Institute for Machine Learning

Seminar in AI , 01-12-2022
Sebastian Lehner

JOHANNES KEPLER
UNIVERSITÄT LINZ

Institute for Machine Learning & LIT AI Lab

JOHANNES KEPLER UNIVERSITÄT LINZ

- Head: Prof. Sepp Hochreiter
- 1 Assistant Professor & 8 Postdocs
- More than 35 PhD Students
- Research focus: Machine Learning, Deep Learning
- Initiated Austria's first AI Study Program at JKU
- LIT AI Lab at JKU:
  - 6 groups (Deep Learning, Computer Vision, Logical Reasoning, Pervasive Computing, Software Engineering, Symbolic Computing)
  - PhD School

Our Collaboration Partners

JKU
JOHANNES KEPLER UNIVERSITÄT LINZ

MERCK

janssen

ucb

ZF

Audi

AstraZeneca

BAYER

BOSCH

TÜV AUSTRIA

AVL

BOREALIS
Keep Discovering

DCS
COMPUTING GMBH

SAL
SILICON AUSTRIA LABS

TGW
LIVING LOGISTICS

enliteAI

scch
software competence center
hagenberg

Fill YOUR FUTURE

7LYTIX

ANYLINE

RISC
Software GmbH

here

zalando

Google

IARAI
institute of advanced
research in artificial
intelligence

ellis

European Laboratory for Learning and Intelligent Systems

European Artificial Intelligence Initiative

About

The ELLIS mission is the creation of a network to advance breakthroughs in AI, a pan-European PhD program to educate the next generation of AI researchers, and to boost economic growth in Europe by leveraging AI technologies.

info@ellis.eu

http://www.ellis.eu

Board Members

- Barbara Caputo (Italy)
- Nuria Oliver (Spain)
- Bernhard Schölkopf (Germany)
- Max Welling (Netherlands)

Board Member Deputies

- Matthias Bethge (Germany)
- Andreas Geiger (Germany)
- Sepp Hochreiter (Austria)
- Josef Sivic (France)

Deep Learning – Deep Neural Networks

- Multiple (hidden) layers in artificial neural network

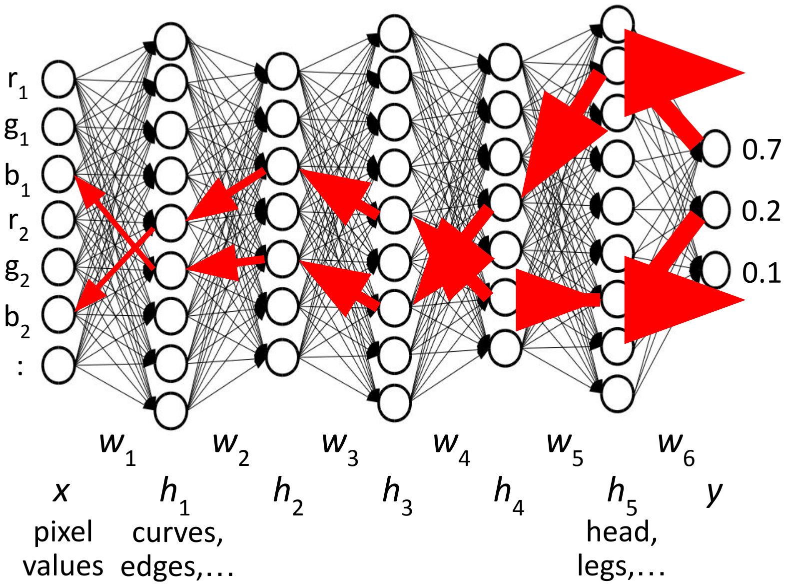

Horse
Dog
Cat

- Gradient descent minimizes error

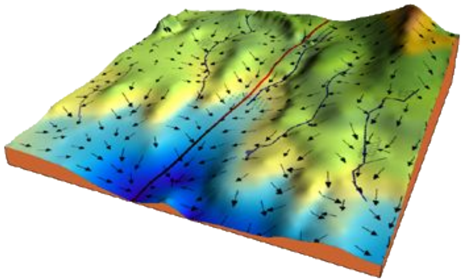

  - Feedback = total gradient: proportional to product of all individual gradients
  - The more layers, the smaller the total gradient
  - “Vanishing Gradient Problem”

Generative Adverserial Networks (GANs)

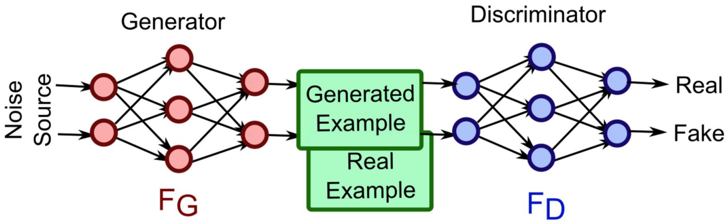

NVIDIA (2018)

- Generator network tries to fake images
- Discriminator network tries to distinguish between real and fake
- Both networks get better and better while training against each other

Quantum Optical Experiments Modeled by LSTM

- Reverse problem in quantum optics: which setups produce interesting (maximally entangled, high-dimensional) quantum states?
- Brute-force solution: Random search of setups
- Improvement: train LSTM network with millions of examples

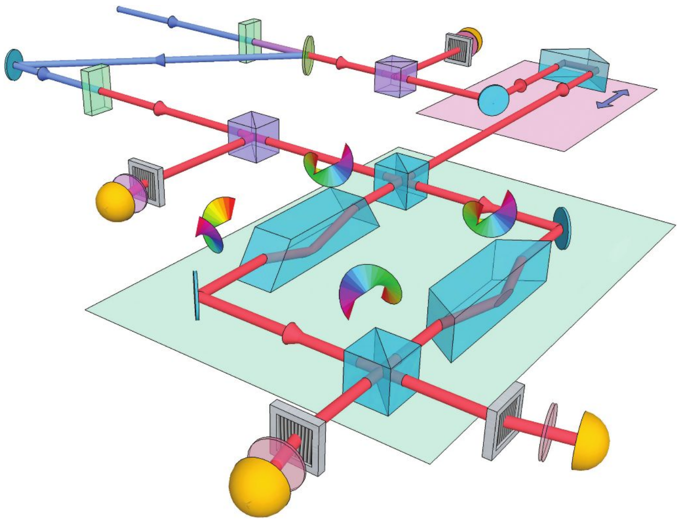

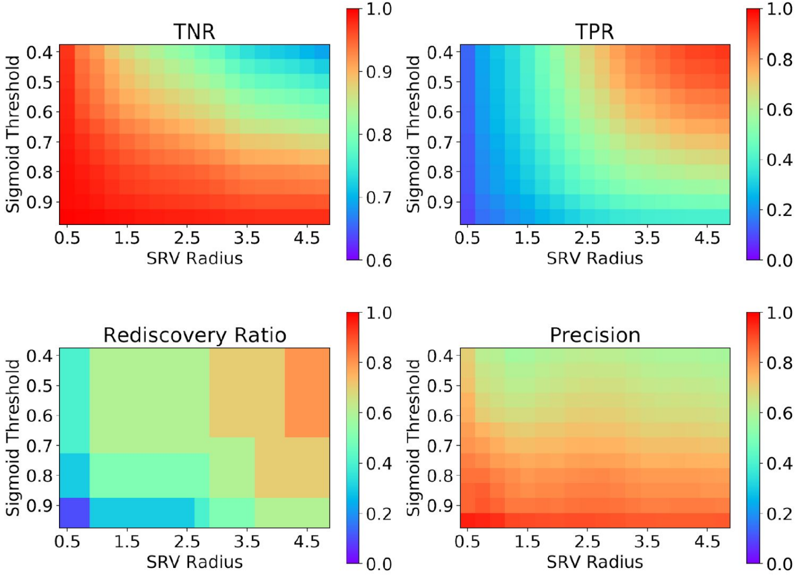

Reinforcement Learning

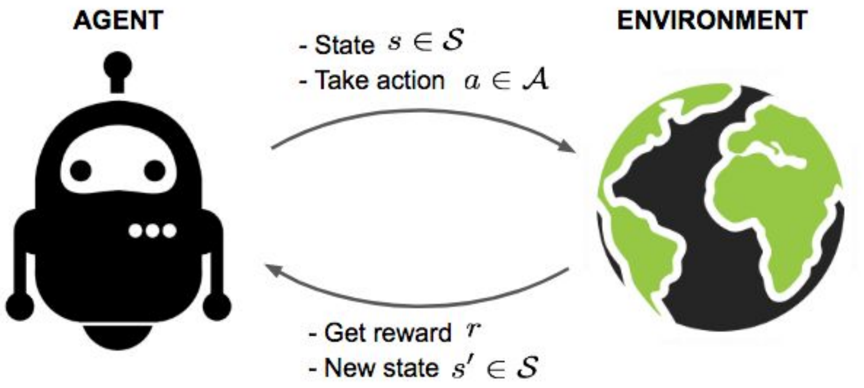

ETH Zürich (2017)

- Agents take actions in an environment to maximize reward (optimal control theory), e.g.

  for drone: not crashing

  in Go: winning the game

- At JKU: RUDDER (Return Decomposition for Delayed Rewards) [NeurIPS 2019]

Recent Publications: SELU, RUDDER, FID, Modern Hopfield Networks, Coulomb GAN, CLOOB …

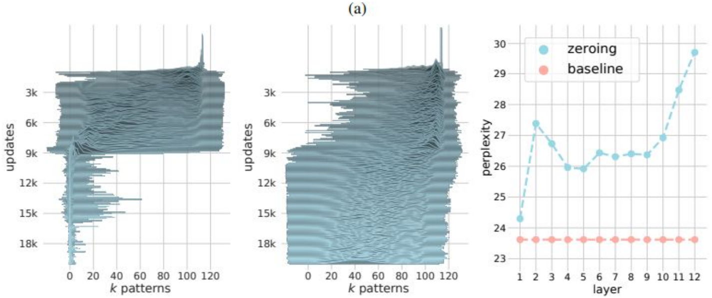

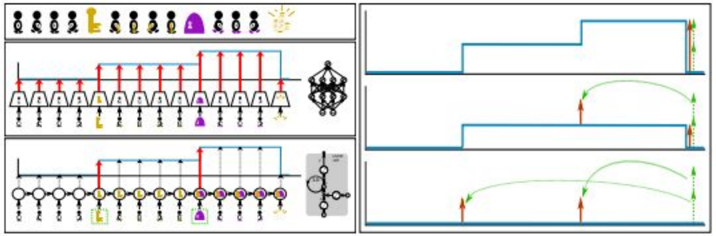

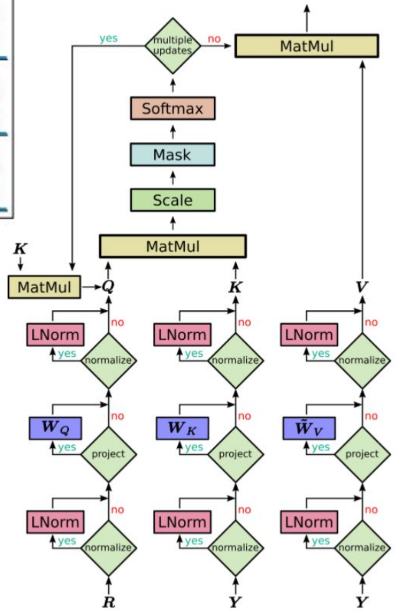

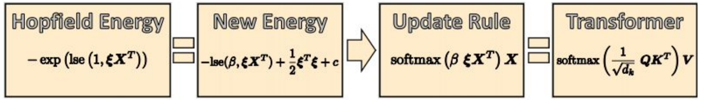

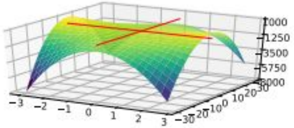

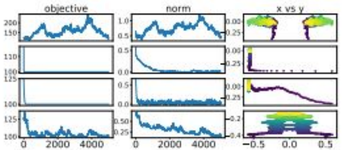

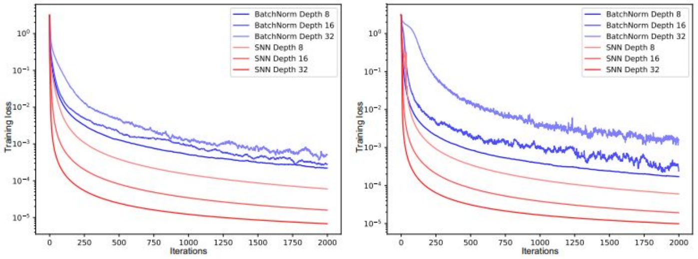

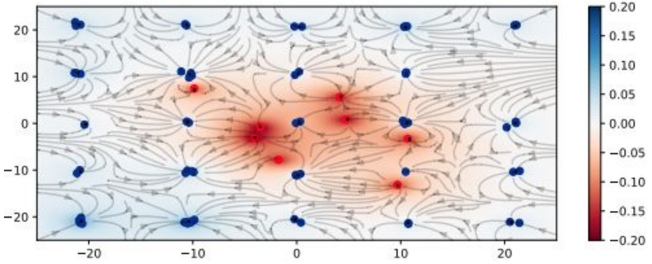

https://www.jku.at/en/institute-for-machine-learning/research/publications/overview/

AI in Life Science: Drug design, chemistry, molecular biology, medicine, healthcare

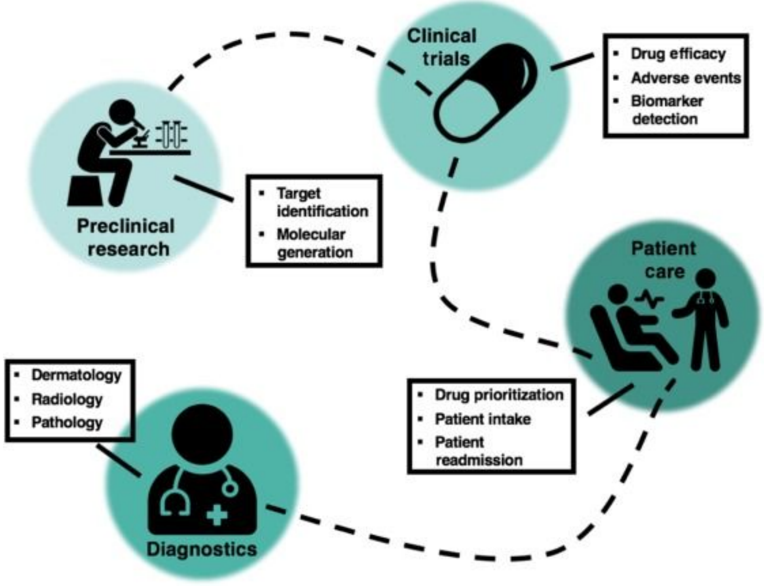

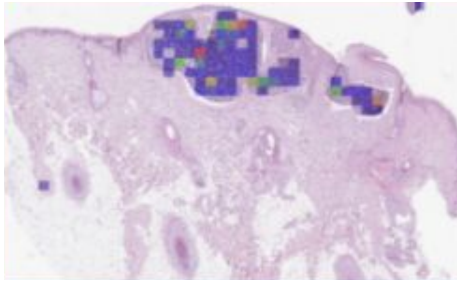

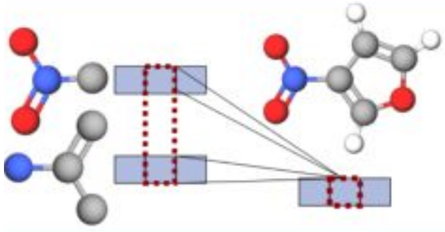

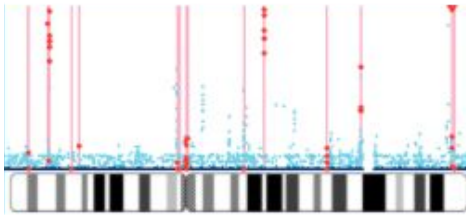

https://www.jku.at/en/linz-institute-of-technology/research/research-labs/artificial-intelligence-laboratory/ai-in-life-sciences-group/

Topics of Research

- Deep Learning
- Generative Models (GANs, VAE)
- Reinforcement Learning
- Transformers
- Modern Hopfield Networks
- Few-Shot Learning
- Meta-Learning

- Climate / Earth Science
- Planetary Science
- Physics: Classical & Quantum
- Autonomous Driving
- Drug Discovery / Life Science
- Industrial Applications
- Manufacturing
- Signal Processing (SAL)
- Chip Design (SAL)
- Certification of ML (TÜV)

Bachelor Theses

- Ideally an extension of the Practical Work project
- Familiarize with the relevant literature
- Formulate project objectives, research goals
- Design and conduct computational experiments
- Analyze experimental results and interpret them
- Write a Bachelor thesis:
  - scientific style & structure
  - 15-30 pages
  - standardized layout
- No oral presentation

Bachelor Theses Topics from last year

		Institute for Machine Learning
GAN	LG, PW	GANs in Theory and Practice
Reinforcement Learning	Mhm	RUDDER on PROCGEN
Reinforcement Learning	Mhm	Credit Assignment for Continuous Control Problems
3D Scene Understanding	JL	
LSTM	TA	RTRL for LSTM in PyTorch
LSTM, Differential Equations	MHI	Integral Equations with LSTM
Recommender Systems	AV	Recommender systems
Hydrology, LSTM	DK, MG	Pareto & rainfall-runoff LSTM models
Hydrology, LSTM	DK, MG	Probing of LSTM models in hydrology
LSTM	BN	Virtual Sensors Cooperation with Bosch Linz
Autonomous Vehicle	BN	DEEP SLAM for JO
Autonomous Vehicle	BN	JO, JKU's interactive AI on wheels
GAN	BN	Generative vs. Discriminative Learning
LSTM, Hopfield	TR	Prediction of unplanned ICU readmission
Contrastive Learning	ER, AF	Contrastive Learning for Ophthalmology
Cheminformatics; Bioactivity Modelling	PS	Libary for Chemical Modeling; Software Engineering
Reinforcement Learning; Cheminformatics	PS	RL on ChemRXN
Cheminformatics; Bioactivity Modelling	PS	Reaction Type Classification
Bioactivity Modelling	PS	Set Transformer for Bioactivity Modelling
Molecular modeling, geom DL	GK	Generative models for molecules in 3D space
Bioactivity Modelling, molecular modeling, geom DL	GK	Improving property prediction and generative models for molecules with geometry information
Bioactivity Modelling	GK	Dissecting successful DL architectures in drug discovery
Medical imaging	HB	Retinal image analysis with DL
Medical imaging	HB	Semantic segmentation in OCT
Medical imaging	TR	Self-supervised representation learning
Bioactivity Modelling	JS	Proteochemometrics
Meta-Learning	JS	Metalearning for Drug Discovery
Interpretability, Bioactivity Modelling	JS	Interpretability Methods for QSAR models
Uncertainty estimation; Bioactivity Modelling	AV	Uncertainty estimation in drug discovery
ML in Healthcare	TR	Prediction of subjective sleep quality
ML in Healthcare	TR	metalearning for recalibration of emg-based upper limp prostheses
ML in Healthcare	TR	machine learning for clinical care
Deep Learning: DNNs, Architectures	AM	Theoretical-Practical Aspects of Deep Learning
Deep Learning: DNNs, Architectures	AM	Bias Unit Sharing in Deep Neural Networks
Deep Learning: LSTM	PH	MC-LSTM on collision data
Deep Learning: LSTM	PH	MC-LSTM on traffic data
Deep Learning: LSTM	AM	Restricted LSTM
Deep Learning: DNNs, Training	AM	Combined BATCH-and-LAYER-Normalization
Deep Learning: Domain Adaptation	MHI, LG	Stability of Domain Adaptation
Deep Learning: Autoregressive modelling	PRe	Autoregressive modelling of molecules using Transformers
Deep Learning: Reaction modelling	PRe	Retrosynthesis prediction
Deep Learning: Hopfield	BS, LG	Deep Learning on Tabular Datasets (Classification/Regression)
Deep Learning: Hopfield	BS, LG	Deep Learning on Tabular Datasets (Data Imputation)
Deep Learning	FP, AF	Deep Learning for UAV Control
Deep Learning: Language Models	FP	Zero-Shot Transfer of Pretrained Language Models

Seminar in AI - Hochreiter, Lehner (3 ECTS)

- Choose one paper from a list of recent ML-Papers
  - NeurIPS (Neural Information Processing Systems)
  - ICML (Int. Conf. on Machine Learning)
  - ICLR (Int. Conf. on Learning Representations)
- Read, understand and analyse the paper
- Gather additional information as needed (!!) from textbooks or other papers
- Prepare slides and present the paper in a talk
- Explain the theory and mathematics in an understandable way
- Highlight novelty and achievements, pinpoint weaknesses
- Be prepared for a deep discussion about the paper
- Hand in a short report in scientific writing style

Seminar Papers from last year

Variational Autoencoder:

1. Very Deep VAEs Generalize Autoregressive Models and Can Outperform Them on Images.
   https://arxiv.org/abs/2011.10650
2. Inference Suboptimality in Variational Autoencoders
   https://arxiv.org/abs/1801.03558
3. VAE with a VampPrior
   http://arxiv.org/abs/1705.07120
4. Hierarchical Quantized Autoencoders
   http://arxiv.org/abs/2002.08111
5. Variational Memory Addressing in Generative Models
   https://arxiv.org/abs/1709.07116
6. NVAE: A Deep Hierarchical Variational Autoencoder
   https://arxiv.org/abs/2007.03898
7. Exemplar VAE: Linking Generative Models, Nearest Neighbor Retrieval, and Data Augmentation
   https://arxiv.org/abs/2004.04795

Attention, Transformers, modern Hopfield Networks, Hopfield Layers:

8. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale
   https://arxiv.org/abs/2010.11929
9. TransGAN: Two Transformers Can Make One Strong GAN
   https://arxiv.org/abs/2102.07074
10. (DO NOT SELECT!) Generative Adversarial Transformers
   http://arxiv.org/abs/2103.01209
11. Generating Long Sequences with Sparse Transformers
   https://arxiv.org/abs/1904.10509
12. Hopfield Networks is All You Need
   https://arxiv.org/abs/2008.02217
13. Dense Associative Memory for Pattern Recognition
   http://arxiv.org/abs/1606.01164
14. Dense associative memory is robust to adversarial inputs
   https://arxiv.org/abs/1701.00939
15. Attention Is All You Need
   https://arxiv.org/abs/1706.03762
16. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding
   http://arxiv.org/abs/1810.04805
17. Modern Hopfield Networks and Attention for Immune Repertoire Classification
   https://arxiv.org/abs/2007.13505
18. Large Associative Memory Problem in Neurobiology and Machine Learning
   https://arxiv.org/pdf/2008.06996
19. Visual Transformers: Token-based Image Representation and Processing for Computer Vision
   https://arxiv.org/abs/2006.03677

Reinforcement Learning

20. Mastering the game of Go with deep neural networks and tree search (incl. Appendix of pdf)
   https://www.nature.com/articles/nature16961
21. Mastering the game of Go without human knowledge (incl. Appendix of pdf)
   https://www.nature.com/articles/nature24270
23. Playing Atari with Deep Reinforcement Learning
   https://arxiv.org/abs/1312.5602
24. Proximal Policy Optimization Algorithms
   https://arxiv.org/abs/1707.06347
25. RUDDER: Return Decomposition for Delayed Rewards
   https://arxiv.org/abs/1806.07857
26. Align-RUDDER: Learning From Few Demonstrations by Reward Redistribution
   https://arxiv.org/abs/2009.14108

Other current topics:

27. Adam: A Method for Stochastic Optimization
   https://arxiv.org/abs/1412.6980
28. Cross-Domain Few-Shot Learning by Representation Fusion
   https://arxiv.org/abs/2010.06498
29. Learning to Simulate Complex Physics with Graph Networks
   http://arxiv.org/abs/2002.09405
30. Learning Mesh-Based Simulation with Graph Networks
   http://arxiv.org/abs/2010.03409
31. Implicit Generation and Generalization in Energy-Based Models
   http://arxiv.org/abs/1903.08689
32. A Style-Based Generator Architecture for Generative Adversarial Networks
   http://arxiv.org/abs/1812.04948
33. MC-LSTM: Mass-Conserving LSTM
   https://arxiv.org/abs/2101.05186
34. Coulomb GANs: Provably Optimal Nash Equilibria via Potential Fields
   https://arxiv.org/abs/1708.08819
35. Overcoming catastrophic forgetting in neural networks
   https://arxiv.org/abs/1612.00796
36. Overcoming catastrophic forgetting with hard attention to the task
   https://arxiv.org/abs/1801.01423
37. A Simple Framework for Contrastive Learning of Visual Representations
   https://arxiv.org/abs/2002.05709
38. Pointer Networks
   https://arxiv.org/abs/1506.03134
39. Deep Sets
   https://arxiv.org/abs/1703.06114
40. Deep Anomaly Detection with Outlier Exposure
   https://arxiv.org/abs/1812.04606

THANK YOU

Questions?

JOHANNES KEPLER
UNIVERSITÄT LINZ