Syllabus
Planning & Relational Learning, 6hp

Issued by the WASP graduate school management group 2025-04-07.

Main Field of Study
AI/mlx, AI/math, AS, software

Course Level
Advanced course for PhD students

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

Course offered for
PhD Students in the WASP graduate school

Entry Requirements
The course requires solid programming experience in a high-level language. The participants are assumed to have a basic background in propositional logic and first-order logic.

Intended Learning Outcomes
On completion of the course the students should be able to

- explain and analyze state-of-the-art planning methods
- implement planning algorithms for solving sequential decision-making problems
- model simple planning problems in a suitable symbolic representation
- use basic methods of inductive logic programming for learning and inference
- apply the concept of "probabilities on domains" to learn a parametrized graphical model from a relational structure
- apply the concept of "probabilities on possible worlds" with parametrized graphical model to make inferences on arbitrary domains.

Course Content
The course provides an introduction to two key fields in Artificial Intelligence: (1) symbolic sequential decision-making and (2) statistical relational learning. The course is structured into independent modules, with each field introduced separately before participants explore their integration in a final group project.

In the sequential decision-making module, the focus is on planning techniques. Participants will explore common planning formalisms and learn how heuristic search can be used to solve planning problems efficiently. The course also covers state-space reduction techniques to optimize planning performance and introduces alternative methods for tackling planning tasks. Through lectures and practical exercises, attendees will develop a solid foundation in AI-driven planning approaches.

The statistical relational learning module introduces the fundamentals of inductive logic programming and presents two key approaches for unifying logic and probability: the "probabilities on domains" approach and the "probabilities on possible worlds" approach. Participants will learn how the first approach enables the construction of graphical models from relational data and how the second, combined with a parameterized graphical model, supports inference over arbitrary finite domains. The course also addresses the challenge of handling large domains without excessive computational costs.

In the final module, planning serves as an application domain for statistical relational learning. Given a set of training instances and their corresponding sequential solutions, participants will use inductive logic programming to characterize the solutions, extracting domain knowledge that can improve planning efficiency. This will be implemented as a collaborative group project, allowing students to apply their knowledge in a practical setting.

Teaching and Working Method

The course consists of three 2-day meetings, one of which is held in Uppsala, the other two in Linköping. During all sessions, there will be lectures that cover the course content. In addition, there are interactive exercise sessions, where students collaborate on solving assignments. In the last block, there will be a joint seminar in which students present their approach and results from a group project.

Examination

One individual hand-in assignment per 2-day meeting and a group project that is presented in a seminar.

Grades

Fail or Pass