



Seminar Wissenschaftliches Arbeiten

Computational Aspects of Machine Learning

Maximilian Thiessen
Prof. Thomas Gärtner
September 20, 2020





Prof. Thomas Gärtner




Maximilian Thiessen


Datalogy, Data Science, ...



Peter Naur (1974):
*The **science of dealing with data**, once they have been established, while the relation of the data to what they represent is delegated to other fields and sciences.*



(Vasily Zubarev, 2018)



(Drew Conway, 2010)



John W. Tukey (?):
The best thing about being a statistician is that you get to play in everyone's backyard

Don't forget that statisticians are the free-est of all scientists — they can work on anything. Take advantage.

John Quackenbush (2014):
Every revolution in science—from Copernican heliocentric model to the rise of statistical and quantum mechanics, from Darwin's theory of evolution and natural selection to the theory of the gene—**has been driven by one and only one thing: access to data**



Application Areas



self-driving cars
human level game AI
recommendation systems
predictive maintenance
personalised medicine
personal assistants
sentiment analysis
(cyber-)security
drug discovery

ML for advanced materials

ML for social good

ML for longer life

ML for demand estimation

ML for digital humanities

ML for production

ML for diagnoses

ML for marketing

ML for products

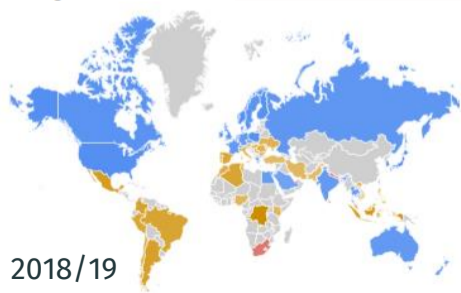


ML for services

ML for ...


...



Machine Learning, a global trend



Machine Learning Tools



Source: <http://xkcd.com/1425/>

In this seminar, you will...

- experience how a typical machine learning conference is organised
- familiarise yourself with a specific area in machine learning
- learn about learning algorithms and their properties
- understand a few research papers
- summarise and present research results

We will use [easychair](#) for submissions and reviews

Format



You will take on the role of **author** and **reviewer** for a machine learning conference:

- submit two short talks and abstracts
- bid for projects of your fellow students
- we will assign you one of your own projects to work on and the projects to review later
- give a progress presentation and submit a report draft
- review the draft reports assigned to you
- give a final talk and submit a final report

Topics

We will focus on computational aspects of machine learning

Algorithmic properties are central, e.g., [computational and sample complexity](#)

We encourage your own creative ideas




(but you will need to have them worked out well)

Check our [teaching homepage](#) for further details

Topics

Example topics include

- active learning
- online learning
- kernel methods
- learning with graphs
- unsupervised learning
- semi-supervised learning
- interpretable machine learning



Attend the *mandatory* first (online) meeting

- date and details will be announced on our [teaching homepage](#)
- propose your own project idea or choose one of ours
 - make sure your idea is **well thought through** and developed according to our guidelines
 - we will give preference to students who can **present and discuss** the project to some detail in the first meeting

Examination Modalities

Your final grade results from your...

- short spotlight presentations
- first abstracts
- progress presentation
- draft report
- reviews for your fellow students
- final presentation
- final report