Syllabus
Graphical Models and Bayesian Learning, 6hp

Issued by the WASP graduate school management group

Main field of study
AI/math, AI/mlx

Course level
Advanced course for PhD students

- AI track: elective
- AS track: elective
- Joint curriculum: advanced

Course offered for
PhD Students in the WASP graduate school

Entry requirements
The participants are assumed to have a background in mathematics corresponding to the contents of the WASP-course “Introduction to Mathematics for Machine Learning".

The course requires basic understanding of probability theory and graph theory. In particular, the participants need to be comfortable with reasoning about (in-)dependence of random variables and factorizing joint probability distributions/density functions using the chain rule of probability, as well as reasoning about basic properties of graphs.

Module 2 requires familiarity with Bayesian statistics (as well as graphical models and basic message-passing algorithms, covered in module 1).

Intended learning outcomes
After completing the course, the students will have functional knowledge of the theory of graphical models, graphical model structure learning methods as well as methods for Bayesian inference with graphical models. Specifically, students should be able to

- Distinguish between the use-cases of different classes of graphical models (undirected and directed acyclic graphical models).
- Solve typical graphical model structure learning problems.
- Formulate the Bayesian inference problem for graphical models and distinguish between scenarios in which this problem can be solved exactly and when approximate inference algorithms are needed.
- Solve typical Bayesian inference problems for graphical models.
- Compare and contrast between different graphical model learning algorithms.

Course content
In module 1 we will introduce undirected graphical and directed acyclic graphical models and explore how these models encode independencies between variables in complex systems via the graph structure. Methods for learning graphical representations of the independences in a data-generating distribution will be discussed, along with exact algorithms for the inference of posterior probabilities.

In module 2 we consider Bayesian inference in graphical models where the involved distributions and/or the structure of the graph prevents exact solutions. We introduce several approximate inference algorithms to tackle this issue, such as expectation propagation, variational inference, and Markov chain Monte Carlo. We highlight some key differences, but also similarities, between these types of algorithms.

Teaching and working methods

The course includes two 3-day meetings with intense teaching on-site, typically a mixture of lectures and exercises.

Module 1 will consist of lectures, exercise sessions, and homework problem sets. The exercise sessions will follow the lectures, giving the students an opportunity to work with the newly introduced theory.

Module 2 will consist of lectures and practical sessions. The latter will be centered around implementing and evaluating approximate inference algorithms using numerical software (preferred language is Python).

Examination

Participants in the course will be assessed using hand-in assignments. To pass each individual module the student must receive a passing grade for their work on the module's hand-in assignments and have actively participated in the in-person meetings for both modules. To pass the course the student must pass both modules.

A re-examination is prepared about 6 months after the course covering completion of missing parts.

Grades

Fail or Pass