

Deep Learning

A Disruptive Technology?

Tobias Würfl Pattern Recognition Lab, Friedrich-Alexander University of Erlangen-Nürnberg March 21, 2017





About Me



Tobias Würfl

- · PhD candidate
- CT reconstruction
- Consistency Conditions
- Deep Learning



What is Deep Learning?









Apples vs Pears







Apples vs Pears







Apples vs Pears







Feature Space





Feature Space













Apples vs Pears revisited







Apples vs Pears revisited







Artifical Neural Networks









Artifical Neural Networks





Deep Neural Networks





Evolution of Deep Learning

Algorithms



- Architectures
- Training



Evolution of Deep Learning

Algorithms

- Architectures
- Training

Graphics Cards



- Training speed
- Affordability



Evolution of Deep Learning

Algorithms



- Architectures
- Training

Graphics Cards



- Training speed
- Affordability

Frameworks



- Development speed
- Developer training



Deep Learning is

- very powerful
- extremely generic
- cheap to develop



What is Deep Learning good for?





Many applications

DEEP LEARNING EVERYWHERE





INTERNET & CLOUD

Language Processing

MEDICINE & BIOLOGY

ancer Cell Detection Diabetic Grading Drug Discovery

Video Captioning Video Search Real Time Translation SECURITY & DEFENSE

Face Detection Video Surveillance Satellite Imagery AUTONOMOUS MACHINES

Pedestrian Detection Lane Tracking Recognize Traffic Sign

Source: NVIDIA, Deep Learning on GPUs, 2016



Large Scale Image Classification 2012





ILSVRC

- \approx 1,4 million images
- labeled using mechanical turk
- dominated by Deep Learning



Large Scale Fine Grained Classification 2014



Deep Face

- Facebook
- Recognizing over 4000 persons
- Over 97% accuracy



Automatic Image Captioning 2014



a large elephant standing in a field of grass logprob: -8.37



a baseball player swinging a bat at a ball logprob: -5.24



Automatic Image Captioning 2014



a herd of elephants walking across a lush green field logprob: -6.48



a young boy is holding a baseball bat logprob: -7.65



System Specification

Deep Learning is bad at

- knowing it's limits
- explaining it's decisions
- failing gracefully



Playing Games 2013



Deep Q Learning

- 2600 games by one algorithm
- 29 out of 49 games better then humans

Source: Mnih et al., Human-level control through deep reinforcement learning 2013



Playing Incredibly Hard Games 2016



Alpha Go

- beat one of the best Go players
- beat him 4 : 1



Putting it to Use



Data Center Energy Consumption

- trained on history sensor data
- 40% decrease

Source: DeepMind.com



What about Healthcare?



What about Healthcare?

Siemens and Pattern Recognition Lab



Anatomical Landmark Detection 2016

- Works in 3D
- Across a range of modalities

Source: Ghesu et al., An Artificial Agent for Anatomical Landmark Detection in Medical Images, 2016



Mammography



Screening

- Very relevant
- Large false positive rate
- · Huge amounts of data



Towards Better Interpretability



Weakly Supervised Learning

- For free
- Shows the relevant areas
- Very active topic

Source: lunit.io



Skin Cancer

Epidermal lesions



Melanocytic lesions (dermoscopy)



Melanocytic lesions

Direct application of ImageNet architecture

- 70 % Accuracy
- At the level of dermatologists
- Works on smartphone images

Source: Esteva et al., Dermatologist-level classification of skin cancer with deep neural networks, 2017



Deep Learning for Reconstruction?



Deep Learning for Reconstruction?

Pattern Recognition Lab

Traditional Reconstruction



Neural Network



- Reconstruction mapped to network
- Iterating to train not for reconstruction

Source: Würfl et al., Deep Learning Computed Tomography, 2016



Conclusion





Risks

· Wide variety of applications

• ... but not always the best solution



- · Wide variety of applications
- Cheap to develop

- ... but not always the best solution
- ... which invites ad hoc solutions



- Wide variety of applications
- Cheap to develop
- Cheap to apply

- ... but not always the best solution
- ... which invites ad hoc solutions
- ... but tuned for graphics cards



- Wide variety of applications
- Cheap to develop
- Cheap to apply
- High accuracy

- ... but not always the best solution
- ... which invites ad hoc solutions
- ... but tuned for graphics cards
- ... but doesn't fail gracefully



Risks

- Wide variety of applications
- Cheap to develop
- Cheap to apply
- High accuracy

- ... but not always the best solution
- ... which invites ad hoc solutions
- ... but tuned for graphics cards
- ... but doesn't fail gracefully

Solves some previously intractable problems



Thanks for listening. Any questions?