

Recent trends in Automated Machine Learning (AutoML)

Summer semester 2021

Tim Meinhardt and Prof. Dr. Laura Leal-Taixé

Outline





- What is AutoML?
- Organization
 - General information
 - Presentation
- Paper and date matching

What is AutoML?

Machine and Deep Learning

Inputs





- Tasks (Classification, Regression, etc.)
- Datasets (research, real, non-vision)



What is AutoML?

Learn a task/dataset specific model:

- Data processing
- Architecture design
- Optimization




Hyperparameter optimization!

What is AutoML?

Enhance progress (state-of-the-art) on inputs

- Research
- Industry



Machine learning experts (or graduate student descent)!



Automated Machine Learning (AutoML)



How to AutoML?

Classic

- Grid or random searches
- Bayesian optimization (TPE, Spearmint, SMAC, etc.)
- ...



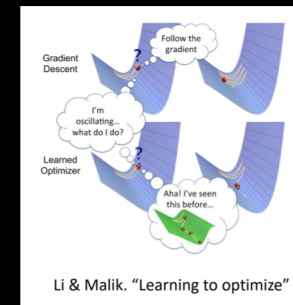




Follow modern trend and apply machine learning.

Learning to learn or Meta Learning

Meta Learning

Leverage power of learning methods to improve learning:

- Few shot learning
- Improve transfer learning (ImageNet)
- Multi-task initialization learning
- Faster Reinforcement Learning
- Optimal architectures
- Improved optimizers



General information

- Website: https://dvl.in.tum.de/teaching/automl_ss21/
- Contact: tim.meinhardt@tum.de
- Room:
 - Virtual event via Zoom
 - Links will be shared via email
- Attendance is mandatory!

Schedule:

- Introduction: 14th April 11 am – 12 pm
- Presentations: Wednesdays 11 am – 12 pm, TBD

Presentation

- Read and work through the paper
- Note questions and difficulties

Three weeks before:

Arrange meeting to discuss and clarify paper.

One week before:

Arrange meeting to discuss slides.

One week after:

Submit presentation slides as PDF via email.

- Duration
 - 20 minutes talk + 10 minutes discussions
 - Rule of thumb: 1-2 minutes per slide, i.e., 10-20 slides
 - Finish talk on time!
- Content
 - Explain the paper in your own words
 - Highlight strengths and weaknesses
 - Complement with additional material and explanations (from an l2DL perspective)
 - Put your paper into perspective with respect to previous presentations

Paper and date matching



Paper categories

- Reinforcement Learning
- Data
- Network architectures
- Optimization