Recent trends in Automated Machine Learning (AutoML)

Summer semester 2019

Tim Meinhardt and Prof. Dr. Laura Leal-Taixé

Outline

- What is AutoML?
- Organization
  - General information
  - Course and paper matching
  - Presentations
- Paper preview

What is AutoML?

Machine and Deep Learning

Inputs

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

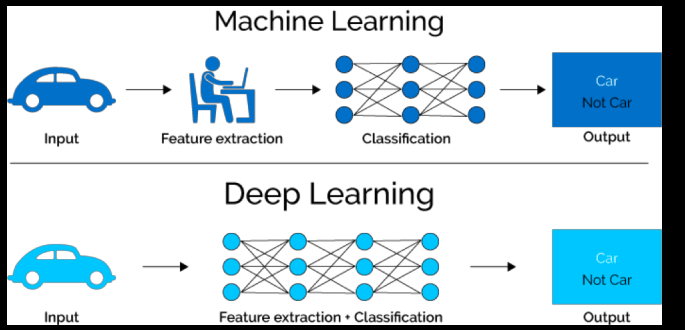

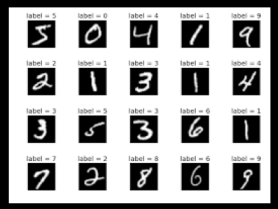

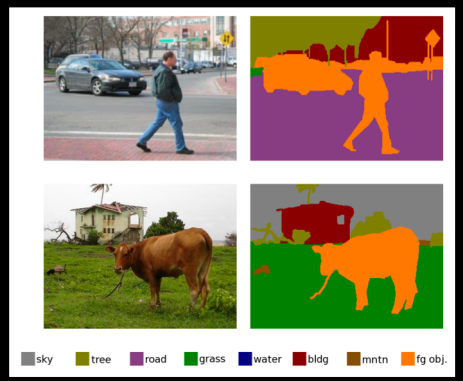

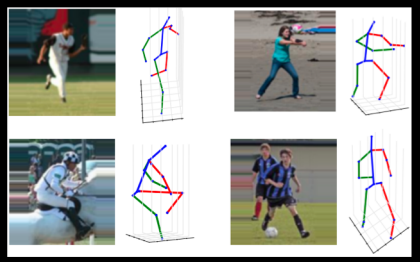

What is AutoML?

Learn a task/dataset specific model:

- Architecture design
- Data processing
- Optimization

Hyperparameter optimization!

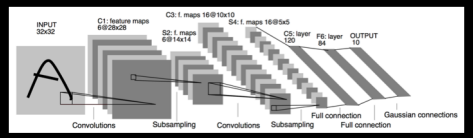

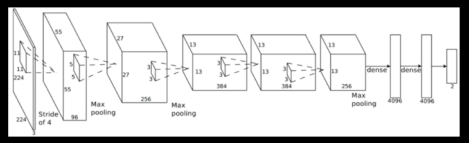

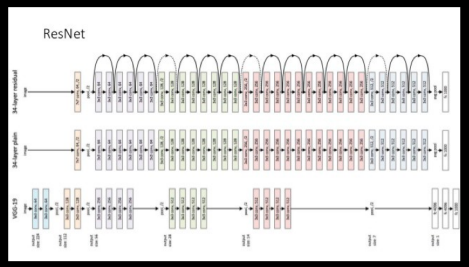

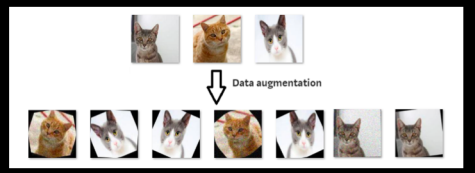

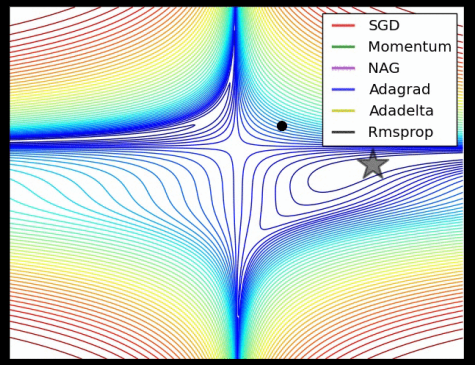

What is AutoML?

- Enhance progress on existing inputs
- Produce state-of-the-art outputs for new inputs
  - Research
  - Industry

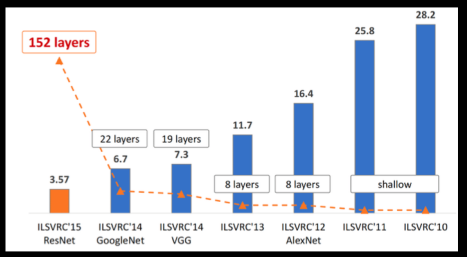

Machine learning experts (or graduate student descent)!

Automated Machine Learning (AutoML)

AUTO KERAS

Auto-Sklearn

Google Cloud

How to AutoML?

Classic optimization:

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

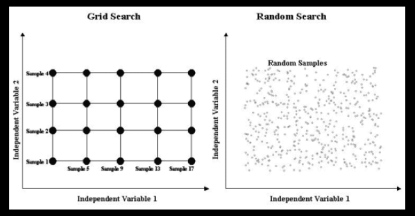

SMAC

Learning to learn or Meta Learning

Meta Learning

Leverage power of learning methods to improve learning:

- Few shot learning
- Pretraining on ImageNet
- Multi-task initialization learning
- Fast Reinforcement Learning
- Learning architectures
- Learning optimizers

AutoML

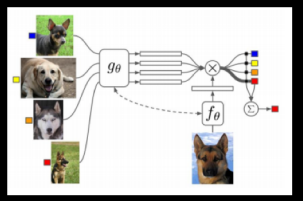

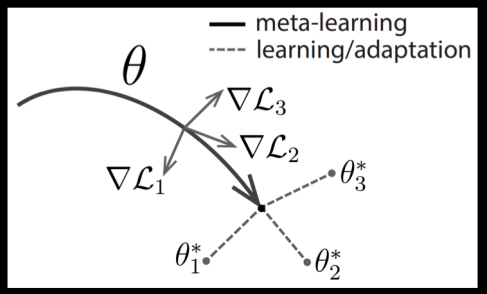

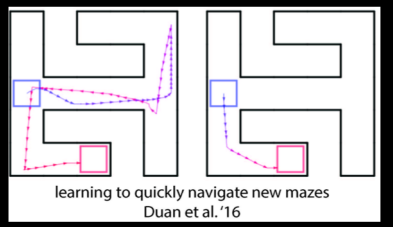

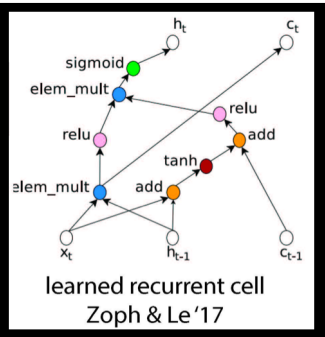

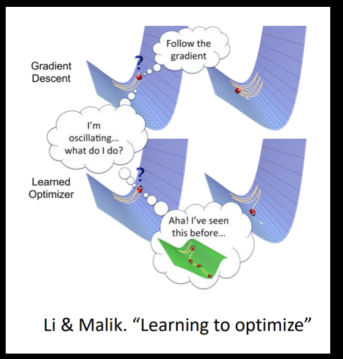

Organization

General information

- Website: https://dvl.in.tum.de/teaching/automl_ss19/
- Contact: tim.meinhardt@tum.de
- Room: MI 02.09.023
- Time: 12 participants -> 6 sessions
- Attendance is mandatory!

Schedule:

- Pre-course meeting: 25th January 1 – 3 pm
- Paper matching: 25th April 2 – 4 pm
- Presentations: Thursdays 2 - 4 pm, TBD

Matching

- Course matching (https://matching.in.tum.de/)
  - See FAQ for details
  - Registration period: 8th - 13th February
  - Preference: I2DL or DL4CV grade (contact us if external student)
  - Announcement: 20th February

- Paper matching
  - Study our list of suggested papers (website 8th February)
  - Propose own paper until 20th April
  - On the 25th April
    - Match paper based on preferences (toss a coin if necessary)
    - Fix dates for the presentations

Before the 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

Presentation

- Duration: 20 minutes talk + 10 minutes discussions
- Finish talk on time!
- Explain in own words
- Complement paper content with additional material and explanations (from an I2DL perspective)
- Rule of thumb: 1-2 minutes per slide, i.e., 10-20 slides
- Submit PDF until 1 week after presentation

Paper preview

Asynchronous Methods for Deep Reinforcement Learning. Mnih et al.

- Q-Learning
- Advantage Actor-Critic

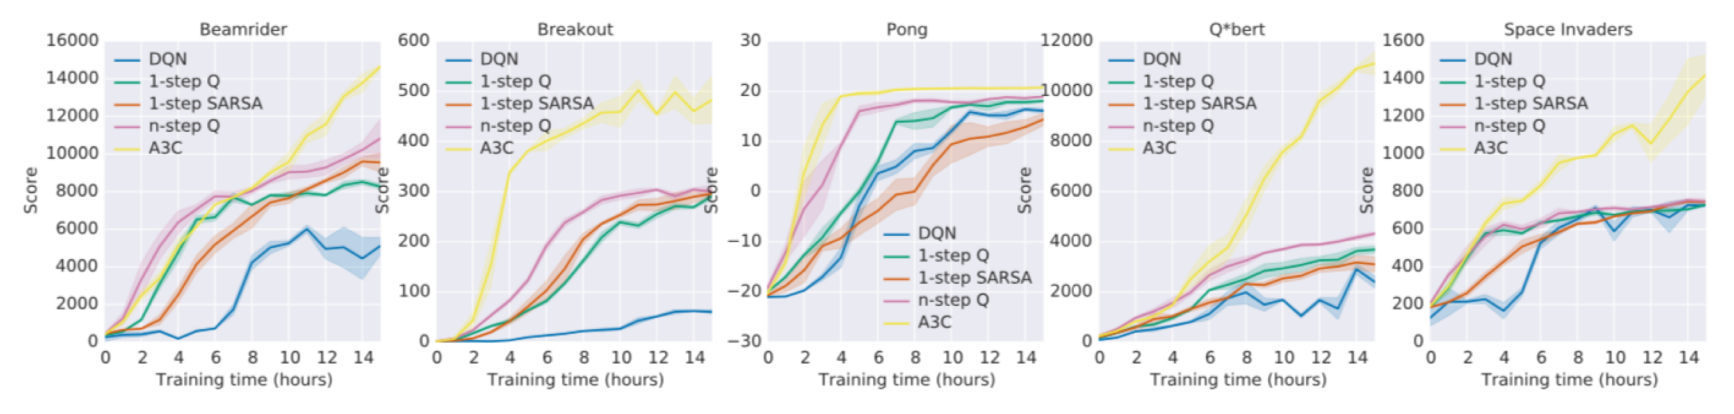

Paper preview

Proximal Policy Optimization Algorithms. Schulman et al.

- (Proximal)Policy gradient methods
- Trust region methods

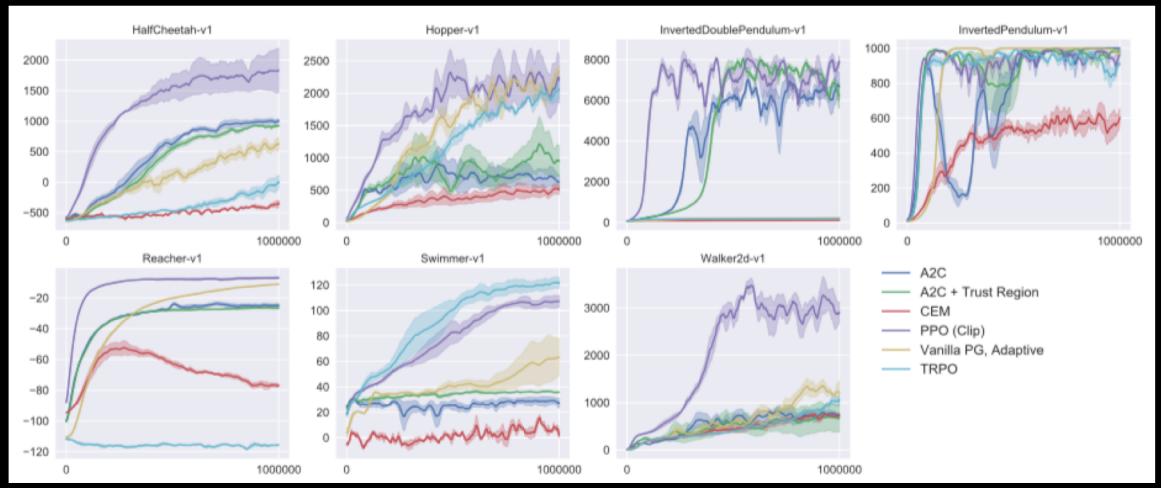

Paper preview

Neural Architecture Search with Reinforcement Learning. Zoph et al.

- Recurrent network to predict architectures (NAS)
- Trained with RL

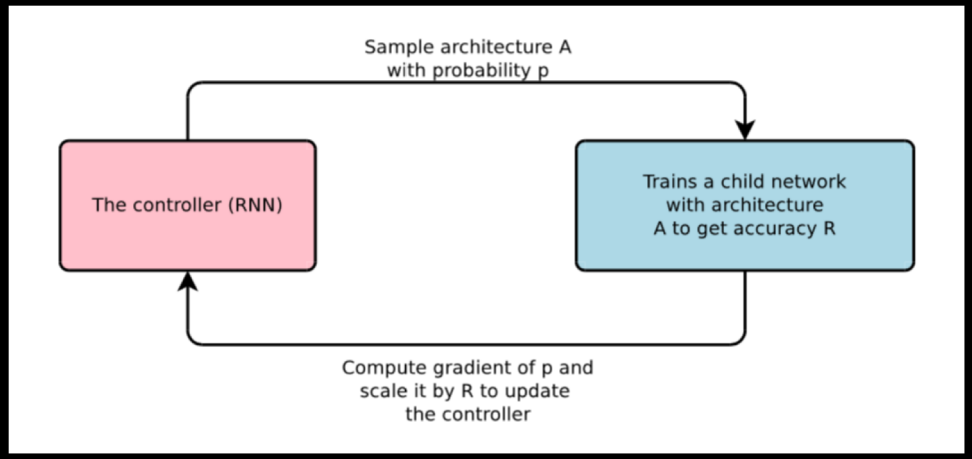

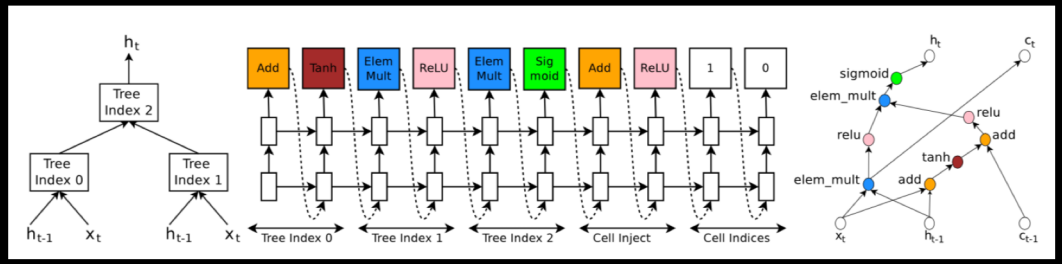

Paper preview

Learning Transferable Architectures for Scalable Image Recognition. Zoph et al.

- Extension of NAS with new architecture search space
- Applicable to large datasets

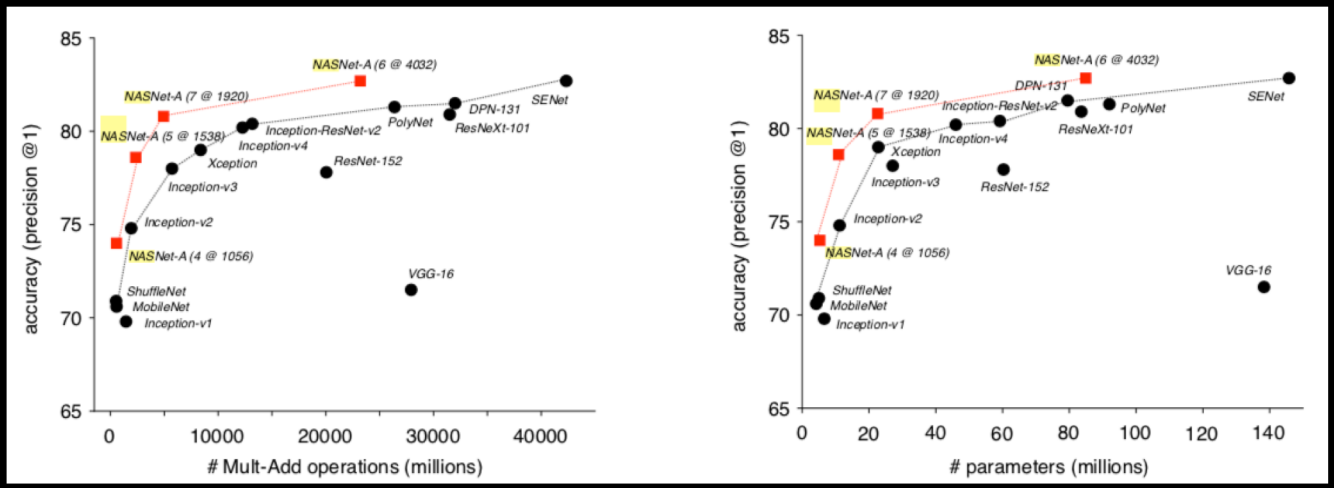

Paper preview

Learning to learn by gradient descent by gradient descent. Andrychowicz et al.

- Design of optimizer casted as a learning problem
- Generalizes to unseen tasks

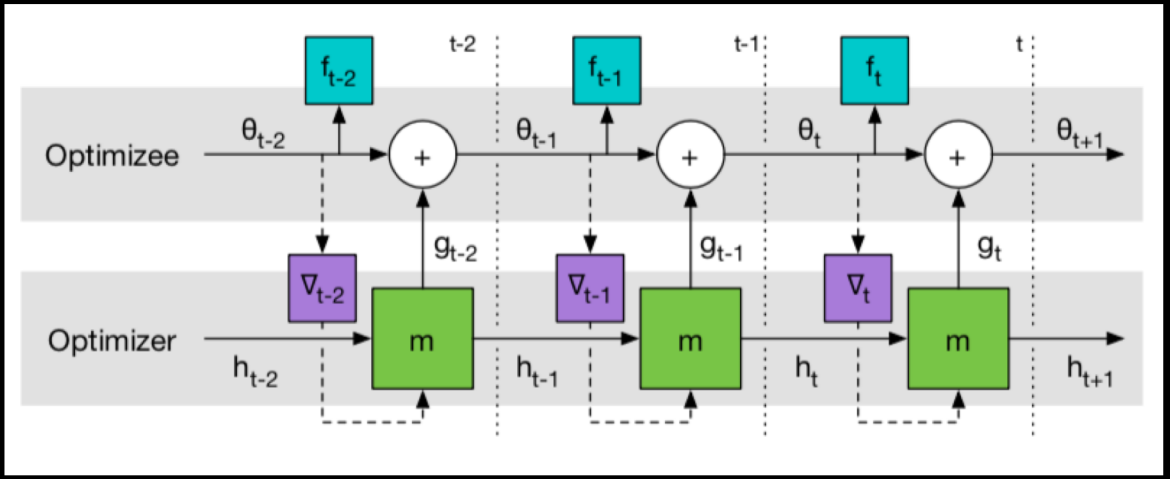

Paper preview

Searching for Activation Functions. Ramachandran et al.

- Apply reinforcement learning to discover new activation
- New activation function *Swish*

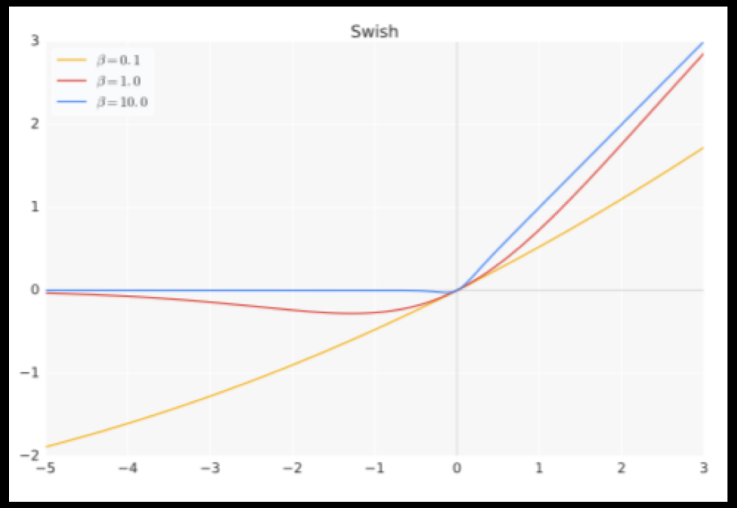

Paper preview

Learning Step Size Controllers for Robust Neural Network Training. Daniel et al.

- Learned learning rate scheduler
- Reinforcement Learning