

Seminar: Efficient Inference and Large-Scale Machine Learning

Pre-course meeting

Stephan Günnemann
Aleksandar Bojchevski
Oleksandr Shchur


Technische Universität München
Department of Informatics
Data Mining and Analytics
kdd.in.tum.de

January 27, 2017

Introduction

Most problems in machine learning can be formulated as inference or optimization

- Regression
- Neural networks
- SVM
- Mixture of Gaussians
- ...



Example: Linear Regression

Question

- How can we estimate the weights \mathbf{w} in a linear regression model given the data \mathcal{D} ?

$$y = \mathbf{w}^T \mathbf{x} + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma^2)$$

$$\mathcal{D} = \{\mathbf{x}_i, y_i\}_{i=1}^N$$

Possible solutions

- Minimize loss
- Maximize data likelihood
- Full Bayesian approach


$$p(\theta|\mathcal{D}) = \frac{p(\mathcal{D}|\theta)p(\theta)}{p(\mathcal{D})}$$

Benefits of Bayesian approach

- + Provides uncertainty measure
- + Avoids overfitting
- + Allows to incorporate prior beliefs

Main challenge

- Computational complexity



Bayesian inference

- Variational inference
- Sampling
- Particle filters
- Exact inference in graphical models

Large-scale optimization

- Advanced gradient-based optimization
- Bayesian optimization
- Probabilistic numerics
- Large-scale learning systems (MXNet, TensorFlow)

You are going to learn

- about the state of the art machine learning techniques for probabilistic inference and large-scale optimization
- to read and understand scientific publications
- to write a scientific report
- how to prepare and give a technical talk

Requirements

Paper

- 5 - 8 pages
- Latex template on the course webpage

Presentation

- 30 minutes talk
- 15 minutes discussion

Reviews

- Everyone has to review 2 papers by other students

- **1 week before the talk** - submission of extended abstract and slides
- **Day of the talk** - submission of preliminary paper for review
- **1 week after the talk** - receiving comments from reviewers
- **2 weeks after the talk** - submission of the final paper

The grade is determined based on

- Report
- Presentation (slides and speech)
- Reviews written by **you**
- Involvement in the class
- Interactions with the supervisor
- Extra bonuses for own contributions (e.g. visualizations, demos, experiments)

- **Before 08.02.** - fill out the pre-course survey
<https://goo.gl/forms/4M1obFN0yePuUq2Q2>
- **03.02. - 08.02.** - registration via the matching system
- **After 15.02.** - notification of the participants and selection of topics
- **April - June** - weekly sessions every Monday 12:30 - 14:00

Questions?