

Reading group in data mining and machine learning a course intro

Jiří Kléma, Filip Železný

Department of Computer Science,
Czech Technical University in Prague



<http://cw.felk.cvut.cz/wiki/courses/xp36rgm/start>

The course content

- Modern and recent data mining methods,
- modern and recent machine learning methods,
- follows the earlier course Selected parts of data mining
 - a better name for the content that did not change,
 - a sister ML course planned, always joined with the DM course,
- Related bachelor and master courses (not direct prerequisites)
 - Pattern recognition,
 - Symbolic/statistical machine learning,
 - Statistical data analysis,
 - many other (optional) courses
 - * Big data technologies,
 - * Vision for robotics (deep learning).

Resources

- Books for traditional indispensable topics
 - Hastie, T., Tibshirani, R., Friedman, J.: The Elements of. Statistical Learning: Data Mining, Inference, and Prediction. Springer, 2009.
 - Rajaraman, A., Leskovec, J., Ullman, J. D.: Mining of Massive Datasets, Cambridge University Press, 2011.
 - Goodfellow, I., Bengio, Y., Courville, Y.: Deep Learning, MIT Press, 2015.
- research/review papers
 - see examples on the course webpage,
- recent issues in relevant journals
 - Journal of Machine Learning Research, Machine Learning,
 - Data Mining and Knowledge Discovery
- recent tutorials, major ML/DM conferences
 - KDD, Discovery Science, ICML, ECML, NIPS/NeurIPS.

Expectations

- Reading class
 - no "labs", presentations/discussions only, most time spent on homework,
- each participant will prepare and give two presentations
 - around 60min each + discussion,
 - the first one a broader review
(explore, learn a new topic that appealed you for long time),
 - the second one can touch/extend your current research if DM/ML related
(exploit, show us what you do or/and boost your current research),
- you are expected to regularly attend the class
 - and read review papers recommended for topics presented by others,
 - and prepare a few questions for discussion,
- the course will be concluded with an exam
 - rather an informal one,
 - still, every student should review the presented topics before the exam.

The main issues for the first class

- Introduce each other
 - name, supervisor, dissertation topic, research interests,
- The initial discussion about the first assignments
 - the topics should be settled at least two weeks in advance ...
 - ...and publicly agreed by the audience.