IA Planning
Course Syllabus

Halim Djerroud

revision 1.0

Introduction

The **Task Planning** module is organized into **15 sessions**, each consisting of **30 minutes of lecture** followed by **1h30 of lab work**. The goal is to understand the **theoretical foundations of AI planning**, practice **modeling with PDDL**, and apply this knowledge in a **final project**.

The assessment consists of a **project** (done in pairs) and a **written exam**. **Continuous assessment** is ensured through selected lab assignments.

Final Grade = 0.5 × Project Grade + 0.3 × Written Exam + 0.2 × Continuous Assessment.

- **Project**: development of a mini planning system for an application scenario (logistics, robotics, or game).
- **Written Exam**: theoretical and practical notions covered in lectures/labs.
- **Continuous Assessment**: 2 selected lab assignments will be graded.

Objectives

- Understand the foundations of AI planning (states, actions, goals).
- Master problem modeling in PDDL.
- Discover main planning algorithms (progression, regression, heuristics, Graphplan, HTN, etc.).
- Apply these concepts to concrete cases (robotics, logistics, games).

Prerequisites

- Basic programming (Python or C++).
- Basic knowledge in algorithms (graphs, search).

Workload distribution

- **15 sessions of 2h** ⇒ 30 min lecture + 1h30 lab
- **Last 5 sessions** ⇒ Final project

Suggested Bibliography

- Ghallab, Nau, Traverso, *Task Planning: Theory and Practice.*
- Malik Ghallab et al., *PDDL Manual.*
- Recommended reading: Blum & Furst, Graphplan (1997).
- Online resources: https://planning.wiki

Lecture 1: Introduction to Planning

- **Lecture 30 min   Lab 1h30**

1. Motivation and applications (robotics, logistics, games)
2. Definitions: states, actions, goals
3. Concrete examples: 8-puzzle, robot grid navigation
4. **Lab**: first state search problems in Python (BFS/DFS) and A star (A*) algorithm

Lecture 2: STRIPS and Modeling Basics

- **Lecture 30 min   Lab 1h30**

1. STRIPS: preconditions, effects
2. Block world example
3. **Lab**: modeling a simple STRIPS problem

Lecture 3: PDDL — Domains and Problems

- **Lecture 30 min   Lab 1h30**

1. PDDL syntax: `domain.pddl`, `problem.pddl` files
2. Classic examples
3. **Lab**: writing your first PDDL problem

Lecture 4: Forward and Backward Search

- **Lecture 30 min   Lab 1h30**

1. Planning by progression and regression
2. Comparison with graph search
3. Limits: combinatorial explosion, complexity
4. **Lab**: mini Python solver (progression)

Lecture 5: Heuristic Planning and Graphplan

- **Lecture 30 min   Lab 1h30**

1. Heuristics, Graphplan
2. Existing solvers
3. **Lab**: using Fast Downward

Lecture 6: Planning as Reduction (SATPlan, CSP)

- **Lecture 30 min   Lab 1h30**

1. Reduction to SAT/CSP
2. Use cases
3. **Lab**: solving via a SAT solver

Lecture 7: Temporal and Resource-based Planning

- **Lecture 30 min   Lab 1h30**

1. Durative actions, limited resources
2. Temporal extensions of PDDL
3. **Lab**: modeling a task schedule

Lecture 8: Planning under Uncertainty

- **Lecture 30 min   Lab 1h30**

1. MDP, POMDP
2. Link with temporal planning (uncertainty in durations, probabilistic resources)
3. Link with reinforcement learning
4. **Lab**: stochastic grid simulation

Lecture 9: Hierarchical Task Network (HTN) Planning

- **Lecture 30 min   Lab 1h30**

1. Hierarchical decomposition
2. Example: travel organization
3. **Lab**: simple HTN implementation

Lecture 10: Multi-Agent Planning and Robotics

- **Lecture 30 min   Lab 1h30**

1. Multi-agent coordination
2. Robotic applications
3. **Lab**: delivery robot simulation

Lecture 11–15: Final Project

- **Lecture 30 min   Lab 1h30**

1. Topic of choice (logistics, robotics, game)
2. Goals: PDDL modeling, solver usage, evaluation
3. Work in pairs, final submission and presentation