Lecture slides for

Automated Planning: Theory and Practice

### Chapter 1 Introduction

Dana S. Nau

### CMSC 722, AI Planning University of Maryland, Fall 2004

### **Textbook**

 M. Ghallab, D. Nau, and P. Traverso *Automated Planning: Theory and Practice*  Morgan Kaufmann Publishers May 2004 ISBN 1-55860-856-7

• Web site: http://www.laas.fr/planning



### **Class Background**

- Things I hope you know
  - Asymptotic complexity, e.g.,  $O(n^3)$ ,  $\Theta(n^5)$ ,  $\Omega(n^2)$ 
    - » Differences in meaning: worst case, best case, average case
    - » Nondeterministic algorithms
    - » P, NP, NP-completeness, NP-hardness
  - Search algorithms
    - » Depth-first, breadth-first, best-first search
    - » A\*, admissible versus admissible heuristics
    - » State-space versus problem reduction (AND/OR graphs)
  - Logic
    - » Propositional logic
    - » First-order logic (predicates and quantifiers)
    - » Horn clauses, Horn-clause theorem proving

Dana Nau: Lecture slides for Automated Planning

### **Plans and Planning**

• Plan:

 A collection of actions for performing some task or achieving some objective

• Planning:

- There are many programs to aid human planners
  - » Project management, plan storage/retrieval, automatic schedule generation
- Automatic plan generation is much more difficult
   Many research prototypes, fewer practical systems
  - » Research is starting to pay off
    - Several successes on difficult practical problems





### **NASA Unmanned Spacecraft**

- Remote Agent eXperiment (RAX)
  - autonomous AI software for planning/control
  - ran on the DS1 spacecraft in May 1998
  - For several days it was allowed to control the spacecraft



### • Mars rover

guided by autonomous AI planning/control software

Dana Nau: Lecture slides for Automated Planning Licensed under the Creative Commons Attribution-NonCommercial-ShareAlike License: http://creativecommons.org/licenses/by-nc-sa/2.0/

### **Other Examples**

- Computer bridge: *Bridge Baron* 
  - Used AI planning to win the 1997 world computer bridge championship
  - Commercial software, thousands of copies sold



- Manufacturing process planning
  - Software included with Amada's sheet-metal bending machines
  - Used to plan bending operations

Dana Nau: Lecture slides for *Automated Planning* Licensed under the Creative Commons Attribution-NonCommercial-ShareAli



### Outline

- Conceptual model
- Restrictive assumptions
- Classical planning
- Relaxing the assumptions
- A running example: Dock Worker Robots

![](_page_7_Figure_0.jpeg)

- Model of the environment: *possible states*
- Model of how the environment can change: *effects of actions*
- Specification of *initial conditions* and *objectives*
- Plans of actions that are generated by a planner
- A model of execution of a plan in the environment
- A model of observation of the environment

Dana Nau: Lecture slides for Automated Planning

### **Conceptual Model**

- State-transition system  $\Sigma = (S, A, E, \gamma)$ 
  - $S = \{\text{states}\}$
  - $A = \{actions\}$  (controllable)
  - $E = \{\text{events}\}$  (uncontrollable)
  - state-transition function  $\gamma: S \times (A \cup E) \rightarrow 2^S$
- Observation function  $h: S \rightarrow O$ 
  - produces observation *o* about current state *s*
- Controller: given observation *o* in *O*, produces action *a* in *A*
- Planner:
  - input: description of  $\Sigma$ , initial state  $s_0$  in S, some objective
  - output: produces a plan to drive the controller

Dana Nau: Lecture slides for Automated Planning

![](_page_8_Figure_14.jpeg)

### **Conceptual Model**

- Possible objectives:
  - A set of goal states S<sub>g</sub>
     Find sequence of state transitions ending at a goal
  - Some condition over the set of states followed by the system
     *e.g.*, reach S<sub>g</sub> and stay there
  - Utility function attached to states
     Optimize some function of the utilities
  - Tasks to perform, specified recursively as sets of sub-tasks and actions

![](_page_9_Figure_6.jpeg)

## **Conceptual Model: Example**

- State transition system  $\Sigma = (S, A, E, \gamma)$ 
  - $\diamond S = \{\mathbf{s}_0, \dots, \mathbf{s}_5\}$
  - A = {move1, move2, put, take, load, unload}
  - $\bullet E = \{\}$
  - γ: as shown
- h(s) = s for every s
- Input to planner:
  - system  $\Sigma$
  - initial state  $s_0$
  - goal state  $s_5$

![](_page_10_Figure_11.jpeg)

Dana Nau: Lecture slides for Automated Planning

### **Planning Versus Scheduling**

- Scheduling
  - Decide how to perform a given set of actions using a limited number of resources in a limited amount of time

![](_page_11_Figure_3.jpeg)

- Typically NP-complete
- Planning
  - Decide what actions to use to achieve some set of objectives
  - Can be much worse than NP-complete
    - » In the most general case, it is undecidable
    - » Most research assumes various collections of restrictions to guarantee decidability
  - We'll now look at some of the restrictions

## Restrictive Assumptions

A0 (finite Σ):
the state space S is finite
S = {s<sub>0</sub>, s<sub>1</sub>, s<sub>2</sub>, ... s<sub>k</sub>} for some k
A1 (fully observable Σ):
the observation function h: S → O is the identity function
i.e., the controller always knows what state Σ is in

![](_page_12_Figure_2.jpeg)

Dana Nau: Lecture slides for Automated Planning

## Restrictive Assumptions

- A2 (deterministic  $\Sigma$ ):
  - for all u in  $A \cup E$ ,  $|\gamma(s, u)| = 1$
  - each action or event has only one possible outcome
- A3 (static  $\Sigma$ ):
  - *E* is empty: no changes except those performed by the controller
- A4 (attainment goals):
   a goal state s<sub>g</sub> or a set of goal states S<sub>g</sub>

![](_page_13_Figure_7.jpeg)

## Restrictive Assumptions

A5 (sequential plans):
 solution is a linearly ordered sequence of actions (a1, a2, ... an)

### • A6 (implicit time):

- no durations,
   instantaneous state-transitions
- A7 (off-line planning):
   planner doesn't know the execution status

![](_page_14_Figure_5.jpeg)

### **Classical Planning**

• Classical planning requires all eight restrictive assumptions

- complete knowledge about a deterministic, static, finite-state system with attainment goals and implicit time
- Reduces to the following problem:
  - Given (Σ, s<sub>0</sub>, S<sub>g</sub>), find a sequence of actions (a<sub>1</sub>, a<sub>2</sub>, ... a<sub>n</sub>) that produces a sequence of state transitions

$$s_{1} = \gamma(s_{0}, a_{1}),$$
  

$$s_{2} = \gamma(s_{1}, a_{2}),$$
  

$$\dots,$$
  

$$s_{n} = \gamma(s_{n-1}, a_{n})$$
  
such that  $s_{n}$  is in  $S_{g}$ .

# **Classical Planning: Example**

- Same example as before:
  - System is finite, deterministic, static
  - Complete knowledge
  - Attainment goals
  - Implicit time
  - Offline planning
- Classical planning is just path-searching in a graph
  - states are nodes
  - actions are edges
- Is this trivial?

![](_page_16_Figure_11.jpeg)

Dana Nau: Lecture slides for Automated Planning

### **Classical Planning**

- Very difficult computationally
  - Generalize the earlier example:
    - » five locations, three piles, three robots, 100 containers
  - Then there are  $10^{277}$  states

![](_page_17_Picture_5.jpeg)

- » More than 10<sup>190</sup> times as many states as the number of particles in the universe!
- The vast majority of AI research has been on classical planning
   Parts I and II of the book
- Too restricted to fit most problems of practical interest
   But the ideas can sometimes be useful in those problems

• Relax A0 (finite  $\Sigma$ ):

- ◆ Discrete, *e.g.* 1st-order logic:
- Continuous, *e.g.* numeric variables
- Sections:
  - » 2.4 (extensions to classical)
  - » 10.5 (control-rule planners)
  - » 11.7 (HTN planning)
- Case study: Chapter 21 (manufacturability analysis)
- Relax A1 (fully observable Σ):
  - If we don't relax any other restrictions, then the only uncertainty is about s<sub>0</sub>

![](_page_18_Figure_11.jpeg)

- Relax A2 (deterministic Σ):
  - Actions have more than one possible outcome
  - Seek policy or contingency plan
  - With probabilities:
    - »Discrete Markov Decision Processes (MDPs)
    - »Chapter 11
  - Without probabilities:
    - »Nondeterministic transition systems
    - »Chapters 12, 18

![](_page_19_Figure_10.jpeg)

• Relax A1 and A2: Finite POMDPs »Plan over *belief states* »Exponential time & space »Section 16.3 • Relax A0 and A2: Continuous or hybrid MDPs »Control theory (see engineering courses) • Relax A0, A1, and A2 Continuous or hybrid POMDPs »Case study: Chapter 20 (robotics)

![](_page_20_Figure_2.jpeg)

- Relax A3 (static  $\Sigma$ ):
  - Other agents or dynamic environment
    - » Finite perfect-info zero-sum games (introductory AI courses)
  - Randomly behaving environment
    - » Decision analysis (business, operations research)
    - » Can sometimes map this into MDPs or POMDPs
  - Case studies: Chapters 19 (space),
     22 (emergency evacuation)
- Relax A1 and A3
  - Imperfect-information games
  - Case study: Chapter 23 (bridge)

![](_page_21_Figure_11.jpeg)

#### Dana Nau: Lecture slides for Automated Planning

- Relax A5 (sequential plans) and A6 (implicit time):
  - Temporal planning
  - Chapters 13, 14
- Relax A0, A5, A5
  - Planning and resource scheduling
  - Chapter 15
- 247 other combinations
  - I won't discuss them!

![](_page_22_Figure_9.jpeg)

#### Dana Nau: Lecture slides for Automated Planning

### A running example: Dock Worker Robots

• Generalization of the earlier example

- A harbor with several locations
  - » e.g., docks, docked ships, storage areas, parking areas
- Containers
  - » going to/from ships
- Robot carts
  - » can move containers
- Cranes
  - » can load and unload containers

Dana Nau: Lecture slides for Automated Planning

### A running example: Dock Worker Robots

- Locations: **|1, |2,** ...
- Containers: c1, c2, ...
  - can be stacked in piles, loaded onto robots, or held by cranes

![](_page_24_Figure_4.jpeg)

- each belongs to a single location
- move containers between piles and robots
- if there is a pile at a location, there must also be a crane there

Dana Nau: Lecture slides for Automated Planning

### A running example: Dock Worker Robots

- Fixed relations: same in all states
   adjacent(l,l') attached(p,l) belong(k,l)
- Dynamic relations: differ from one state to another

![](_page_25_Figure_3.jpeg)

Dana Nau: Lecture slides for Automated Planning