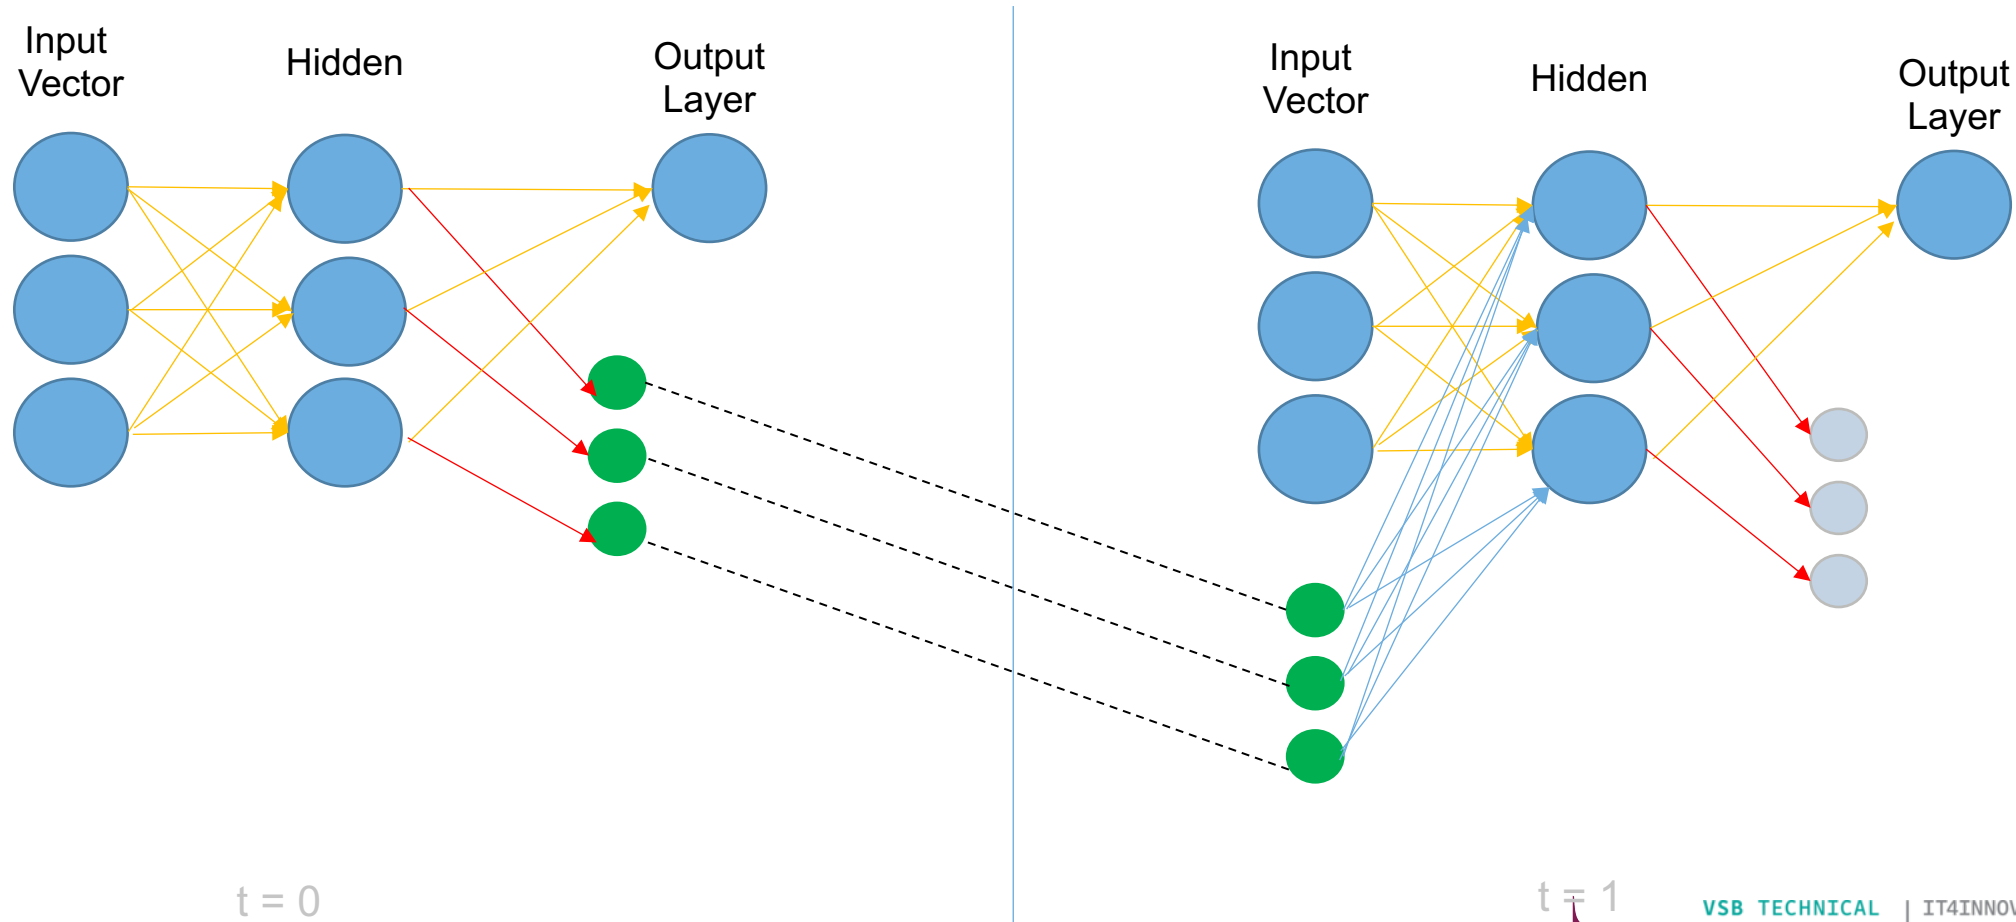
VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

Understanding RNNs



DEEP
LEARNING
INSTITUTE



$t = 0$



VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER

Shortcomings of RNNs



1. Expensive training procedures
 1. A back propagation iteration updates each of the unrolled steps
2. Relationships between a word and the words that have appeared before
 1. Some words in some languages depends on what comes afterwards
 2. Bidirectional Recurrent Neural Networks
3. How many words in the past (or the future) influence the next word
 1. e.g., “The young woman, having found a free ticket on the ground, went to the movies.”
 2. Need of remembering the past across the entire input (young woman -> went)
 3. LSTM (Long Short Term Memory) Cells



Summary



DEEP
LEARNING
INSTITUTE



- Brief introduction to Deep Learning with emphasis in Deep Convolutional Neural Networks
- Review of basic concepts: from perceptron to the learning task
- Debrief of most important concepts of neural network architectures
- Introduction to language modeling and RNNs



VSB TECHNICAL
UNIVERSITY
OF OSTRAVA

IT4INNOVATIONS
NATIONAL SUPERCOMPUTING
CENTER