SSDA 2019 Introduction into Machine Learning and Neural Networks

Marcus Liwicki EISLAB Machine Learning (chair) Luleå University of Technology

Slides at: bit.ly/2019-ssda

LULEÅ UNIVERSITY OF TECHNOL<u>OGY</u>

EISLAB: Embedded Intelligent Systems LAB

What is the Biggest Break Through of AI?



Handa







Overview

- A few words about Luleå and LTU
- Concepts of Machine Learning
- Biological Motivation
- Foundations for Neural Networks

Thursday:

- LSTM Basics and Current Trends
- LSTM for HWR Interactive Session (bring your notebook)

Quiz: Where is Luleå?

Call a la



Marcal

Our strengths

- Education and research in close collaboration with industry and society
- Cutting-edge environments and testbeds for innovation and research
- Research that creates global societal benefit



Luleå university of technology



LULEÅ UNIVERSITY OF TECHNOLOGY

han a st



SSDA 2019, Marcus Liwicki: Introduction into ML and NN 7

A AND A

OF TECHNOLOGY

The Space University of Sweden

- Competence and international network
 - The only Master Programme in Space Engineering in Sweden
- NanoSat Lab
- RIT Space for innovation and growth
- Water on Mars



Our students

Contract education (♀46%, ∂54%) 1.46% Second-cycle Doctoral stu-dents (\$37%, education/Master (♀**4**5%, ♂55%) 4.81% 74% Freestanding First-cycle courses (♀58%, education **∂**42%) (♀46%, ♂54%) 33.09% 56.90%

(Percentage women \bigcirc , men \bigcirc)

When I



LTU AI VISION

Vision

We transform the Luleå region into a sustainable flagship in applied Artificial Intelligence (AI) at both national and European level. A region that will demonstrate innovative and responsible applications of AI in real-life use cases of benefit to industrial, research, educational, social, and health aspects of human life and society.

Mission statement

We aim at creating a strong and active ecosystem in AI, an ecosystem that directly connects fundamental research with real-life applications and demonstrations of AI in the industrial sector and beyond, and by that contribute to a safe and measurable strong impact of AI innovations in everyday life.



We are Well Recognized

- 3 LTU AI-researchers in IVA's 100 list of innovators
- Selected in IVA's top Achievements in Engineering 2018
- Leaders in various machine learning competitions
- Leaders in Aerial Robotics, and many other areas
- Young investigator awards in various Applied AI topics
- World's 1st PhD student in living labs
- Sweden's only NVIDIA DLI ambassador
- Marcus among top 50 AI researchers



Bayer, J., Behnke, S. Sven Bengio, Y., Bertolami, R., Bunke, H., Catanzaro, B. Ciresan, D. C., Coates, A., Couprie, C., Courville, A., Fernandez, S., Gambardella, L. M., Gerstner, W. Giusti, A., Goller, C.; Goodfellow, Ian Graves, A., He, Kaiming Hinton, G. E Hochreiter, S. Huval, B., Karpathy, Andrej Krizhevsky, A., Kilchlan A

1, 0.-1.

Japan Germany Lamblin, P., Germany Larochelle, H. France Canada LeCun, Y. Switzerland Li, Fei-Fei Switzerland Li, X.; JSA Liwicki, M., Romania Manning, Ch JSA Masci, J France Meier, U Canada Mohamed, A.-R., Spain Ng, A.Y Italy Popovici, D. "rise of Al" conference 2019 (Berlin): Germany Italy Marcus Liwicki is on the list of the Germany **Top 50 most famous AI researchers** UK. in the world China Togelius, J UK Waibel, Alex Germany Wang, T., JSA Wierstra, D., Slovakia Wu, D. J., Canada Ukraine

France

Canada

France

China

China

Germany

Australia

Egypt UK (HK)

Switzerland

Romania

Canada

Sweden

China

China

Germany

Netherlands

12

Italy



Reading Systems

Contact: Marcus.Liwic ki@ltu.se

Document Analysis modern & historical





LTU

Natural Language Processing Chatbots, intent classification

60.G

- Open PhD positions - Scholarships for women

Google

Top Start-Ups Micros

oft



IBM

Open

Source

SAP

Impact

- 10 open source frameworks and Web **Services**
- Young Investigator Award
- World's best in OCR. intent, writer, and classification, ...
- **Funding from** industry and country

https://github.com/kumar-shridhar/HackathonLulea

Overview

- A few words about Luleå and LTU
- Concepts of Machine Learning
- Biological Motivation
- Foundations for Neural Networks

Thursday:

- LSTM Basics and Current Trends
- LSTM for HWR Interactive Session (bring your notebook)

Machine Learning is inspired by Human Intelligence

easy

B

Bebop

Bibopo estas muzikstilo, kiu komence de la 1940-aj jaroj en ĵazo anstataŭis la svingon kiel ĉefa stilo kaj per tio formis la bazon por la moderna ĵazo. ructures (Deep





SSDA 2019, Marcus Liwicki: Introduction into ML and NN 15

Concept of Machine Learning



Supervised vs. unsupervised learning





A CAR S

Unsupervised ML algorithm









Intuitive Intro towards Machine Learning

<u>https://experiments.withgoogle.com/ai</u>
 <u>https://quickdraw.withgoogle.com/</u>



millen a Stor

http://bit.ly/liwicki-vdl-17 (all my lecture material)



Key Terms in Machine Learning

Machine Learning: Computers learn without being specifically programmed **Pattern Recognition**: Computer recognizes structures (text, objects) Learning/Training: Give the computer samples to automatically adjust itself **Ground Truth**: Labels generated by human expert – what we expect to get **Training Data:** Representative example patterns (images) with ground truth Validation Data: A separate set to find the best method **Typical Split Test Data:** Another separate set for final evaluation raining **Deep Learning**: Buzzword – to get funding Validation 1 **Typical Neural Networks:** Validation 2

■Zest

Convolutional Neural Networks (CNN) – for images (fixed size) Long Short-Term Memory (LSTM) – for n-dimensional sequences

Overview

- A few words about Luleå and LTU
- Concepts of Machine Learning
- Biological Motivation

Bonus content unlocked: Do we all need higher salaries?

Foundations for Neural Networks

Thursday:

- LSTM Basics and Current Trends
- LSTM for HWR Interactive Session (bring your notebook)

Let me introduce Marvin

Not the real name, but this applies to multiple of my current and former PhD students



- Is a great person who joined our team
- Searches for internships at Google, fb, ...
 rejected because he is not from MIT ;)
- However, she is very skilled and got offers from other companies
 - high salary (\$ 5000-11000 per month)
 - Also I had ~5000 net salary
 - Stepped down to ~4000 in Sweden

LULEÅ VNIVERSITY OF TECHNOLOGY

People need opportunities not money And you can find that in Sweden!

- Stable base
 - Tenured positions
 - Health, social security
- Fruitful environment
 - High-impact researchers
 - Start-up culture ?
- Perfect infrastructure
 - Computing & research facilities
 - Support for travel, organization, collaboration, ...



Overview

- A few words about Luleå and LTU
- Concepts of Machine Learning
- Biological Motivation
- Foundations for Neural Networks

Thursday:

- LSTM Basics and Current Trends
- LSTM for HWR Interactive Session (bring your notebook)

Neural networks as core technology for Al



>10¹¹ neurons

Goal: make computers intelligent Idea: simulate neural behavior on PC



Learning of Neural Networks

- Learning takes place at the synapses
 - Efficiency is increased if more ionic channels open
 - So-called NDMA receptors (N-methyl-d-aspartic acid)
 - Much exitation leads to unblocking (Mg⁺) of receptors
- Stored information needs to be refreshed periodically





Fundamental: Human Mind is More Complex

CHI al



Overview

- A few words about Luleå and LTU
- Concepts of Machine Learning
- Biological Motivation
- Foundations for Neural Networks

Thursday:

- LSTM Basics and Current Trends
- LSTM for HWR Interactive Session (bring your notebook)

Multi-Layer Perceptrons for Object Recognition Object image



Neural Network

Ba







LULEÅ UNIVERSITY OF TECHNOLOGY

How Does Machine Learning Work?



STALL!

Human-designed Features

- Number of circles
- Biggest shape
- Number of dark pixels





attan a So

Feature vector (at timestamp t) $x_1^t, \dots, x_n^t \chi^t$ Perceptrons in the individual layers - Aggregation function $a^t = \sum_{w_i x_i^t} w_i x_i^t$ - Activation function (squashing function) $b_h^t = h(a^t)$ h: $\sigma(x) = \frac{1}{1 + e^{-x}}$ (sigmoid) $\tanh(x) = \frac{e^{2x} - 1}{e^{2x} + 1}$ SSDA 2019, Marcus Liwicki: Introduction into ML and NN 31



all all a sol

æ

Backpropagation - conceptual



CIT & COR



Deep Learning – Various Perspectives

- (Very) deep NN
- Learning of representations

aller a Cor

- Learning of sequences
- Learning of concepts/relations





Question: When did Deep Learning start?

2011 – AlexNet 4 Object Recognition? 2006 – The American Dream Trio?





mite	container ship	motor scooter	leopard
mite	container ship	motor scooter	leopard
black widow	lifeboat	go-kart	jaguar
cockroach	amphibian	moped	cheetah
tick	fireboat	bumper car	snow leopard
starfish	drilling platform	golfcart	Egyptian cat
grille	mushroom	cherry	Madagascar cat
convertible	agaric	dalmatian	squirrel monkey
grille	mushroom	grape	spider monkey
pickup	jelly fungus	elderberry	titi
beach wagon	gill fungus	ffordshire bullterrier	indri
fire engine	dead-man's-fingers	currant	howler monkey

No! – 1980, Fukushima

Bonus Question: When did the Promgramming of Intelligent Machine Algorithms start?

She suggested the data input that would program the machine to calculate Bernoulli numbers

"Supposing, that the fundamental relations of pitched sounds in the science of harmony and of musical composition were susceptible of such expression and adaptations, the **engine** might compose elaborate and scientific pieces of music of any degree of complexity or extent."



Ada Lovelace, mathematician in 1840's

DL Enablers – We need progress in all of them

AI Ethics – Trustworthy AI

GPU

Al architectures

GIT A ST



SSDA 2019, Marcus Liwicki: Introduction into ML and NN 37

BIG DATA

Basic Network Architectures

Deep Convolutional Neural Networks (CNN) for images





1997 – Hochreiter & Schmidhuber

1980 - Fukushima 1998 – LeCun (MNIST)

> LULEÅ JNIVERSITY OF TECHNOLOGY

ImageNet Large Scale Visual Recognition Challenge results

wrong

Prominent Break-Through:





2012 won by CNN (AlexNet)



Data: imageNet

Deep Learning in Production

Speech Recognition Recommender Systems Autonomous Driving Real-time Object Recognition Robotics Real-time Language Translation Many More...







Useful Toolkits (Most Popular)

- Keras: <u>https://elitedatascience.com/keras-tutorial-deep-learning-in-python</u>
- deeplearnjs: <u>https://deeplearnjs.org/</u>

alless a Stor

- Deeplearning4j: <u>https://deeplearning4j.org/</u>
- https://mxnet.incubator.apache.org/how_to/finetune.html
- Tensorflow (and interesting visualizations in tensorboard)
 - <u>https://www.tensorflow.org/get_started/</u>
 - <u>https://www.tensorflow.org/programmers_guide/summaries_and_tensorboard</u>

Caffe2 and Caffe

- https://caffe2.ai
- http://caffe.berkeleyvision.org
- PyTorch & Torch
 - <u>http://pytorch.org/</u>
 - <u>http://torch.ch/</u>

UI's and more for end-users https://cloud.google.com/ml-engine/docs/ https://aws.amazon.com/machine-learning/ https://azure.microsoft.com/en-us/overview/machinehttps://developer.nvidia.com/digits https://developer.nvidia.com/digits https://developer.nvidia.com/digits https://developer.nvidia.com/digits

UNIVERSITY OF TECHNOLOGY

SSDA 2019, Marcus Liwicki: Introduction into ML and NN 41

Slides at: bit.ly/2019-ssda

Research sasy-to-use

Useful Models

The number one place for finding pre-trained models

- <u>https://github.com/BVLC/caffe/wiki/Model-Zoo</u>
- (also gives hints for successful applications)
- A bit easier to understand, because it is curated
 - <u>https://modeldepot.io/</u>
- Small, but with demos
 - http://pretrained.ml/

alless a 2 3 2

- https://github.com/keras-team/keras/tree/master/examples
- Individual task: Look at both websites and try to find a model working in a domain which is interesting for YOU

SSDA 2019, Marcus Liwicki: Introduction into ML and NN 42

Slides at: bit.ly/2019-ssda

Other Useful Links

- https://teachablemachine.withgoogle.com/
- http://playground.tensorflow.org

GP GP

1 And

- https://experiments.withgoogle.com/ai
- https://transcranial.github.io/keras-js/#/imdb-bidirectional-lstm
- https://transcranial.github.io/keras-js/#/mnist-acgan
- https://quickdraw.withgoogle.com/



We can Learn From Failures & Success

- Deep Learning and AI is not the answer to everything
 - https://www.techrepublic.com/article/top-10-ai-failures-of-2016/
- An extension of reinforcement learning is Artificial Curiosity
 - Could (and definitely would) go terribly wrong

atters a 2 20

https://blog.statsbot.co/deep-learning-achievements_4 63e034257



LULEÅ UNIVERSITY OF TECHNOLOGY

(WA)

LARAS D

Thank You + Lab Members & Beyond



Marcus



Priamvada



then som

Pedro



Gustav

György



Fotini



Oluwatosin Incoming 2019 ... Saleha Javed



Kaiserslautern

And colleagues

Fribourg

LTU



Hidden Quest unlocked:

- During lunch sit at a table with at least (one of each):
 - 1. Person not having your gender
 - 2. Person not working in the same city as you
 - 3. One of the lecturers of this summer school
- easy mode (today)
 - I will not sit at a VIP table, if there is one
 - We (lecturers) will sit at different tables (hopefully)

