

# Introduction to Deep Learning

# The Team

Lecturer



Prof. Matthias Niessner

TAs



Yujin Chen



Haoxuan Li



Chandan Yeshwanth

# Student Tutors



Shaotong  
Chen



Han  
Wu



Hongtai  
Yuan



Tom  
Schimansky



Estevao  
Gomes

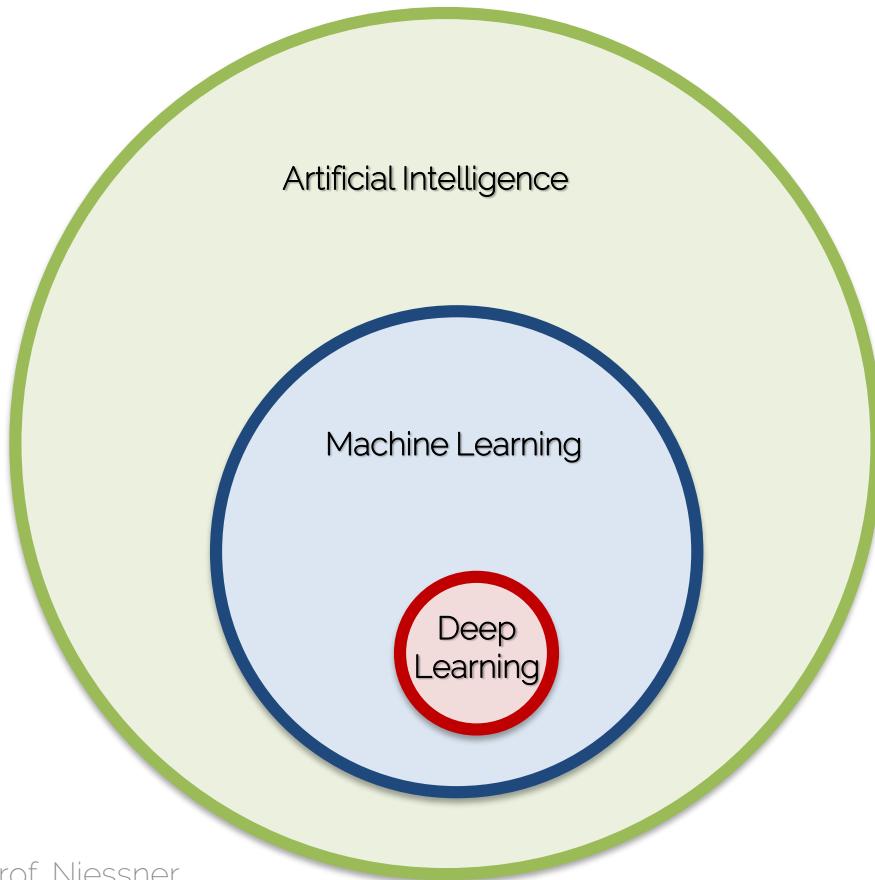


Ignacio  
Dassori

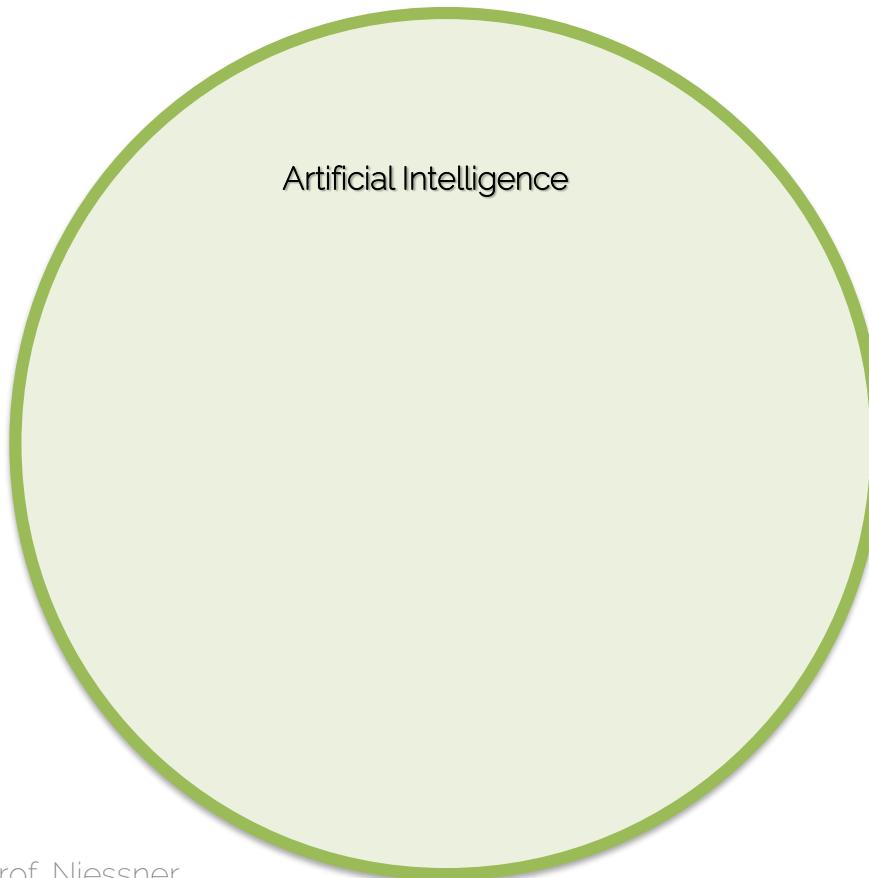


Mats  
Grobe

# What is this lecture about?



# What is this lecture about?

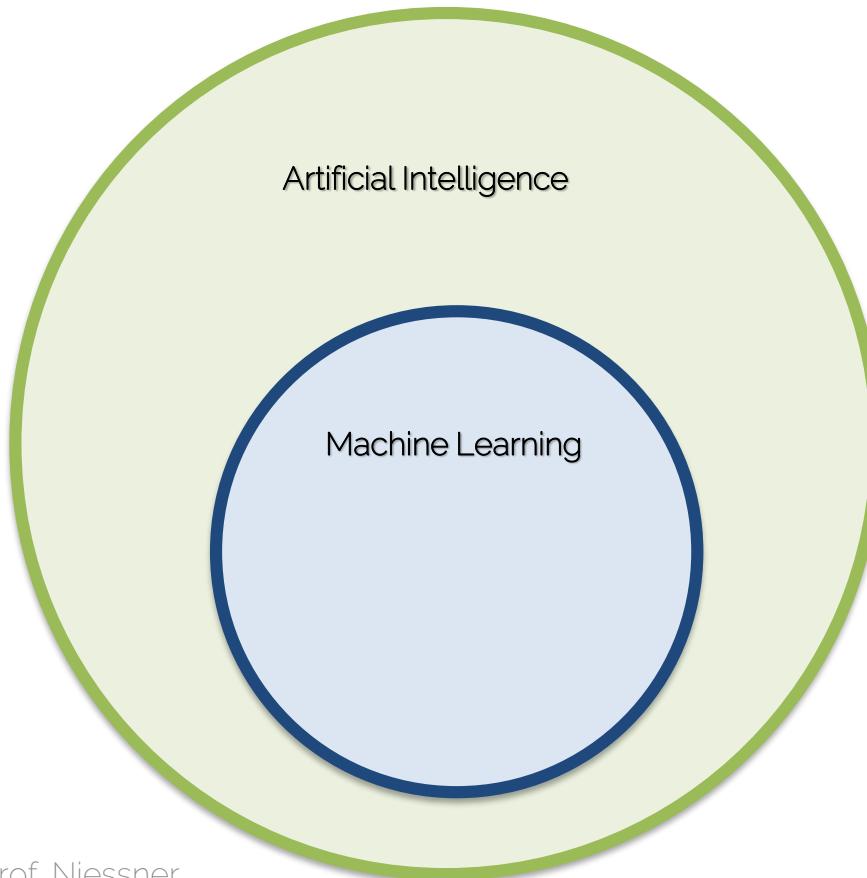


- AI Methods
  - Broad definition!
    - "if" statements

```
if (cold)
    turn_up_heat();
```
    - Binary Search
    - Dijkstra, A<sup>\*</sup>, ...
    - Prime, Kurskal, ...
    - Logic algorithms, etc.

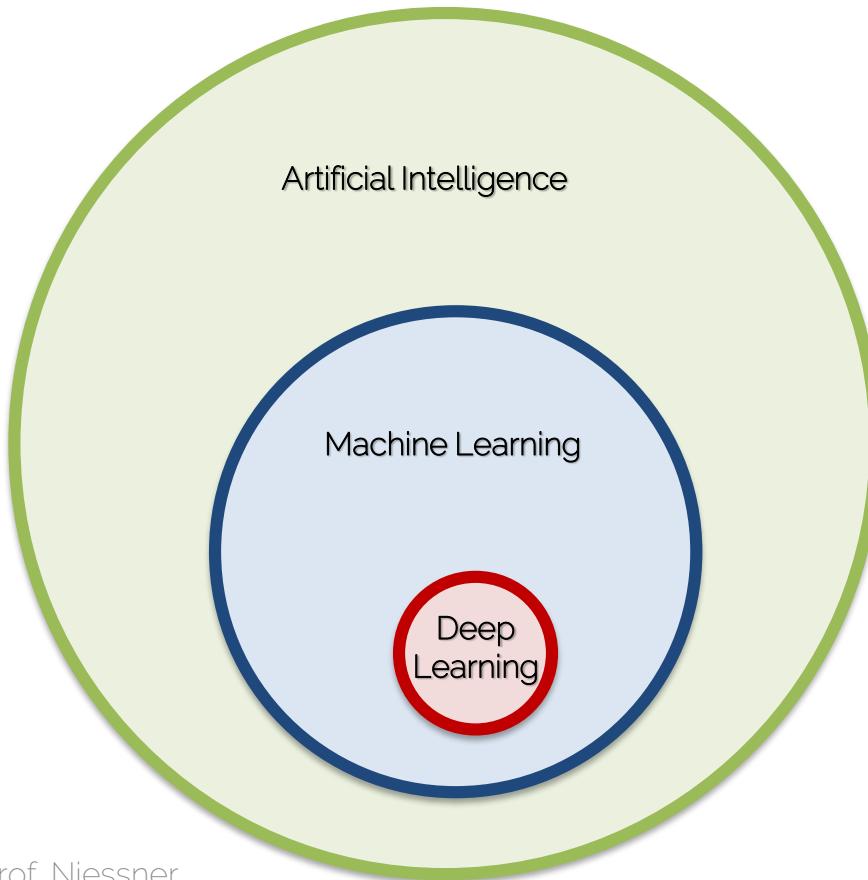
— ...

# What is this lecture about?



- ML Methods
  - Linear/logistic regression
  - Support Vector Machines
  - Random Decision Trees, Forests, Jungles, ...
  - ...

# What is this lecture about?



- Deep Learning
  - ML-methods leveraging neural networks
    - Multi-layer perceptrons
    - Convolutional neural networks
    - Recurrent neural networks
    - Transformers
    - Generative models, etc...

# What do you see?



# What does a machine see?

## HOW TO CONFUSE MACHINE LEARNING



# What does a machine see?



**Q: Is the door of the truck cab open?**



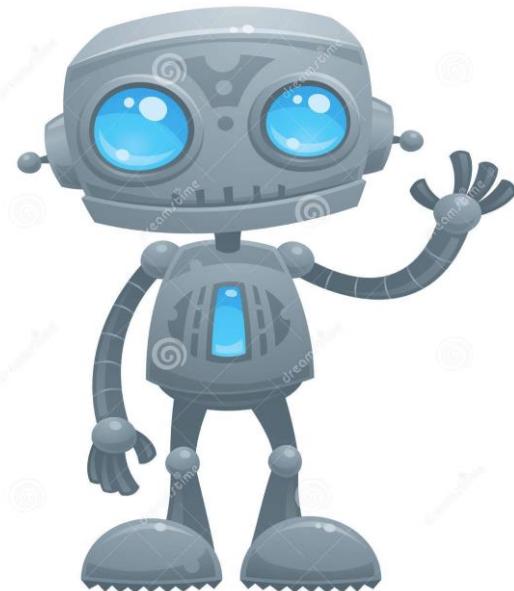
**Q: In this image, how many eyes can you see on the animal?**



**Q: How many wheels can you see in the image?**

# What is Computer Vision?

- First defined in the 60s in artificial intelligence groups
- "Mimic the human visual system"
- Center block of robotic intelligence



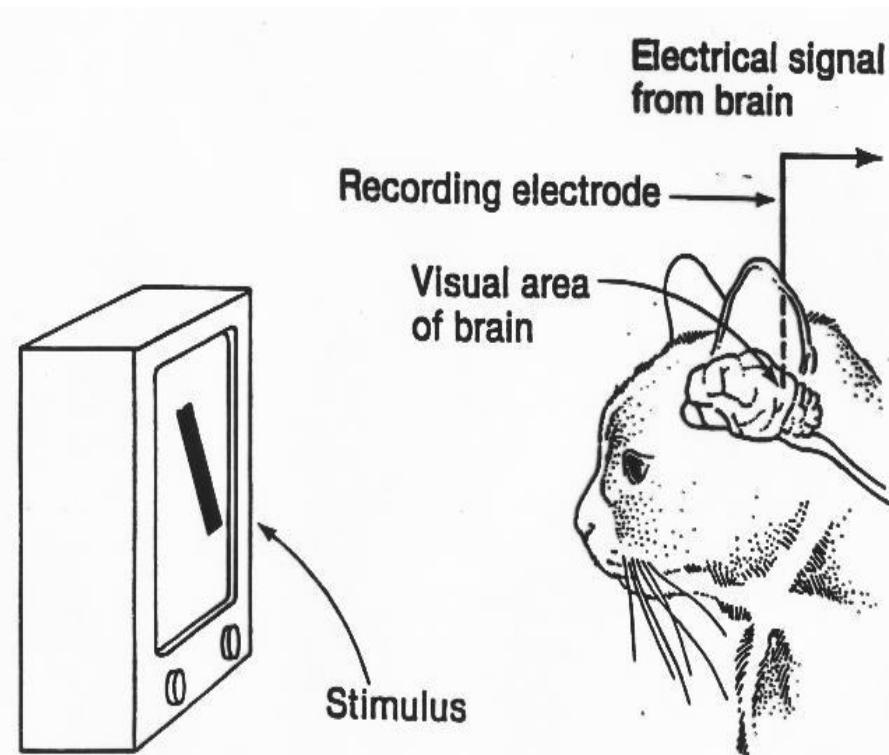
# Hubel and Wiesel

- David Hubel and Torsten Wiesel were neurobiologists from Harvard Medical School
- Experiment revealed insights into the primate visual system
- Nobel prize 1981



# Hubel and Wiesel Experiment

- Recorded electrical activity from individual neurons in the brains of cats.
- Slide projector to show specific patterns to the cats noted specific patterns stimulated activity in specific parts of the brain.
- Results: Visual cortex cells are sensitive to the orientation of edges but insensitive to their position



Artificial Intelligence Group  
Vision Memo. No. 100.

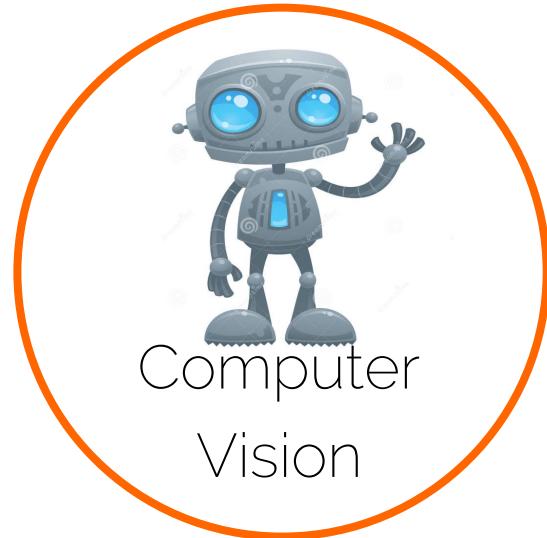
July 7, 1966

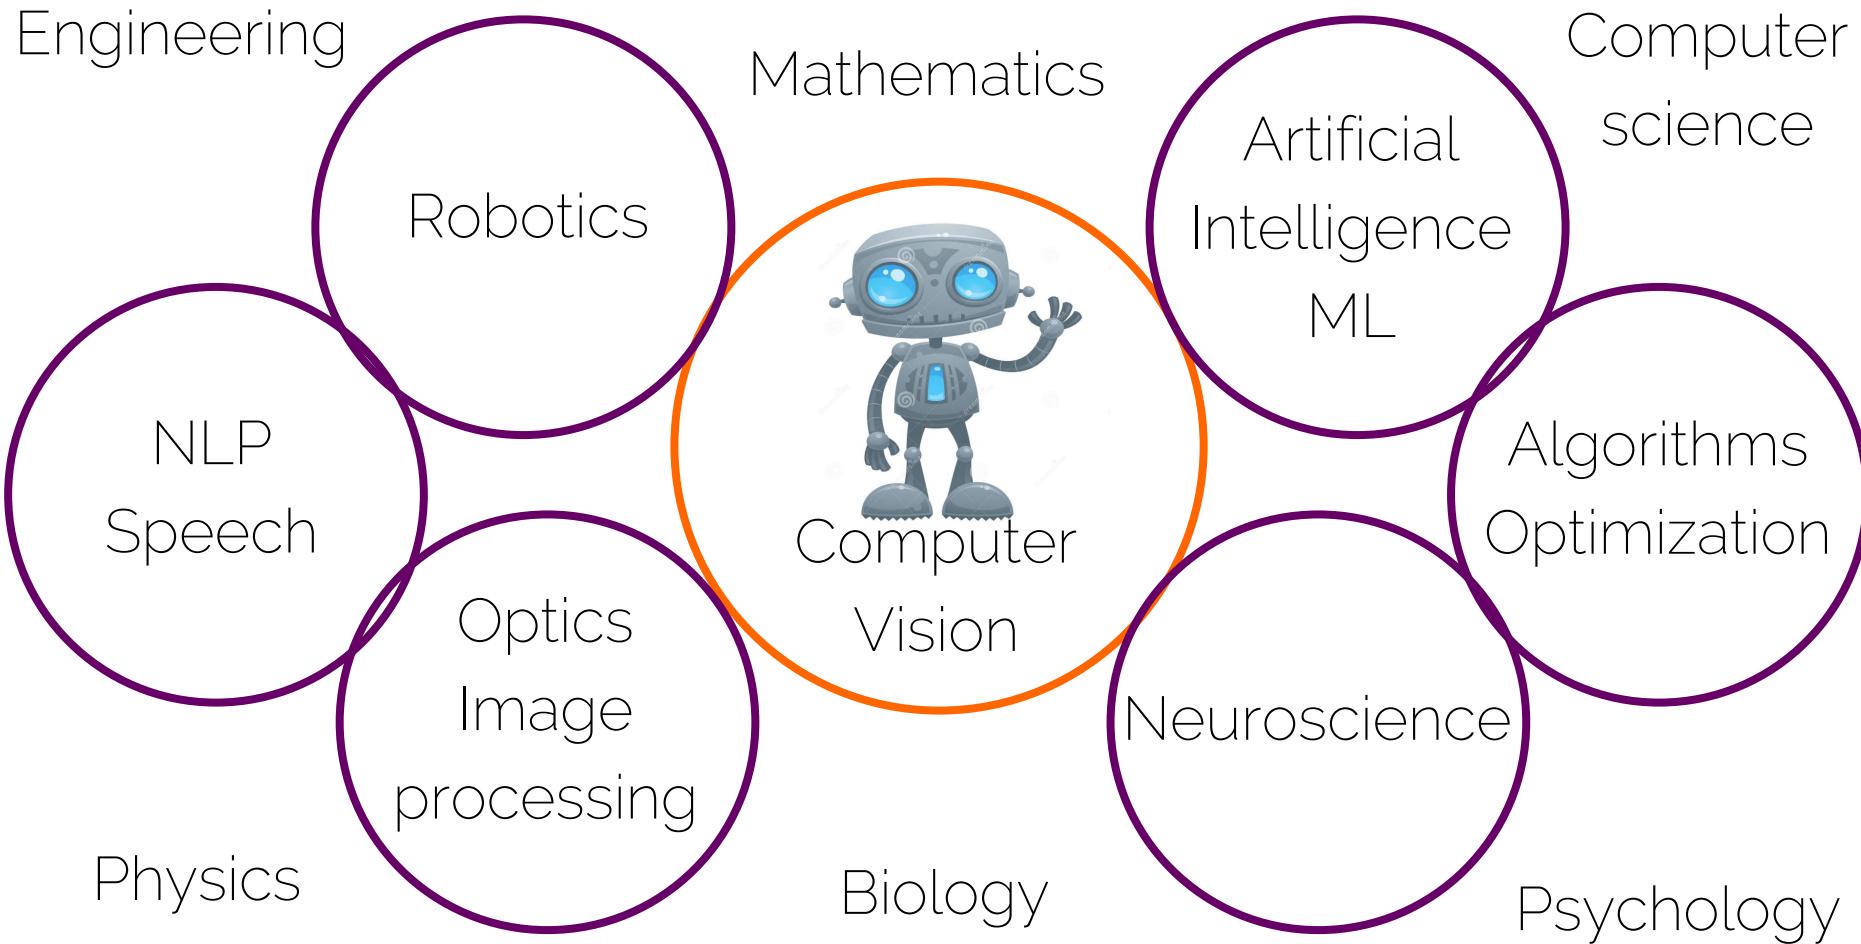
THE SUMMER VISION PROJECT

Seymour Papert

The summer vision project is an attempt to use our summer workers  
effectively in the construction of a significant part of a visual system.  
The particular task was chosen partly because it can be segmented into  
sub-problems which will allow individuals to work independently and yet  
participate in the construction of a system complex enough to be a real  
landmark in the development of "pattern recognition".

# A Few Decades Later...





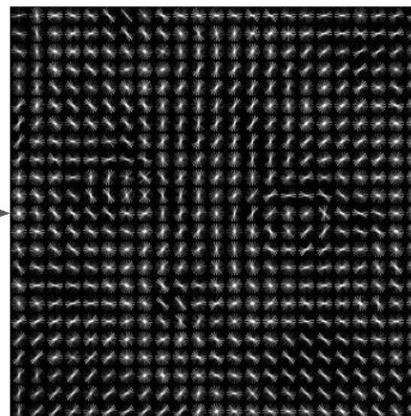
# Image Classification



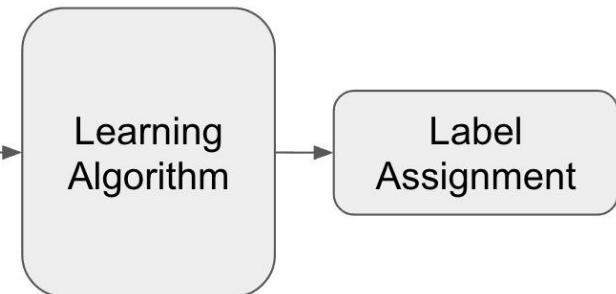
Input image



Preprocessing



Features : HAAR, HOG,  
SIFT, SURF



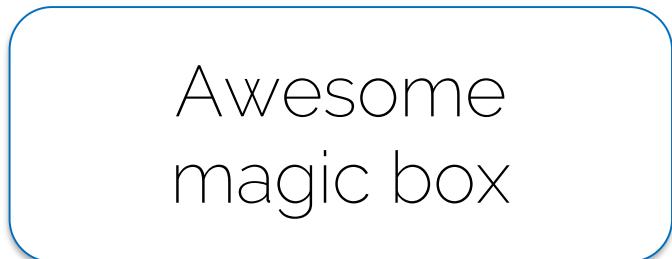
SVM,  
Random  
Forests,  
ANN

Cat or  
Background

# Image Classification



Input image



Open the box

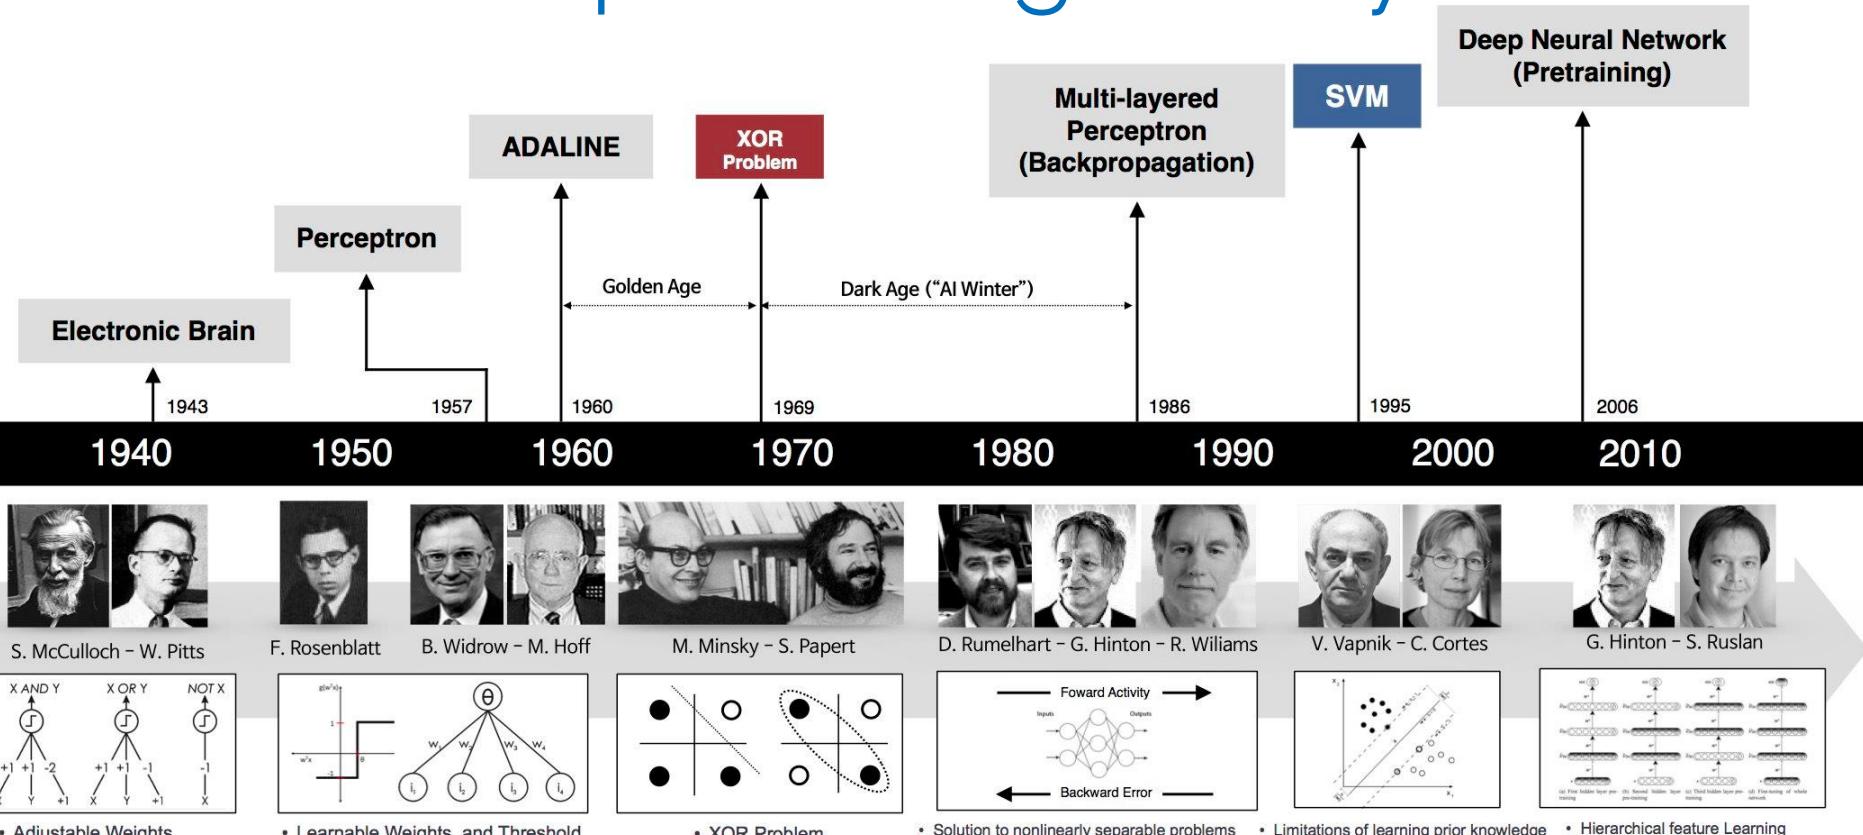


Become magicians

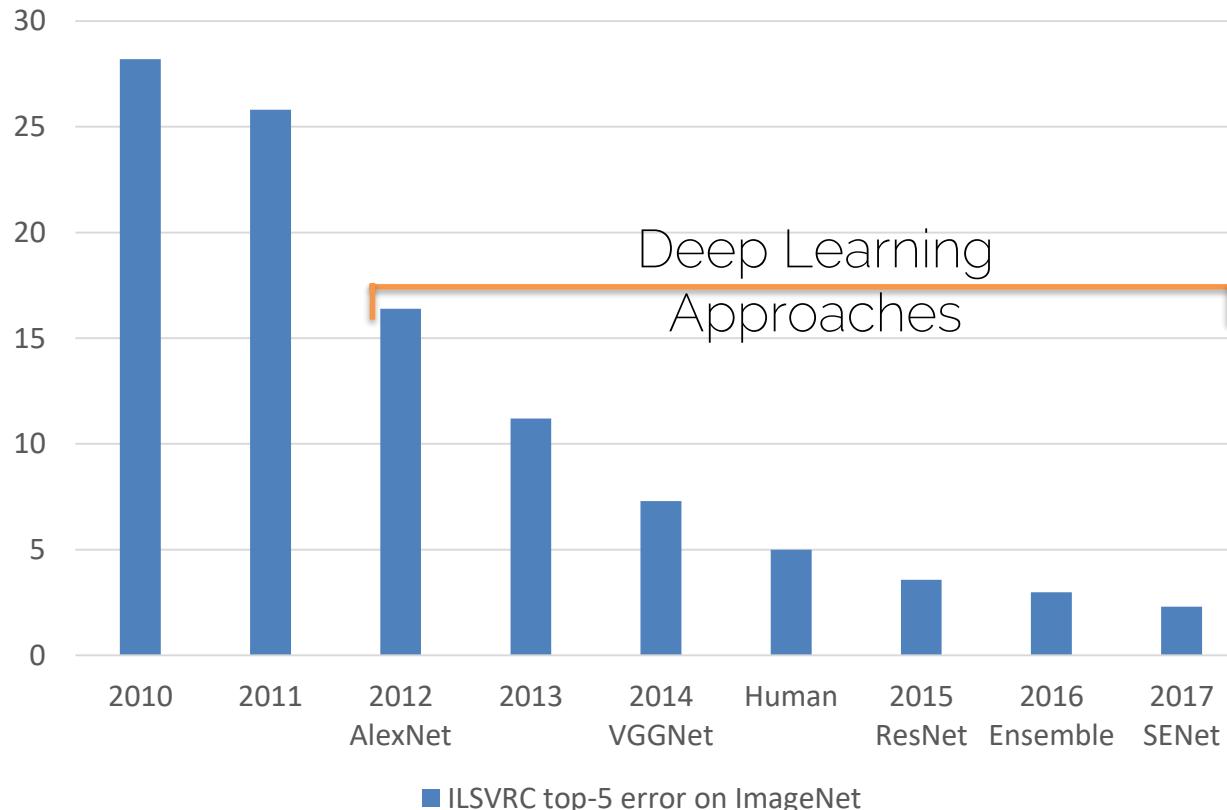
Cat or  
Background

# Why Deep Learning?

# Deep Learning History

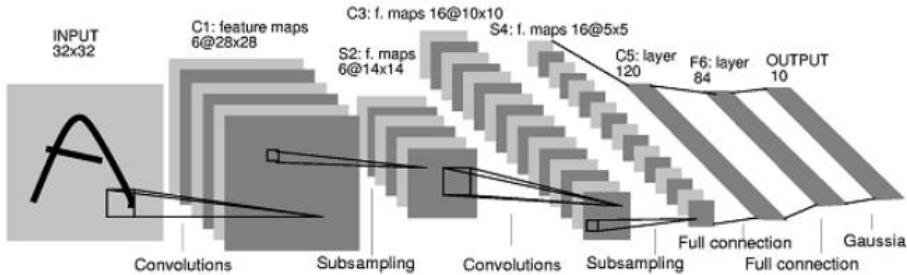


# The Empire strikes Back



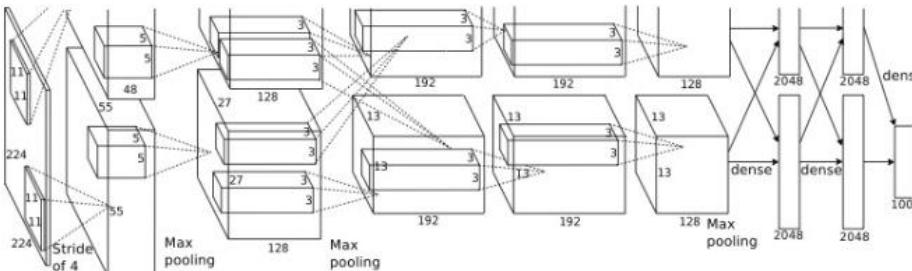
# What has Changed?

1998  
LeCun  
et al.



- MNIST digit recognition dataset
- $10^7$  pixels used in training

2012  
Krizhevsky  
et al.



- ImageNet image recognition dataset
- $10^{14}$  pixels used in training

# What Made this Possible?



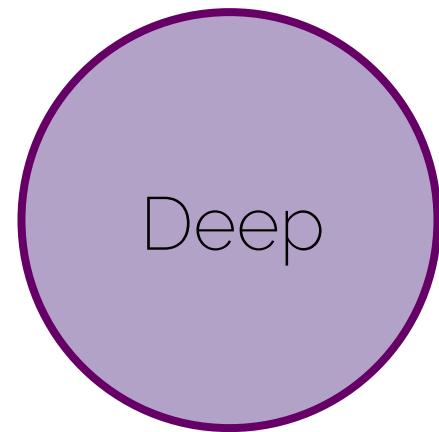
Big Data

Models know  
where to learn from



Hardware

Models are  
trainable



Deep

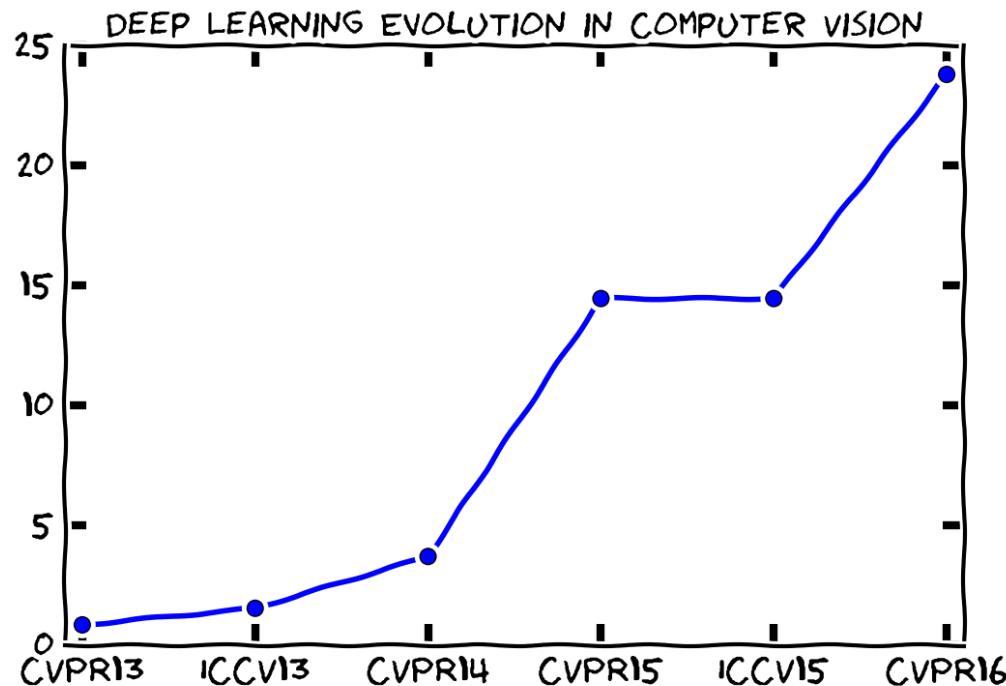
Models are  
complex

# Deep Learning Recognition



ACM Turing Award 2019 (Nobel Prize of Computing)  
Yann LeCun, Geoffrey Hinton, and Yoshua Bengio

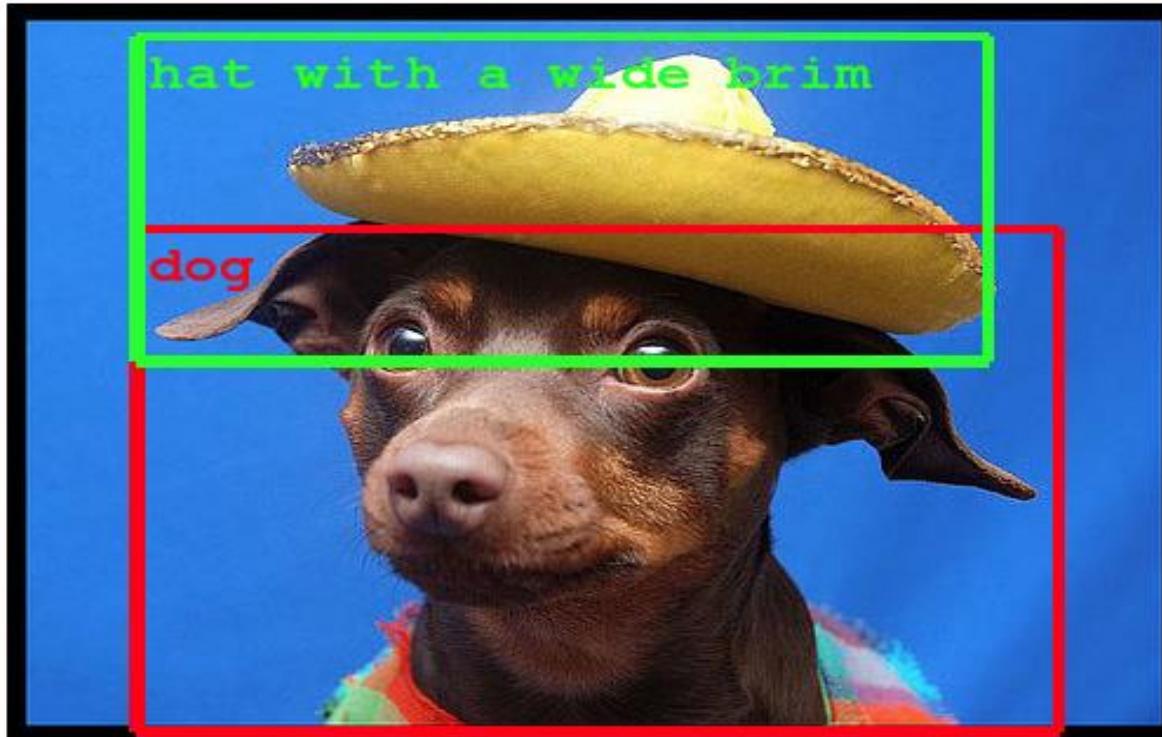
# Deep Learning and Computer Vision



There are now >2,000 papers at almost every top-tier AI conference  
(CVPR/ICCV/ECCV NeurIPS/ICLR/ICML)

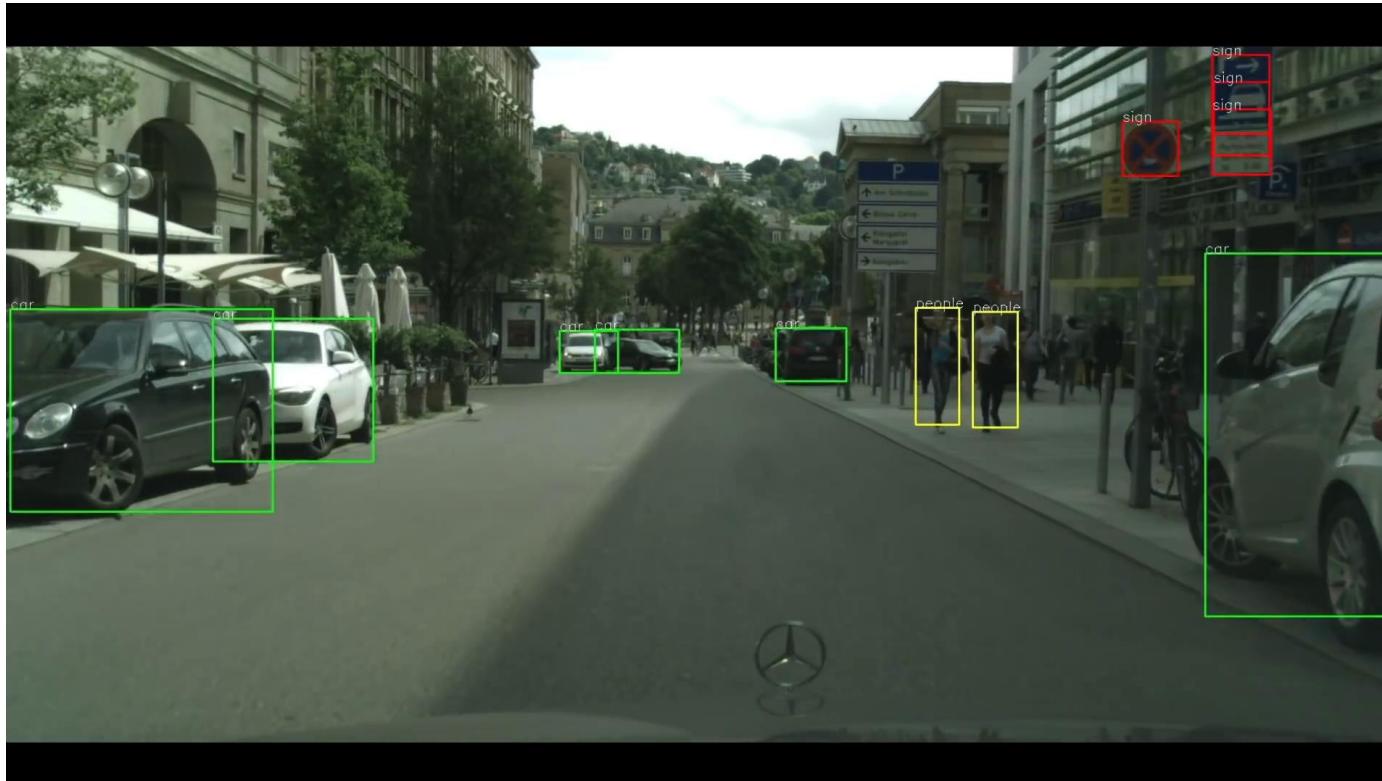
Credits: Dr. Pont-Tuset, ETH Zurich

# Deep Learning Today



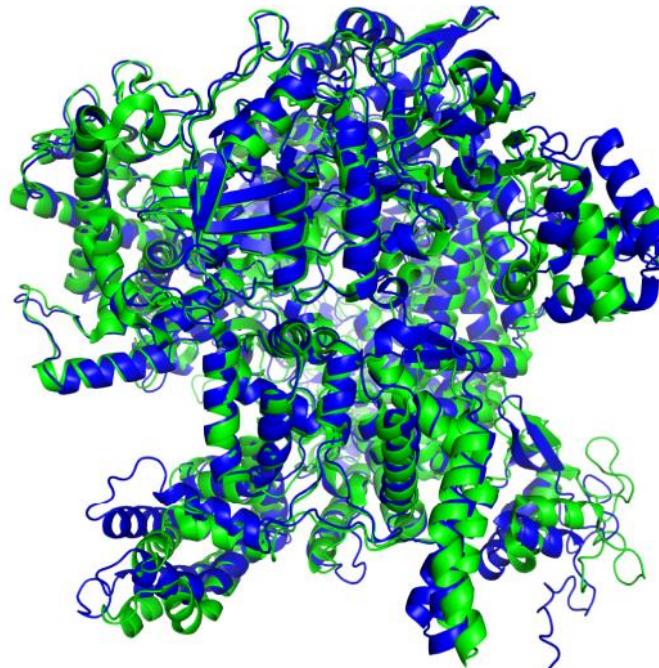
Object Detection

# Deep Learning Today



Self-driving cars

# Deep Learning Today



AlphaFold   Experiment

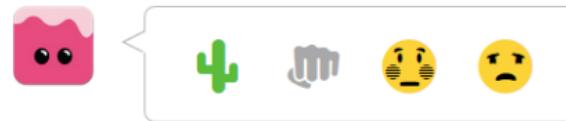
Biological research

# Deep Learning Today



AlphaGo

ever punch a cactus?

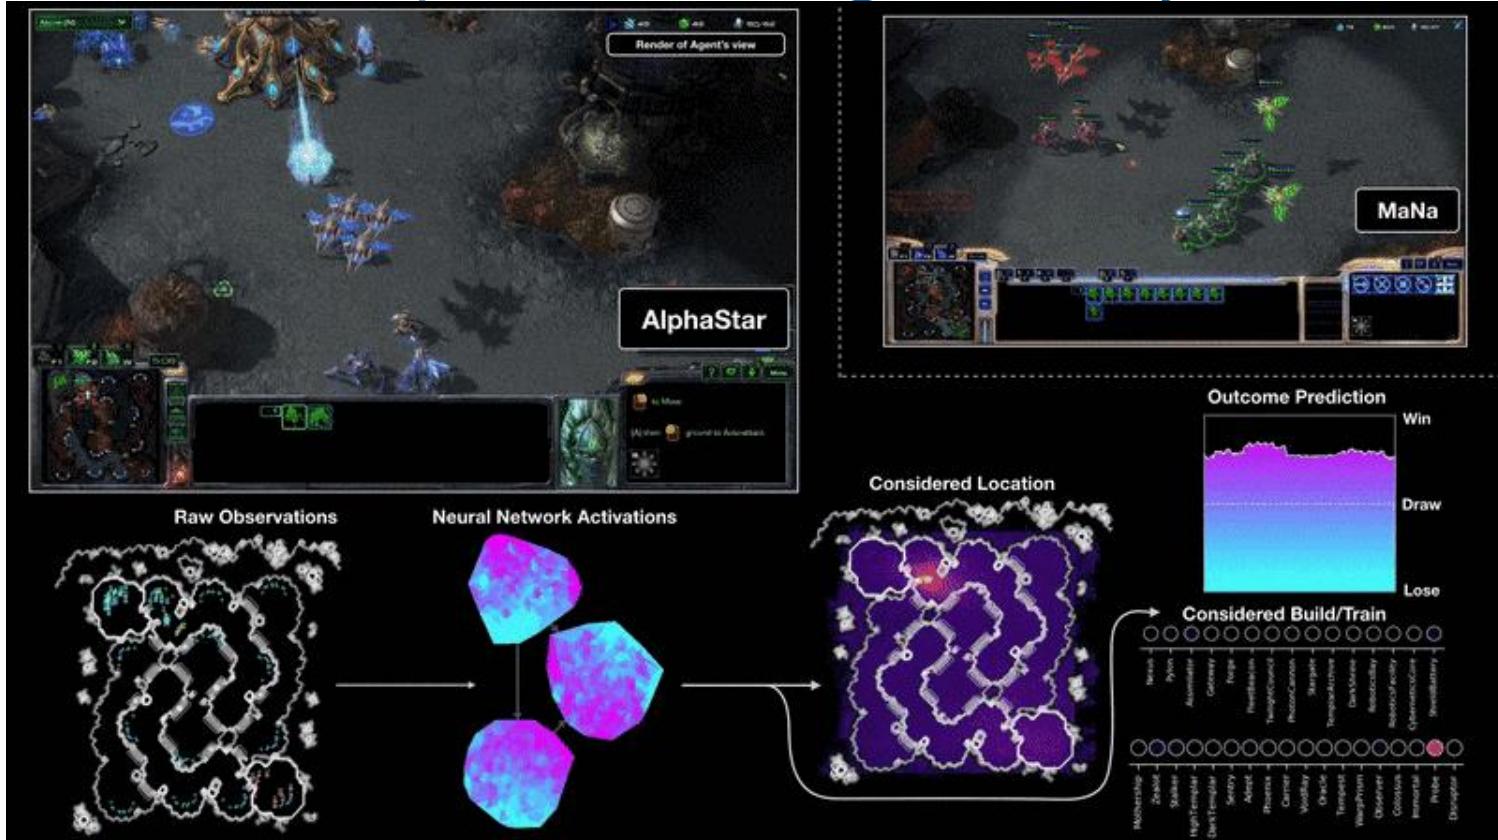


Emoticon suggestion

A screenshot of a machine translation interface. The English input is "Deep Learning rocks" and the German output is "Deep Learning Felsen". The interface includes language selection dropdowns and a speech icon.

Machine translation

# Deep Learning Today



Alpha Star

# Deep Learning Today



Google LaMDA (Google IO'22)

# Deep Learning Today

- Large Language Models (LLMs), e.g., Chat-GPT

NI

Tell me about deep learning and how I can teach it at TUM.



Deep learning is a subset of machine learning that involves training artificial neural networks to recognize patterns and make predictions or decisions based on input data. The term "deep" refers to the fact that these networks have many layers, which allows them to learn increasingly complex representations of the data as it moves through the network.



To teach deep learning at TUM, you could start with an introductory course that covers the basics of machine learning and neural networks. This could include topics such as:

1. Supervised learning: This involves training a model using labeled examples, so that it can make predictions on new, unseen data.

# Deep Learning Today

## TEXT DESCRIPTION

An astronaut **Teddy bears** A bowl of soup

mixing sparkling chemicals as mad scientists **shopping for groceries** working on new AI research

**in the style of ukiyo-e** as a one-line drawing in ancient Egypt



## DALL-E 2



Dall-E 2 Image generation from text

# Deep Learning Today

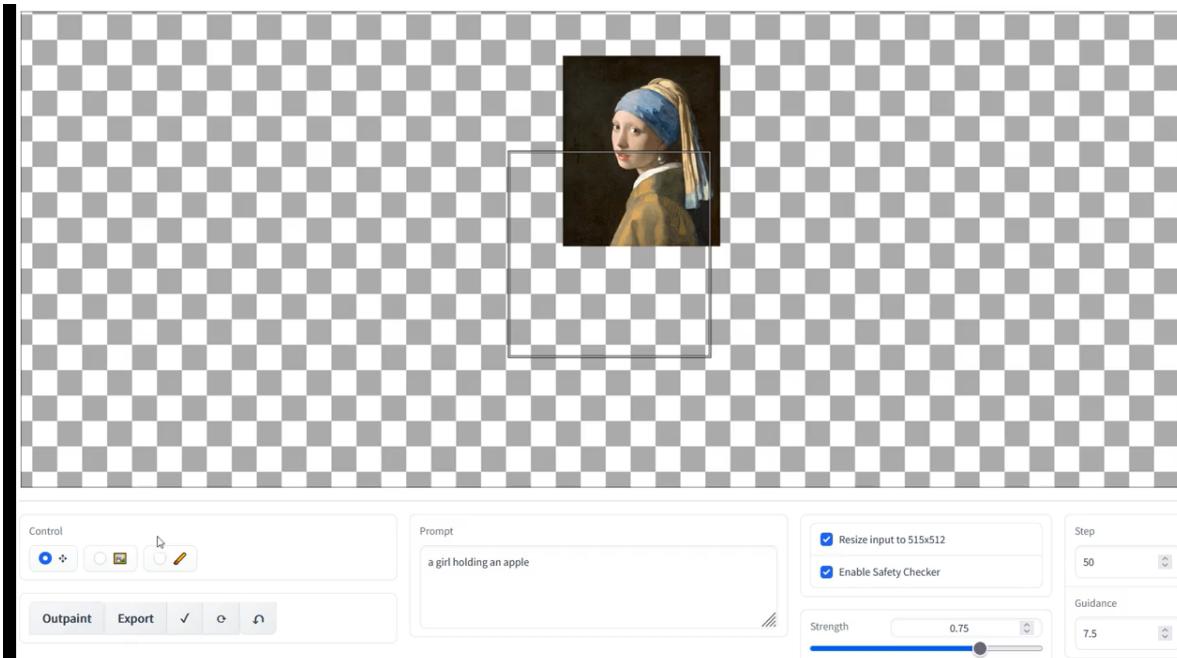


*“a robot painting a colorful sunset on a canvas in a futuristic art studio”*



*“a majestic dragon made of water flying through a colorful, abstract city”*

# Deep Learning Today



StableDiffusion Image Outpainting

# Deep Learning Today



# Deep Learning Today



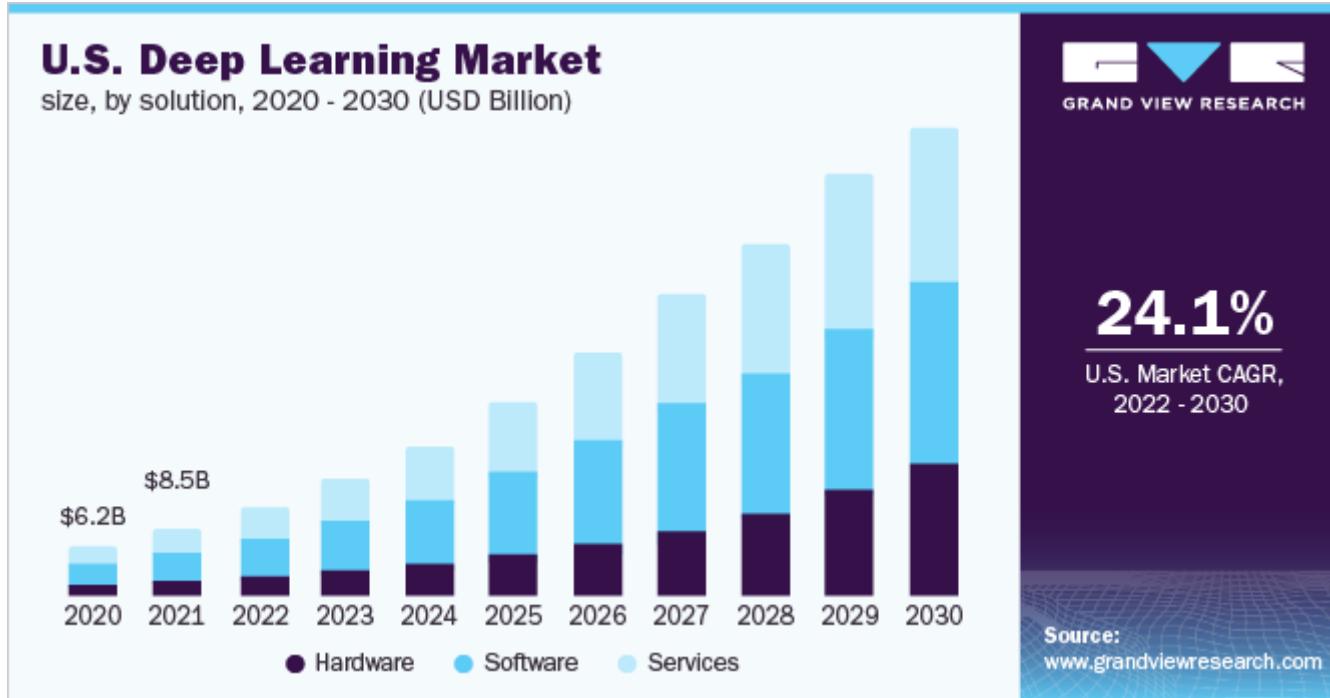
# Deep Learning Today



# Deep Learning Today



# Deep Learning Market



[...] market research report Deep Learning Market [...] " the deep learning market is expected to be worth **USD 415 Billion by 2030**.

# Deep Learning Job Perspective

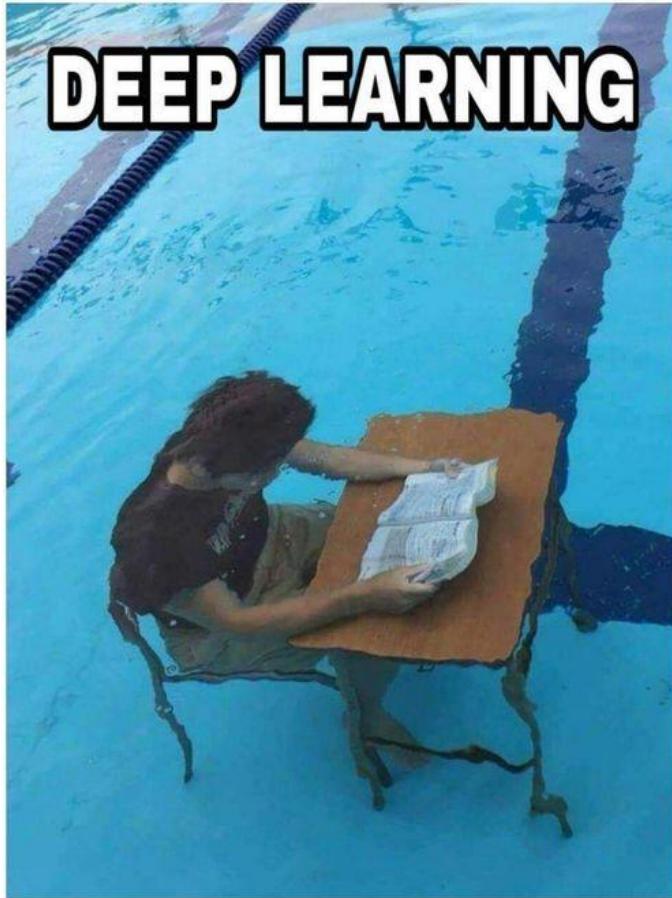
- Excellent Job Perspectives!
  - Automation requires ML/DL -> growth!
  - Top-notch companies will gladly hire you!
- Many industries now:
  - IT-Companies
  - Cars, Logistics, Health Care, etc...
  - Manufacturing / Robotics, etc...

# But: Also Challenging!

- High-level understanding is not enough
  - Need proper theory background
  - Need proper practical skillsets
- Can be competitive!
  - Many good people
  - Education gap widens
  - Downloading scripts / running code not enough ☺
  - Deeper understanding often requires PhDs

# Deep Learning on the Internet

# Deep Learning Memes



# Deep Learning Memes

## Deep Learning



What society thinks I do



What my friends think I do



What other computer  
scientists think I do



What mathematicians think I do

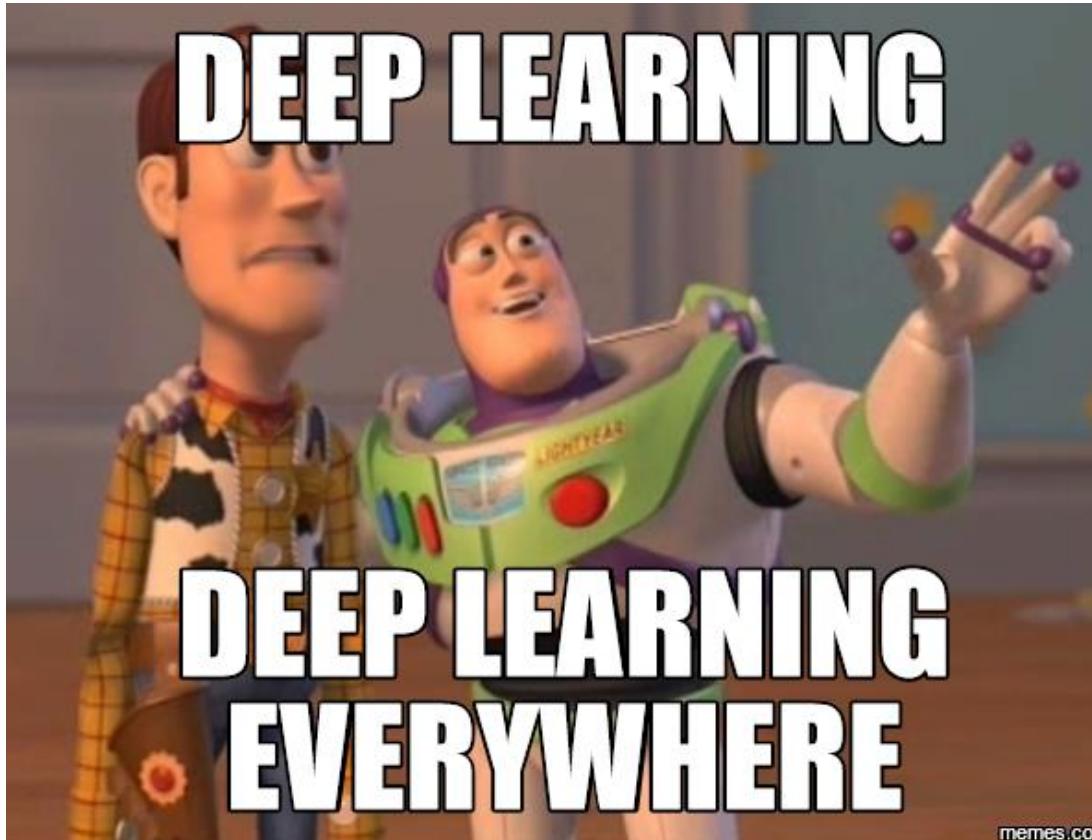


What I think I do

from theano import \*

What I actually do

# Deep Learning Memes



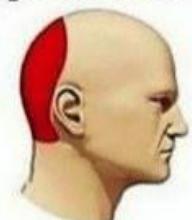
# Deep Learning Memes

## Types of Headaches

**Migraine**



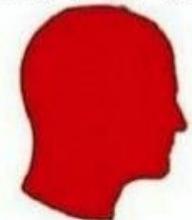
**Hypertension**



**Stress**



**MATH BEHIND DL**



# Deep Learning at TUM

# Many TUM Research Labs use DL

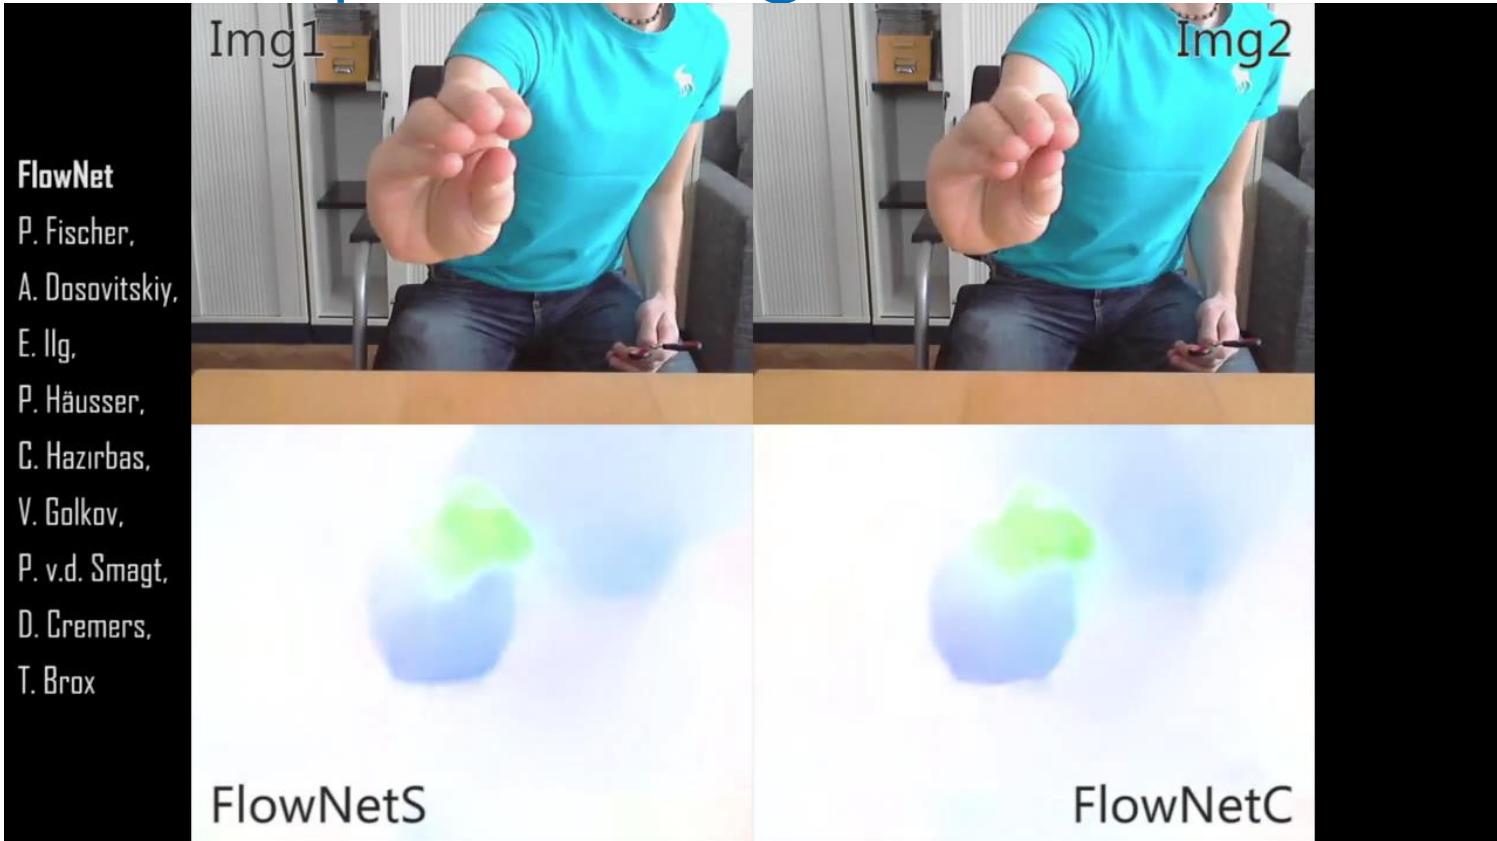
- 3D AI Lab (Prof. Dai)
  - Research in 3D perception, 3D scene understanding
- Visual Computing & AI (Prof. Niessner):
  - Research in computer vision, graphics, and machine learning
- Computer Vision Group (Prof. Cremers)
  - Research in computer vision, machine learning and robotics
- Data Mining and Analytics Lab (Prof. Günnemann)
  - Research methods for robust machine learning
- Computer Aided Medical Procedures (Prof. Navab)
  - Research in machine learning for medical applications
- And many more ☺

# Deep Learning at TUM



[Caelles et al., CVPR' 17] One-Shot Video Object Segmentation

# Deep Learning at TUM



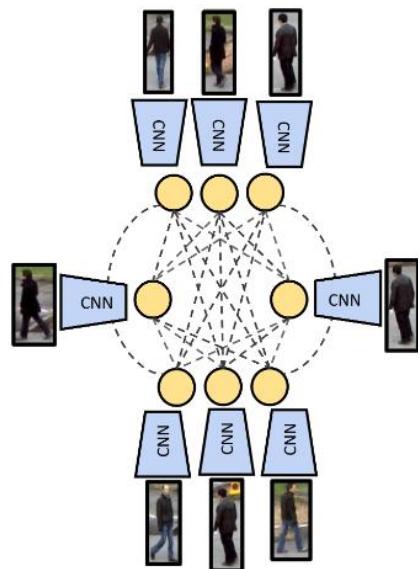
[Dosovitskiy et al., ICCV' 15] FlowNet

# Deep Learning at TUM

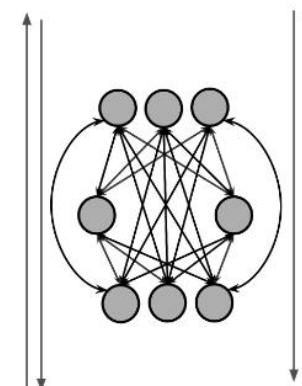
- Multiple object tracking with graph neural networks



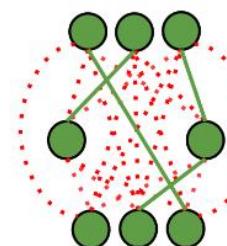
(a) Input



(b) Graph Construction + Feature Encoding



(c) Neural Message Passing



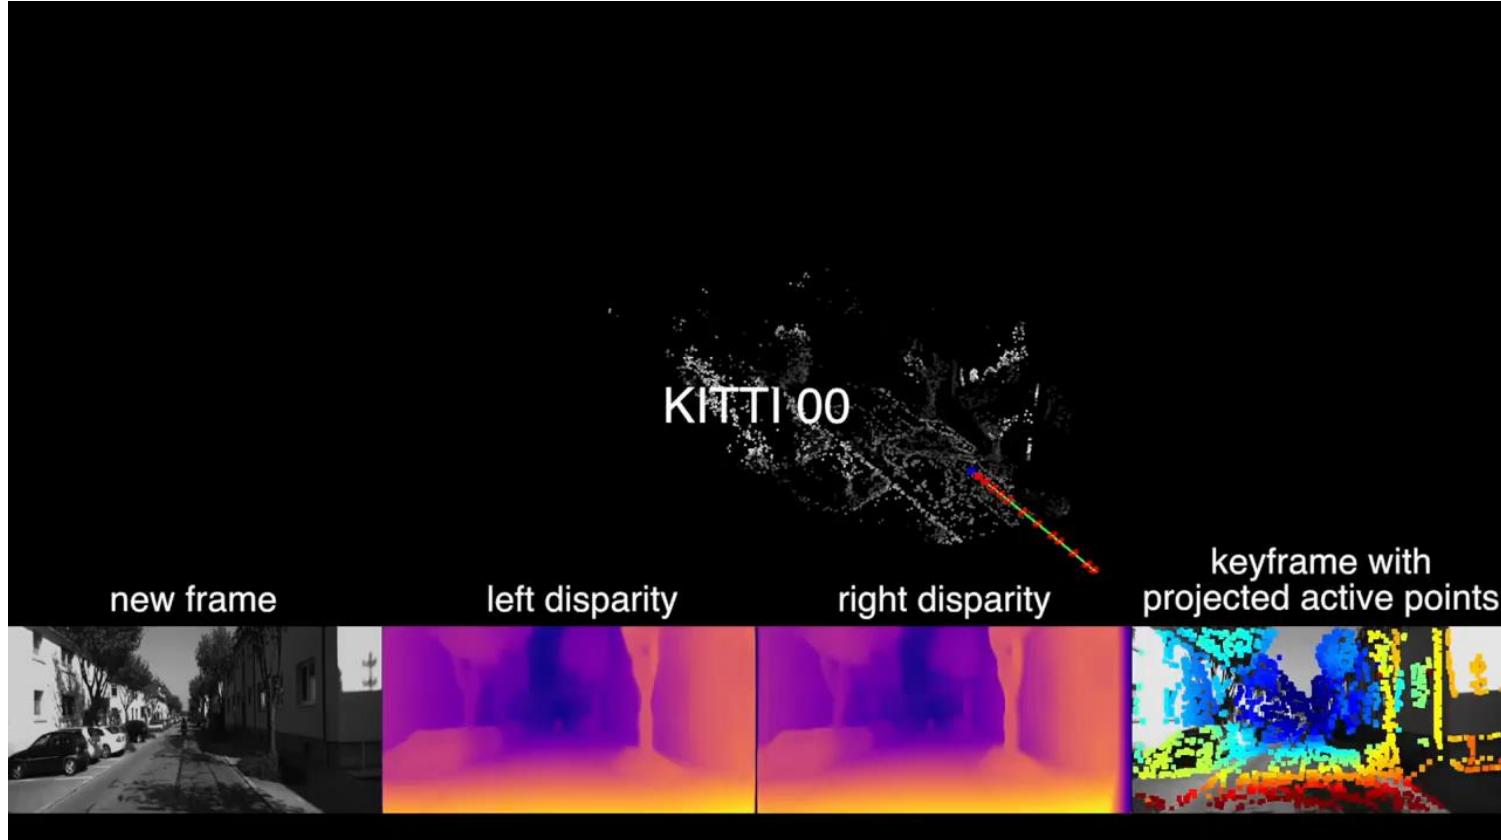
(d) Edge Classification



(e) Output

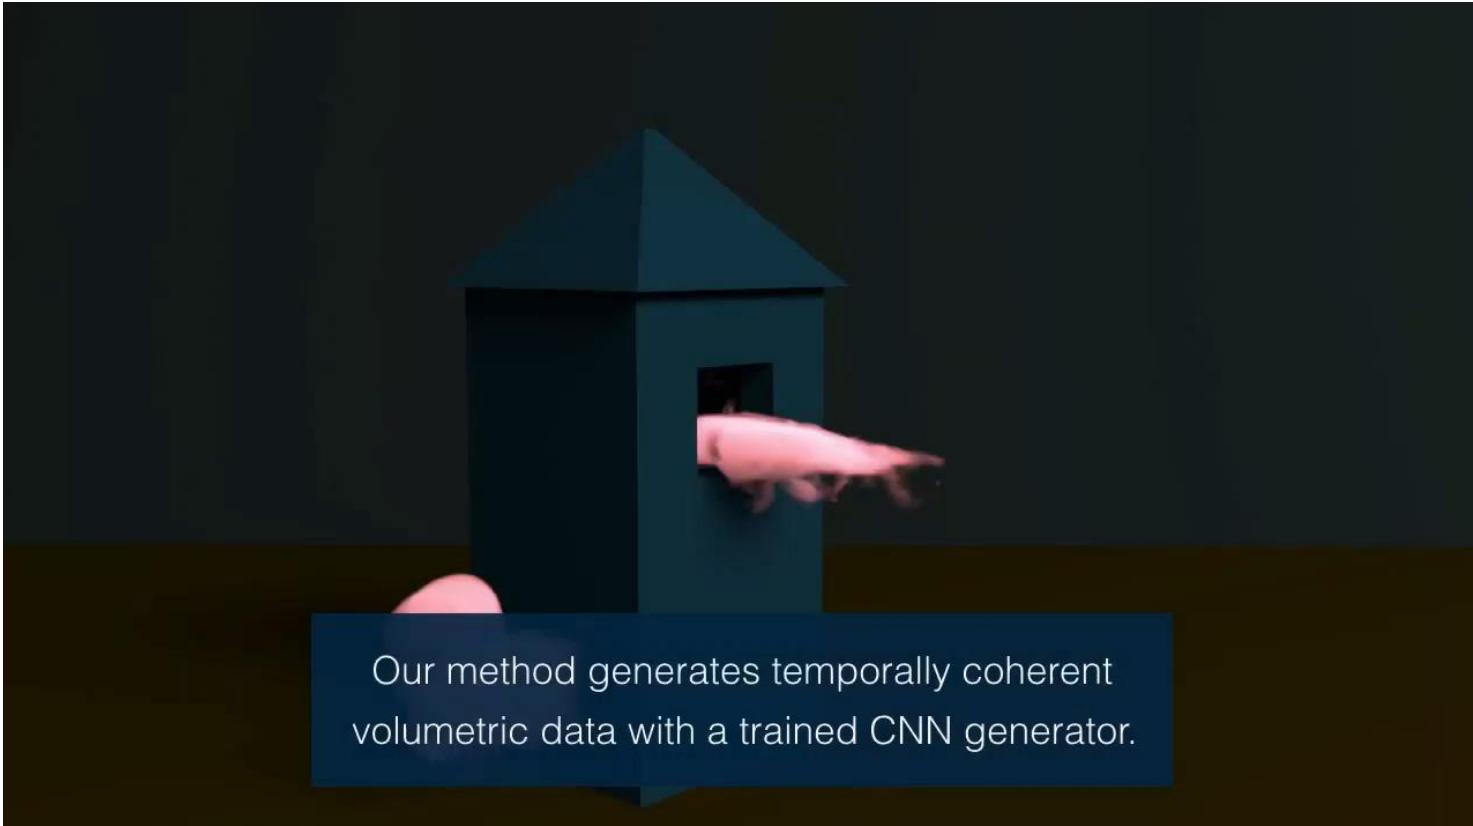
[Brasó and Leal-Taixé, CVPR 2020] Learning a Neural Solver for Multiple Object Tracking.

# Deep Learning at TUM



[Yang et al., ECCV' 18] Deep Virtual Stereo Odometry

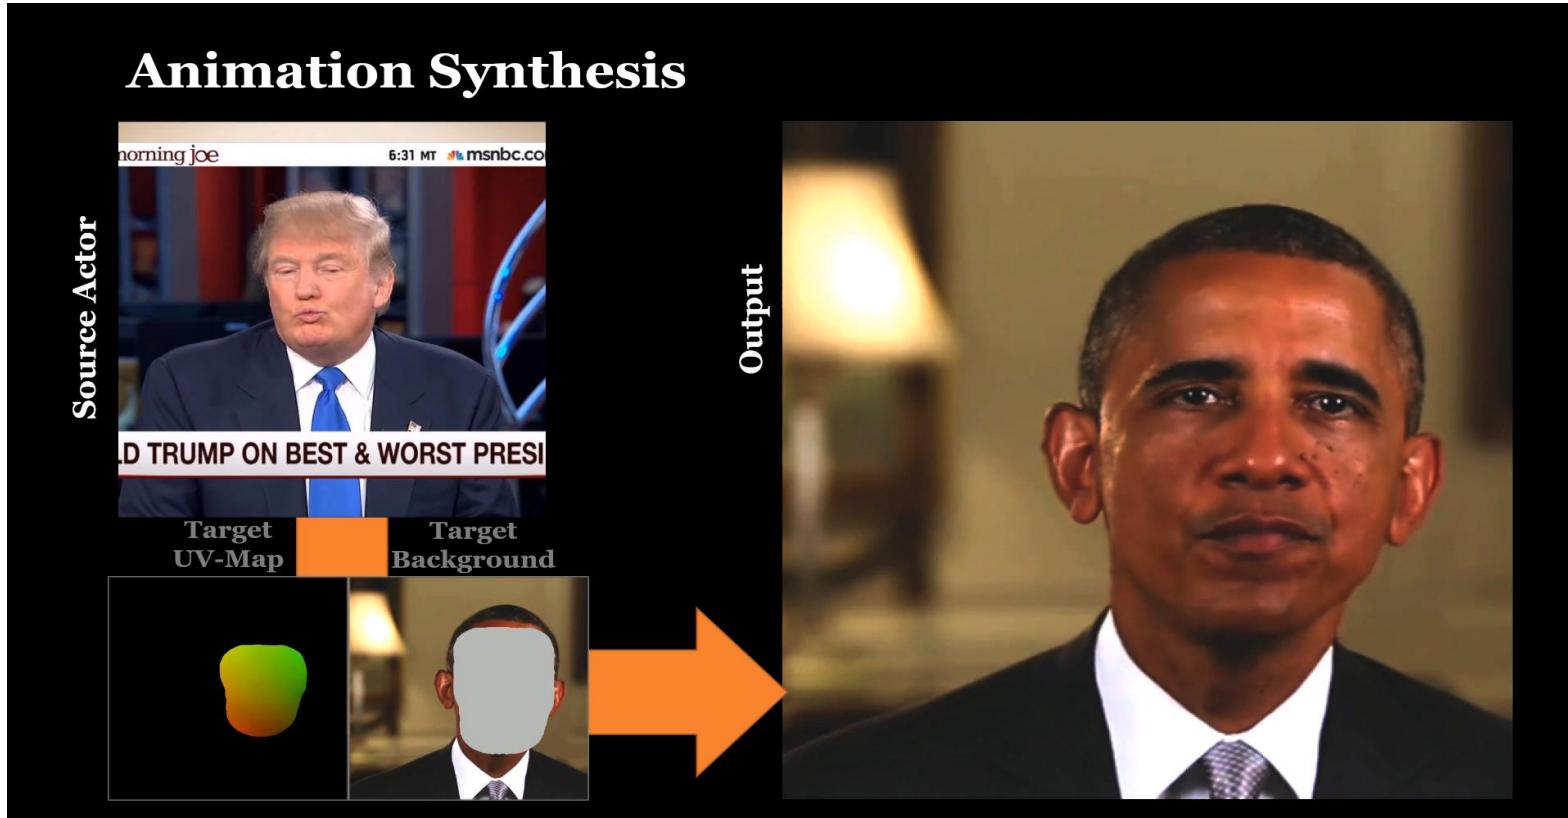
# Deep Learning at TUM



[Xie et al. Siggraph' 18] tempoGAN

# Deep Learning at TUM

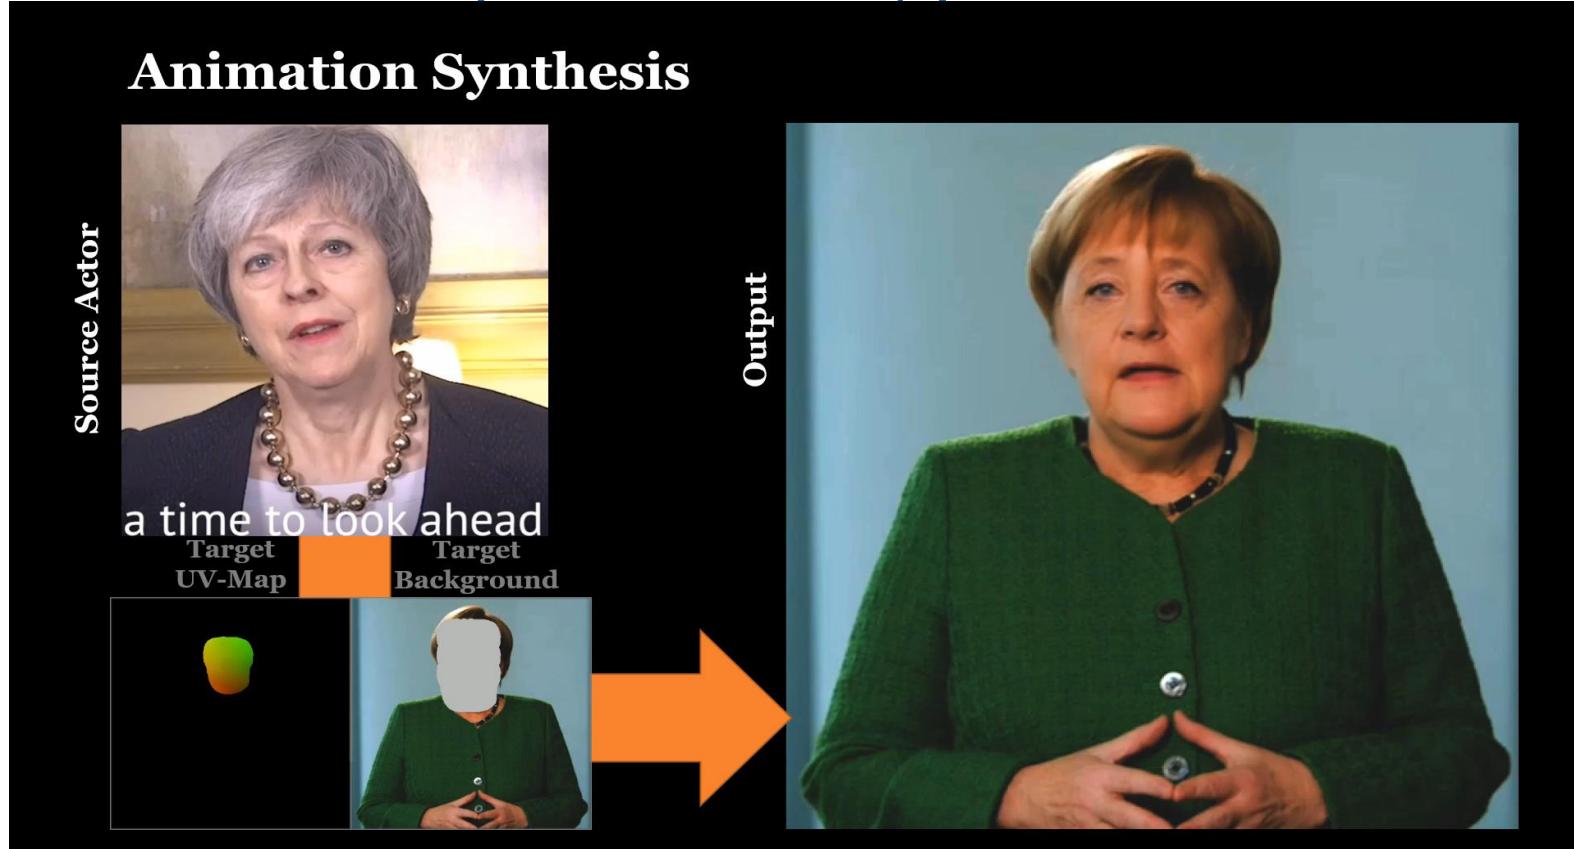
## Animation Synthesis



[Thies et al., Siggraph'19]: Neural Textures

# Deep Learning at TUM

## Animation Synthesis



[Thies et al., Siggraph'19]: Neural Textures

# Deep Learning at TUM



[Kirchstein et al., ToG'23] Nersempble

# Deep Learning at TUM



ScanNet Stats:

- Kinect-style RGB-D sensors
- 1513 scans of 3D environments
- 2.5 Mio RGB-D frames
- Dense 3D, crowd-source MTurk labels
- Annotations projected to 2D frames

# Deep Learning at TUM



[Yeshwanth & Liu et al., ICCV'23] ScanNet++

ScanNet++:  
- 1006 high-fidelity  
Scenes  
- 1mm laser scans  
- DSLR images  
- iPhone RGB-D  
- semantics

# Deep Learning at TUM

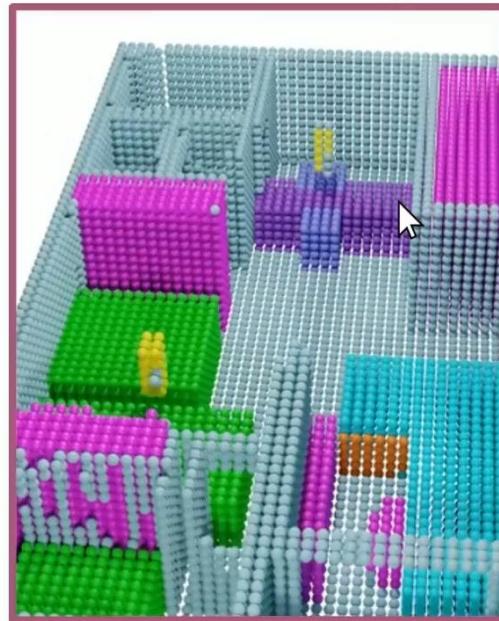


[Bokhovkin et al., CVPR'25] SceneFactor

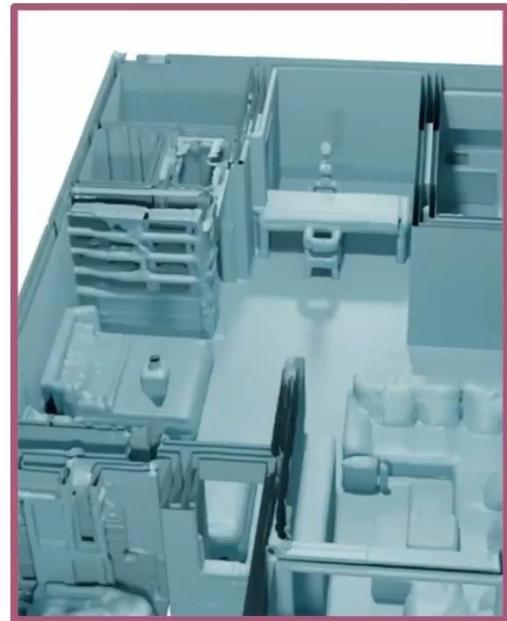
# Deep Learning at TUM



3D Semantic Map

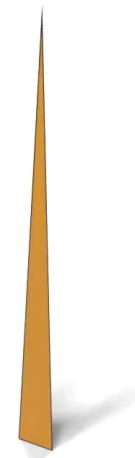


Scene Geometry



[Bokhovkin et al., CVPR'25] SceneFactor

# Deep Learning at TUM



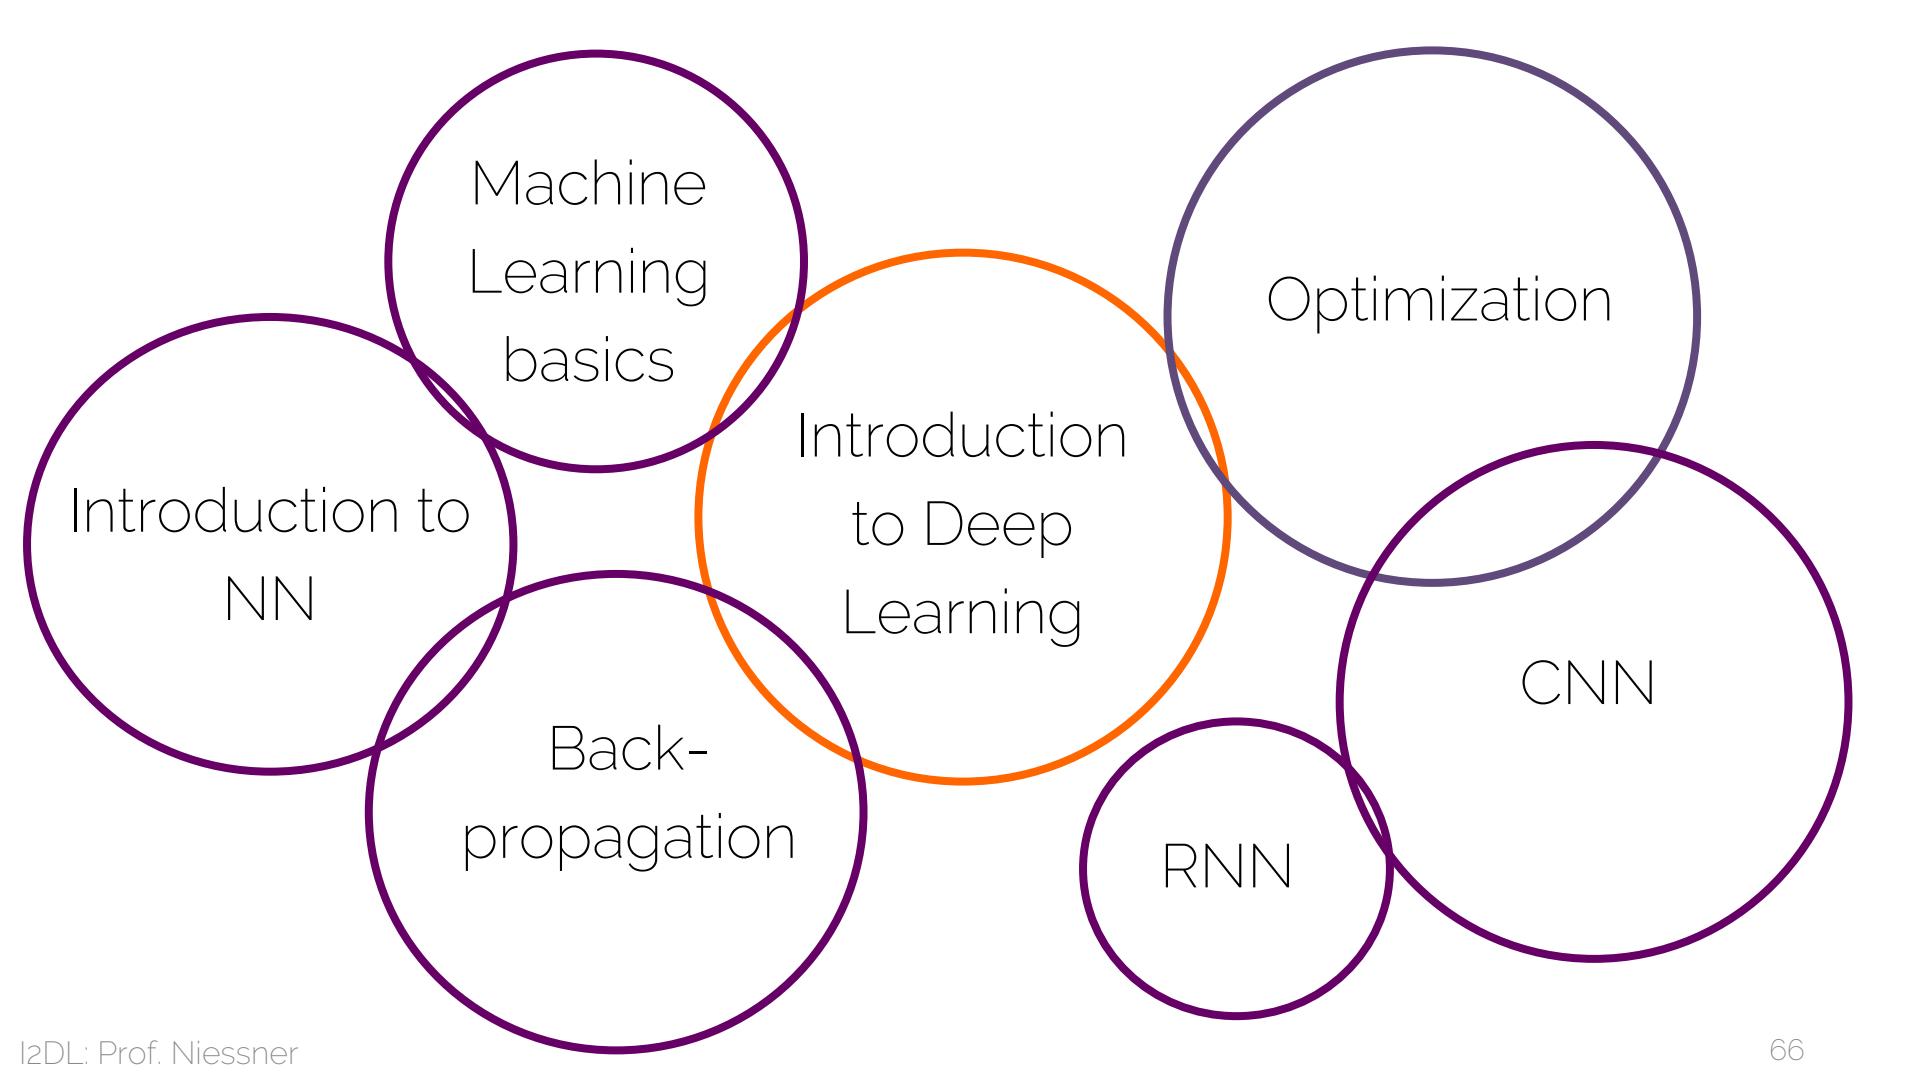
[Siddiqui et al., CVPR'24] MeshGPT

# Deep Learning at TUM

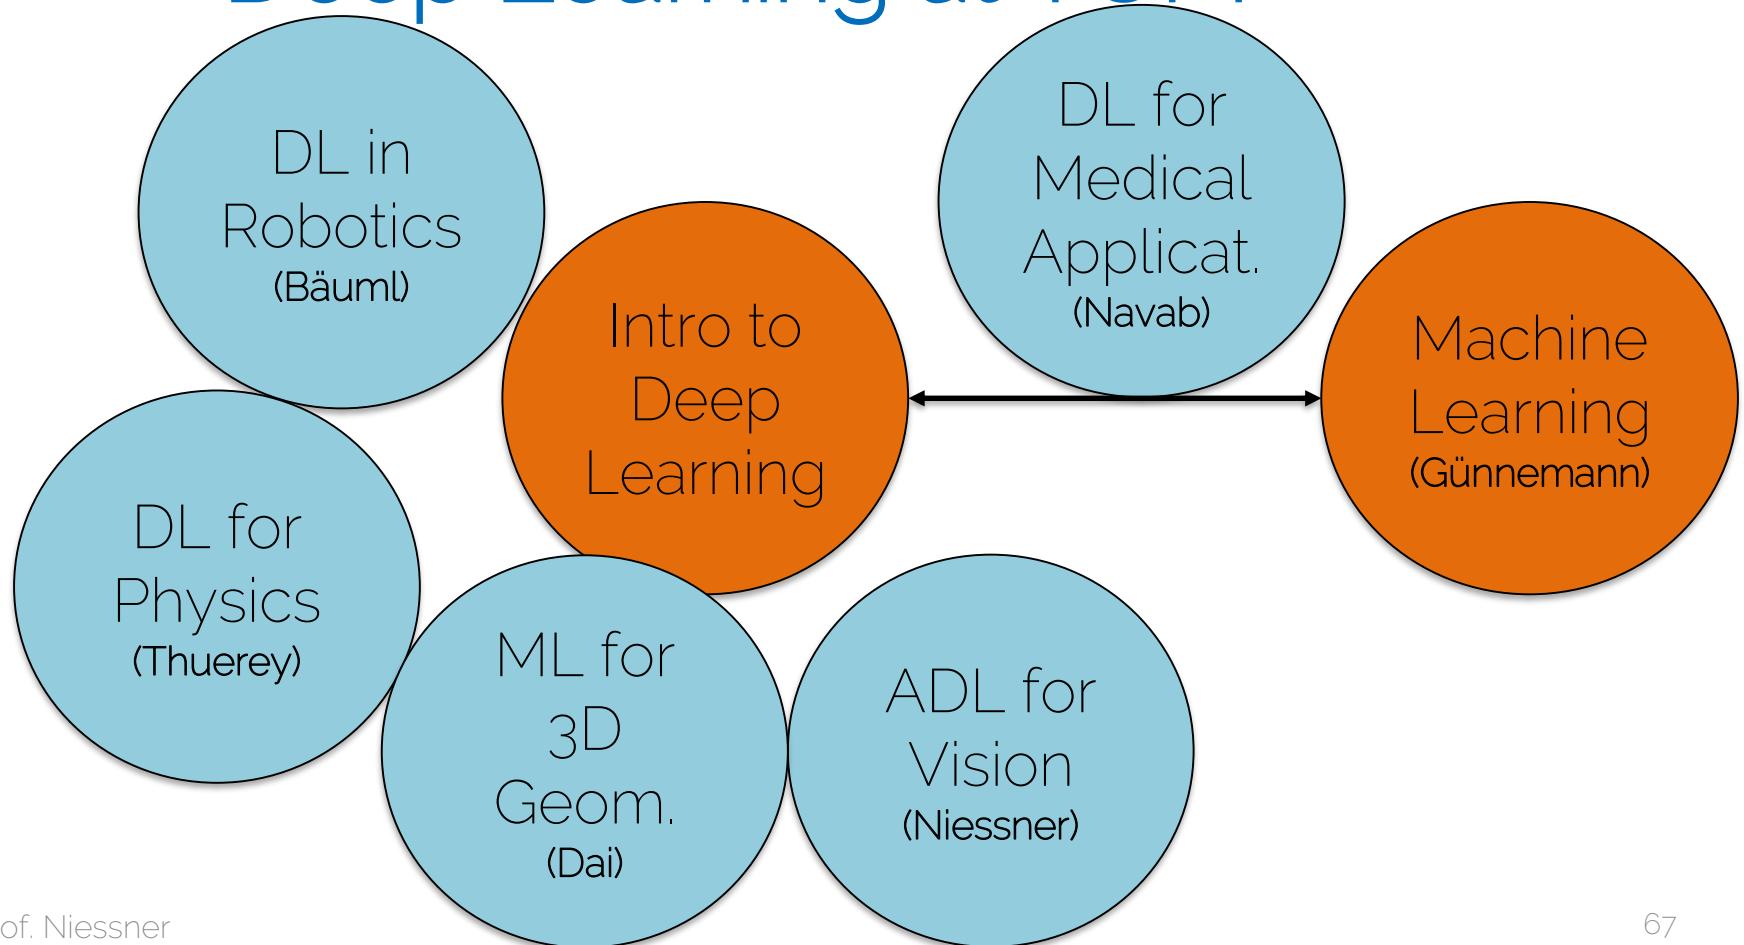


[Schneider et al., SIGGRAPH Asia'25] WorldExplorer

# Relation to other lectures at TUM



# Deep Learning at TUM



# Deep Learning at TUM

- I2DL is critical for the curriculum:
  - often a pre-requisite for many other lectures
  - many lectures build on top of it
- Follow up lectures such as Advanced Deep Learning for Visual Computing (generative models etc.)
- Take I2DL first -> AdvancedDL -> practical/GR/IDP

# Introduction to Deep Learning

Logistics

# About the Lecture

- Course webpage: <https://niessner.github.io/I2DL/>
- Theory lectures (every Tuesday at 14:15)
  - In-person, live-streamed
- Tutorials and exercises (every Thursday at 10:00)
  - Tutorial: Online videos posted to Piazza and the webpage
  - Practical exercises
- Guest Lecture ☺

# Preliminary Syllabus

Lecture 1: Introduction to the lecture, Deep Learning, Machine Learning.

Lecture 2: Machine Learning Basics, Linear regression, Maximum Likelihood

Lecture 3: Introduction to Neural Networks, Computational Graphs

Lecture 4: Optimization and Backpropagation

Lecture 5: Scaling Optimization to large Data, Stochastic Gradient Descent

Lecture 6: Training Neural Networks I

Lecture 7: Training Neural Networks II

Lecture 8: Training Neural Networks III

Lecture 9: Introduction to CNNs

Lecture 10: CNNs architectures;

Lecture 11: Recurrent Neural Networks (RNNs)

Lecture 12: Advanced Deep Learning architectures

# Moodle → Piazza

- Announcements via Piazza - **IMPORTANT!**
  - Sign up online for access: <http://piazza.com/tum.de>
    - Select "Winter 2026" term, search for IN2346
    - Use your @mytum.de email address
  - We will share common information (e.g., regarding exam)
- Forum
  - Ask and discuss questions
    - Tutors will monitor and answer questions
    - You are very welcome to actively participate
  - Please do not post solutions of the exercises
  - You can post private questions visible only to the staff



# Email

- Email list:

i2dl@vc.in.tum.de

- Do NOT email us personally!
  - Cannot handle so many emails / hence will be ignored
- Email list for organizational questions only!
  - Content questions -> Piazza or Office Hours
  - Or post the question/issue in a private thread on Piazza

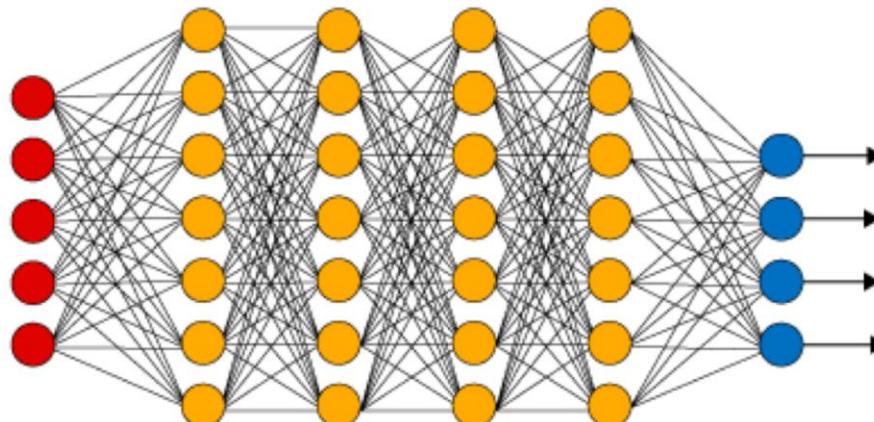
# Website



Visual Computing Group  
Prof. Matthias Nießner

[Home](#) [Publications](#) [Teaching](#) [Openings](#) [Team](#)

## Introduction to Deep Learning (I2DL) (IN2346)



<https://niessner.github.io/I2DL/>

# (Virtual) Office Hours

- We will have dedicated office hours regarding
  - Theoretical help (e.g., specific lecture questions)
  - Help on exercises
- More info in the first tutorial session
- Zoom links will be posted on Piazza

# Exam FAQ

- Final Exam: TBA
- Content: Lecture & exercises
- Important: No retake exam (I2DL is taught every semester)
- Grade Bonus:
  - Solve 8 out of 9 “non-optional” programming exercises
  - Bonus 0.3 on a **passed** final exam
  - Bonus is transferable from previous and future semesters

# Other Administrative

- “External” students welcome (LMU, TUM PhD)
  - Fill out registration form and we will add you to the course
  - Will get Certificate / Schein at the end
- Again:
  - Check announcements on Piazza
  - Check content on website:  
<https://niessner.github.io/I2DL/>

# Practical Exercises

# Exercise – Goal

- Hands-on programming experience (learning by doing)
- Reimplement basic building blocks
- Introduction to common libraries
- (Get grade bonus)
- Ultimately: Gather enough experience to start your own individual (research) deep learning project

# Exercise – Format

- Tutorial:
  - Video only  
Posted on Piazza and course website
  - Video length  
Ex02: Full lecture with written exercise  
Ex03-12: Short (~30min) video and coding exercises
- Programming:
  - Interactive coding notebooks (~4h each)
- Start time: Thursdays 10:00
- Working Time: 1-2 weeks
- Deadline: Wednesdays, **23:59**

# Upcoming Lecture

- Important:
  - No lecture next week, Oct 21<sup>st</sup> (I'll be at ICCV – Hawaii)
  - Next lecture is on Oct 28<sup>th</sup>
- Next Lecture: Lecture 2: Machine Learning basics
- Thursday: Tutorial 1 and Exercise 1

See you next time ☺