ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

inf | Informatik
Computer Science

Seminar: Advanced topics in Machine Learning

Prof. Joachim Buhmann
Prof. Andreas Krause
Prof. Gunnar Rätsch

Goal of this seminar course

Practice giving scientific presentations

Learn about advanced topics in machine learning

Goal of today

- Seminar topics overview
  - ~40 papers in several topic areas
- Assignment of students to these topics
  - Detailed assignment of papers done later
- Discuss course organization, grading

(Reinforcement) Learning and controlling dynamical systems
Adversarial Examples and Robustness
Adversarial learning
Bayesian Nonparametrics
Deep learning, architecture, RNN
Deep learning, memory augmented models
Deep learning, regularization
Deep learning, theory
Deep representation learning
Deep representation learning, discrete latent representations
Deep structure learning, Bayesian networks
Disentangled Representations
Fairness and Interpretability
GANs
GPs and neural networks
Generative model
Implicit Models
Matrix Factorization/Causality
Model interpretability
Neural Networks
Perturbations, Optimization and Statistics
Sequential model
Submodularity and Discrete Optimization
Unsupervised Learning
Unsupervised Learning/Causality
Variational Inference

Course organization

- Two time slots:
  - **Tuesday 16-18 in CAB G56**
  - **Thursday 16-18 in CAB G57**
- You are expected to present in and regularly attend one of these slots
  - n sessions per slot (probably, ~10) with 2 talks each
  - You must attend n-1 sessions in the slot you're presenting.
  - You may switch between slots in one week.
- First talks will start on 9. and 11. October, respectively.

Grading

- Your grade will be determined based on your talk, as well as participation in the discussion
- Criteria:
  - Structure (how well is your talk organized?)
  - Understandability (how understandable is your oral presentation and slide design?)
  - Completeness (how well do you provide right background, and manage to focus on what is important and relevant?)
  - Activity (how engaged are you in class and in the talk preparation?)
  - Independence (how independent are you in preparing the presentation, and in reflecting on the paper?)

Presentation details

- Use electronic slides (ppt, pdf, ...)
- Talk length: 30 min + 15 min discussion
- The talk should provide sufficient background to be understandable to someone who has taken an ML class
- You should present the papers contributions and results, as well as reflect on them

Typical talk outline

- Introduction (Motivation and background)
- Formal problem statement (Notation, ...)
- Technical contribution (algorithm, theoretical result, ...)
- Experimental results (if any)
- Discussion (what are perceived strengths and weaknesses of the paper; what could be done more; ...)
- Conclusion

Advice on presentation design

- Giving compelling presentations is hard!
  - „Minimize words and maximize illustrations“
  - Focus on giving intuition, identifying key insights etc.

- Some pieces of advice:
  - http://www.cs.berkeley.edu/~jrs/speaking.html
  - http://research.microsoft.com/en-us/um/people/simonpj/papers/giving-a-talk/giving-a-talk.htm
  - http://greatresearch.org/2013/10/04/presenting-a-technical-talk/

Topic assignment

- Each student gives ordered preferences for 3 topics
- We match students to topics according to preference

https://goo.gl/forms/bL45KhiM3X9gsD2B3

- Each topic will have a representative, who will assign papers from the topic to students
- For each paper, there will be an advisor whom you can ask for clarification and advice
- Set up at least one meeting with your advisor to discuss the paper / presentation (weeks before talk)
- Send (near-complete) draft of slides to advisor at least 7 days prior to presentation date