

Deep Learning in Graphics & Physics

Informatics 15, N. Thuerey
<https://ge.in.tum.de/>

Bachelor/Master Seminars

One of the deep-learning break-throughs:



The winning solution (AUC 0.99) of the **Camelyon challenge on detecting metastatic cancer** beats the human pathologist benchmark (AUC 0.96)

A CNN designed by a team at the University of Toronto wins the ImageNet Challenge bringing down the error rate to 16% (compared to 25% 2011)

Fei Fei Li and colleagues at Princeton University start to collect a large database of annotated images, the **ImageNet**

A group around Yann LeCun successfully applies a back-propagation algorithm to a multi-layer neural network, **recognizing handwritten ZIP codes**

Frank Rosenblatt develops the **Perceptron**, an early neural network enabling pattern recognition based on a two-layer learning network

2016

2015

2012

2009

2007

2006

1989

1986

1957

A CNN by team from Microsoft beats the human benchmark (5% error rate) by bringing down the error rate to 3% in the ImageNet Challenge

A group around Andrew Ng introduce **Graphics Processing Units (GPUs)** for Deep Learning making them applicable on a large scale

Hinton summarizes ideas of **multilayer neural networks** and training them to generate sensory data rather than to classify it

Rumelhart, Hinton, and Williams introduce **backpropagation** as a learning procedure for "networks of neuron-like units"

A blue bowl filled with white, fluffy dumplings, with a whisk above them.

DL for other applications?
I.e. graphics & simulations.

Overview

- Purpose
- Topics
- Requirements & sign-up

Purpose of a Seminar

- Introduction to **scientific work**
- Get **practical experience** for basic research techniques (which are hopefully useful later on!)
- Actual content - hopefully interesting as well

Content

- 16 topics each from all areas of **machine learning** (plus some foundation papers)
- Focus on *computer graphics* (seminar instance 1) and *physical simulations* (seminar instance 2)
- Check previous years for examples...
- 1 Paper per person (talk & short report)

Content

- No “introduction to deep learning”, if unsure about DL or physics content - please check examples from previous years
- Implementations not required

Additional Purpose of this Seminar


- Language: English (talk, discussion & report)
- Quality of expressions won't be graded!
- Primary goal: exposure to English language
- Practice, practice, practice...



Solver-in-the-Loop paper:

<https://ge.in.tum.de/publications/2020-um-solver-in-the-loop/>

Test Case $Re=273.4$ with SOL_{16}




Source

MAE=0.167

SOL_{16}


MAE=0.130

Reference




Preview

<https://www.youtube.com/watch?v=Ih8EfvOzBOY>




tempoGAN paper:

<https://www.youtube.com/watch?v=i6JwXYypZ3Y>




Requirements



- **Short report** of 4-6 pages (**only pdf**), **2 weeks before** talk, meeting with supervisor
- **Talk**
 - Duration ca. 30 minutes (hard limit)
 - Send your semi-final slides (**pdf**) **1 week before** your talk (**otherwise, the talk will be cancelled**)
- Be present at all talks and active in discussions (max. 1 miss)
- Final grade: **presentation/participation (75%), report (25%)**

Requirements



- Sign up via [TUM matching system](#)
- No [additional](#) documents required
- Post Student ID == Immatrikulationsnr. “*now*” in chat for preferred matching
- Typical times (tentative)
 - [Weekly, Tuesdays 14:00: Graphics](#)
 - [Weekly, Wednesdays 14:00: Physics](#)

Questions?
