Deep Learning & Artificial Intelligence

WS 2018/2019



Teaching Team



Prof. Dr. Matthias Schubert



Dr. Florian Buettner



Dr. Markus Geipel



Pankaj Gupta



Dr. Denis Krompaß



Dr. Sigurd Spieckermann



Prof. Dr. Volker Tresp

Schedule

17.10.2018	Introduction
24.10.2018	Basic Neural Networks
31.10.2018	Training Neural Networks
07.11.2018	Convolutional Neural Networks
14.11.2018	Recurrent Neural Networks
21.11.2018	Deep Learning and Uncertainty
28.11.2018	Representation and Distributional Learning
05.12.2018	Deep Learning Tools
12.12.2018	Generative Models
19.12.2018	Sequential Decision Problems and Autonomous Agents
09.01.2019	Model-free Reinforcement Learning
16.01.2019	Value Function Approximation
23.01.2019	Policy Gradient and Actor-Critic Learning
30.01.2019	Knowledge Graphs in Al
06.02.2019	Q&A

Deep Learning & Artificial Intelligence

Many people are talking about it, few people <u>really</u> know what they are talking about.

SOURCE: http://www.clickode.com/en/2016/02/01/google-rilascia-gratis-lezioni-di-deep-learning/

2

9

10

C

Dara

-

Darae

.................

Oata

Dara



















https://affinelayer.com/pixsrv



Input

Input

Output

https://www.youtube.com/watch?v=qc5P2bvfl44

H. Kim, P. Garrido, A. Tewari, W. Xu, J. Thies, M. Nießner, P. Pérez, C. Richardt, Michael Zollhöfer, C. Theobalt, Deep Video Portraits, ACM Transactions on Graphics (SIGGRAPH 2018)



Das Wirtschaftswachstum hat sich in den letzten Jahren verlangsamt .















Success! Success!

But there are also some issues ...











Goodfellow, I. J., Shlens, J., & Szegedy, C. Explaining and harnessing adversarial examples (2014). arXiv preprint arXiv:1412.6572.



Azulay, A., & Weiss, Y. (2018). Why do deep convolutional networks generalize so poorly to small image transformations?. arXiv preprint arXiv:1805.12177.



Guo, C., Pleiss, G., Sun, Y., & Weinberger, K. Q. (2017). On calibration of modern neural networks. arXiv preprint arXiv:1706.04599.



- Startups Apps
- , ibbe
- Gadgets
- **Events**
- Videos
- Crunchbase More
- Search Q
- zuckerberg testimony Privacy TC Sessions: Robotics 2018 Developer

Microsoft silences its new A.I. bot Tay, after Twitter users teach it racism [Updated]

Sarah Perez @sarahintampa / Mar 24, 2016



 \times



Microsoft's • newly launched A.I.-powered bot called Tay, which was responding to tweets and chats on GroupMe and Kik, has already been shut down due to concerns with its inability to recognize when it was making offensive or racist statements. Of course, the bot wasn't *coded* to be racist, but it "learns" from those it interacts with. And naturally, given that this is the Internet, one of the first things

entire ware tought Tay was how to be racist, and how to spout back ill-informed or

https://techcrunch.com/2016/03/24/microsoft-silences-its-new-a-i-bot-tay-after-twitter-users-teach-it-racism/



Brown, T. B., Mané, D., Roy, A., Abadi, M., & Gilmer, J. (2017). Adversarial patch. arXiv preprint arXiv:1712.09665.



Elsayed, G. F., Goodfellow, I., & Sohl-Dickstein, J. (2018). Adversarial Reprogramming of Neural Networks. arXiv preprint arXiv:1806.11146.

Deep Learning & Artificial Intelligence

in a nutshell



Machine Learning

SUPERVISED LEARNING

Classification Regression

UNSUPERVISED LEARNING

Clustering Dimensionality Reduction Recommendation

REINFORCEMENT LEARNING

Reward maximization

Deep Learning vs. Classic Data Modeling



Example: Sentiment Analysis

This movie was awesome!

This movie was not too bad.

This movie was really bad!

Hierarchical Feature Extraction



SOURCE: http://www.eidolonspeak.com/Artificial_Intelligence/SOA_P3_Fig4.png

Why is Deep Learning Taking Off?



Demystifying Deep Learning & Al



Demystifying Deep Learning & Al























Development Cycle



Deep Learning = Many Layers?

NO ... YES ... NOT ONLY

3

Flexibility due to computation graph



Little feature engineering = Less domain knowledge = Automated learning

Automated learning

Common framework for many applications

What's new?







More and faster computational resources (CPUs, GPUs, ...)



Improved optimization algorithms Improved software + open source ecosystem



Improved neural network architectures + building blocks

See you next week