FEP3260 Fundamentals of Machine Learning Networks 10.0 credits

Grunderna i maskininlärning över nätverk

This is a translation of the Swedish, legally binding, course syllabus.

Establishment

Course syllabus for FEP3260 valid from Spring 2019

Grading scale

P, F

Education cycle

Third cycle

Specific prerequisites

Basic knowledge of convex optimization and probability theory is required to follow the course.

Language of instruction

The language of instruction is specified in the course offering information in the course catalogue.

Intended learning outcomes

After the course, the student should be able to:

· give new tools and training to model basic ML problems by optimization

· present basic theories of large-scale ML, distributed ML, and MLoNs

· provide a thorough understanding of how such problems are solved, pros and cons of various approaches, and some experience in solving them

- review on recent topics in ML and MLoNs, including communication-efficiency, security, and MLoNs with partial knowledge
- give students the background and skills required to do research in this growing field

Course contents

- Lecture 1: Introduction
- Lecture 2: Centralized Convex ML
- Lecture 3: Centralized Nonconvex ML
- Lecture 4: Distributed ML
- Lecture 5: ADMM, guest lecturer
- Lecture 6: Communication Efficiency
- Lecture 7: Deep Neural Networks
- Lecture 8: Computer Assignment Session and Homework
- Lecture 9: Special Topic 1: Large-scale ML
- Lecture 10: Special Topic 2: Security in MLoNs
- Lecture 11: Special Topic 3: Online MLoNs
- Lecture 12: Special Topic 4: MLoNs with partial knowledge
- Lecture 13: Special Topic 5: Application Areas and Open Research Problems

Disposition

Lectures, homework problems, computer assignments, presentations on selected topics by the participants, and final project

Course literature

[1] Bubeck, Sébastien. "Convex optimization: Algorithms and complexity." Foundations and Trends in Machine Learning, vol. 8, no.3-4 (2015): 231-357.

[2] L. Bottou, F. Curtis, J. Norcedal, “Optimization Methods for Large-Scale Machine Learning”, SIAM Rev., 60(2), 223–311.

[3] Boyd, Stephen, et al. "Distributed optimization and statistical learning via the alternating direction method of multipliers." Foundations and Trends in Machine learning 3.1 (2011): 1-122.

[4] Goodfellow, Y. Bengio, A. Courville, “Deep Learning”, MIT press 2016

[5] Jordan, Michael I., Jason D. Lee, and Yun Yang. "Communication-efficient distributed statistical inference," Journal of the American Statistical Association, 2018.

[6] Smith, Virginia, et al. "CoCoA: A general framework for communication-efficient distributed optimization." Journal of Machine Learning Research 18 (2018): 230.

[7] Alistarh, Dan, et al. "QSGD: Communication-efficient SGD via gradient quantization and encoding." Advances in Neural Information Processing Systems. 2017.

[8] Schmidt, Mark, Nicolas Le Roux, and Francis Bach. "Minimizing finite sums with the stochastic average gradient." Mathematical Programming 162.1-2 (2017): 83-112.

[9] Boyd, Stephen, et al. "Randomized gossip algorithms," IEEE Transactions on Information Theory, 2006.

[10] Scaman, Kevin, et al. "Optimal algorithms for smooth and strongly convex distributed optimization in networks," ICML, 2017.

Equipment

Personal laptop/computer

Examination

- EXA1 - Examination, 10.0 credits, grading scale: P, F

Based on recommendation from KTH’s coordinator for disabilities, the examiner will decide how to adapt an examination for students with documented disability.

The examiner may apply another examination format when re-examining individual students.

If the course is discontinued, students may request to be examined during the following two academic years.

Other requirements for final grade

- Attending at least 11 lectures (out of 13)
- 45 min oral presentation of a selected topic in one of the Special Topic lectures
- 80% on homework problems and computer assignments
- Project (preferably on a problem related to the student’s own research)

Ethical approach

- All members of a group are responsible for the group's work.
- In any assessment, every student shall honestly disclose any help received and sources used.
- In an oral assessment, every student shall be able to present and answer questions about the entire assignment and solution.