Seminar: Robust Data Mining Techniques

Kickoff meeting

Stephan Günnemann
Aleksandar Bojchevski
Oleksandr Shchur
Amir Moin
Roberto Alonso

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

Match 1, 2017

Organization

Timeline

March

- **Until 08.03** - send us your preferred topics via email.
- **10.03** - get assigned a topic and a supervisor
- **After 10.03** - work on your topic, meet with your supervisor

April - July

- **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 *reviews* from your peers
- **2 weeks after the talk** - submission of the *final paper*

Deliverables

Extended abstract

- 1 page, documentclass article

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

Goal

Your paper and presentation should

- Introduce the problem setting.
- Provide a summary of the topic.
- Describe main ideas and important results.
- Mention applications and connections to other methods.

Grading

The grade is determined based on

- Extended abstract
- 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)
- Penalties for missed deadlines

Topics

Robust Extensions

- Generalizations of classical methods for handling noisy data.
- Different data mining tasks on vector / graph / temporal data.
- Probabilistic models, additive decomposition, ...

Attacks on Classifiers & Fooling Deep Networks

- Machine learning algorithms can be easily fooled.
- Different attack types: poisoning, evasion
- Human-imperceptable noise breaks classifiers.

Learning in the Adversarial Setting

- Sparse and directed corruptions in the data.
- Adversary adapting to the defense strategy.
- Game-theoretic view of adversarial classification.

Learning from Crowds

- Harnessing the Big Data.
- Several unreliable sources of information.
- Multiple possibly contradicting labels per instance.

Differential Privacy

- Maximize the accuracy while restricting identification of the individuals.
- Mathematical formulation of privacy.
- Each instance has little effect on the decision boundary.

Robustness of Complex Networks

- Complex networks sustain their functions even when components fail.
- Role of networks in the cascading failures.
- Connections to statistical mechanics / percolation theory.

Recap

Most important points

- Send us your preferred topics until 08.03.
- Let us know if you want to deregister until 15.03.
- Do not work on your topics completely on your own. Reach out to your supervisors.

Questions?