

CSC2542

Introduction to Planning

Sheila McIlraith and Rick Valenzano

Department of Computer Science

University of Toronto

Fall 2016

Administrative Announcements

- Room change (and it might change again)
- Poll on new tutorial time
- Suggested readings for next week:
 - Hoffmann paper (more on this from Rick)
- Slides will be posted by the weekend
- Assignment description will be provided next week
- Send us email regarding papers or topics you'd like to cover in class (csc2542profs@cs.toronto.edu)

Acknowledgements

Some of the slides used in this course are modifications of Dana Nau's lecture slides for the textbook *Automated Planning*, licensed under the Creative Commons Attribution-NonCommercial-ShareAlike License:

<http://creativecommons.org/licenses/by-nc-sa/2.0/>

Other slides are modifications of slides developed by Malte Helmert, Bernhard Nebel, and Jussi Rintanen.

I have also used some material prepared by P@trick Haslum.

I would like to gratefully acknowledge the contributions of these researchers, and thank them for generously permitting me to use aspects of their presentation material.

plan *n.*

1. A scheme, program, or method worked out beforehand for the accomplishment of an objective: *a plan of attack.*
 2. A proposed or tentative project or course of action: *had no plans for the evening.*
 3. A systematic arrangement of elements or important parts; a configuration or outline: *a seating plan; the plan of a story.*
 4. A drawing or diagram made to scale showing the structure or arrangement of something.
 5. In perspective rendering, one of several imaginary planes perpendicular to the line of vision between the viewer and the object being depicted.
 6. A program or policy stipulating a service or benefit: *a pension plan.*
- Synonyms:** *blueprint, design, project, scheme, strategy*

Dictionary.com/plan

Address: <http://dictionary.reference.com/search?q=plan>

Dictionary.com Search

Dictionary Thesaurus Web

[Home](#) Premium: [Sign up](#) | [Login](#)

plan *n.*

1. A scheme, program, or method worked out beforehand for the accomplishment of an objective: *a plan of attack.*
 2. A proposed or tentative project or course of action: *had no plans for the evening.*
 3. A systematic arrangement of elements or important parts; a configuration or outline: *a seating plan; the plan of a story.*
 4. A drawing or diagram made to scale showing the structure or arrangement of something.
 5. In perspective rendering, one of several imaginary planes perpendicular to the line of vision between the viewer and the object being depicted.
 6. A program or policy stipulating a service or benefit: *a pension plan.*
- Synonyms:** *blueprint, design, project, scheme, strategy*

[a representation] of future behavior ...
 usually a set of actions, with temporal and
 other constraints on them, for execution by
 some agent or agents.

- Austin Tate

[MIT Encyclopedia of the Cognitive Sciences, 1999]

```

00T V01 20 4.8 01Tbd tmen M1
00A E1 00 2.2 01Peelanbar dsch adwad
02Dyboadn o a 5dg.F
00B E1300 0.8 01Sep
02Spadphrest 6m1000RMspner
00C E1300 2.0 01Sep
02Pbblthgap of p oesst
   usg p o o in ealges
00D E1300 0.0 01Sep
00T E1900 5.7 02Et hngo fopp
01Tbd tmen E
00A M1300 4.8 01Sep
02Pearebad 6 sdera
00B M
  
```

A portion of a manufacturing process plan

02Scenpitsodr eponbad

02Cup oad
 03Estblshdtupoinatblse(08,100)
 demilig 00
 (-08,125
 ,dph050
 at(025 18)
 ,dph050
 (-08,300
 ,dph050
 at(025 30)
 ,dph050
 dmilingtol

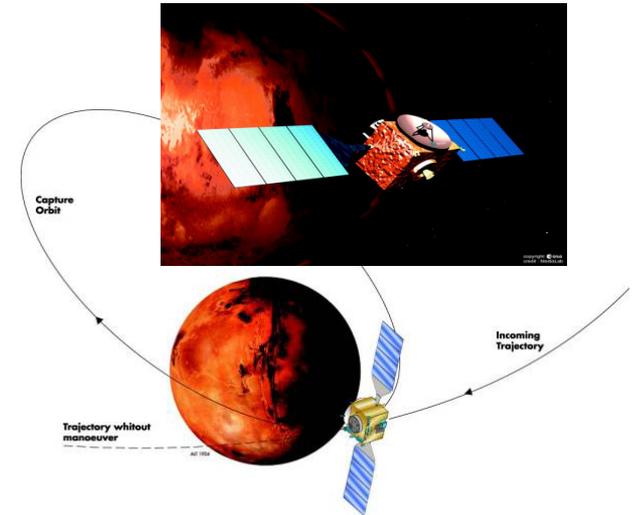
Modes of Planning

- Mixed Initiative Planning
- Automated Plan Generation

Example Planning Applications

Autonomous Agents for Space Exploration

- Autonomous planning, scheduling, control
 - NASA: JPL and Ames
- Remote Agent Experiment (RAX)
 - Deep Space 1
- Mars Exploration Rover (MER)



Other Autonomous Systems

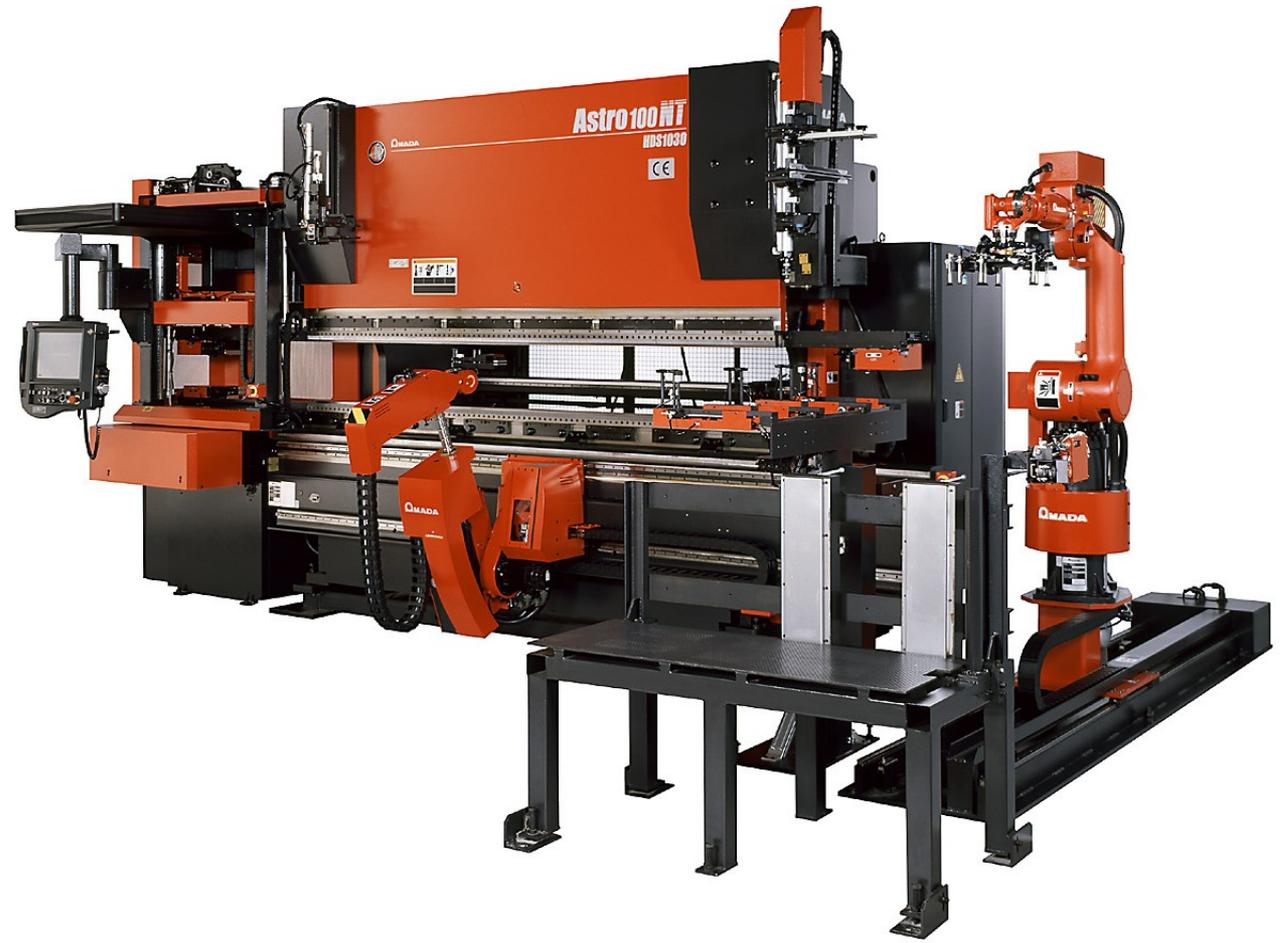


Not necessarily embodied! →



Manufacturing Automation

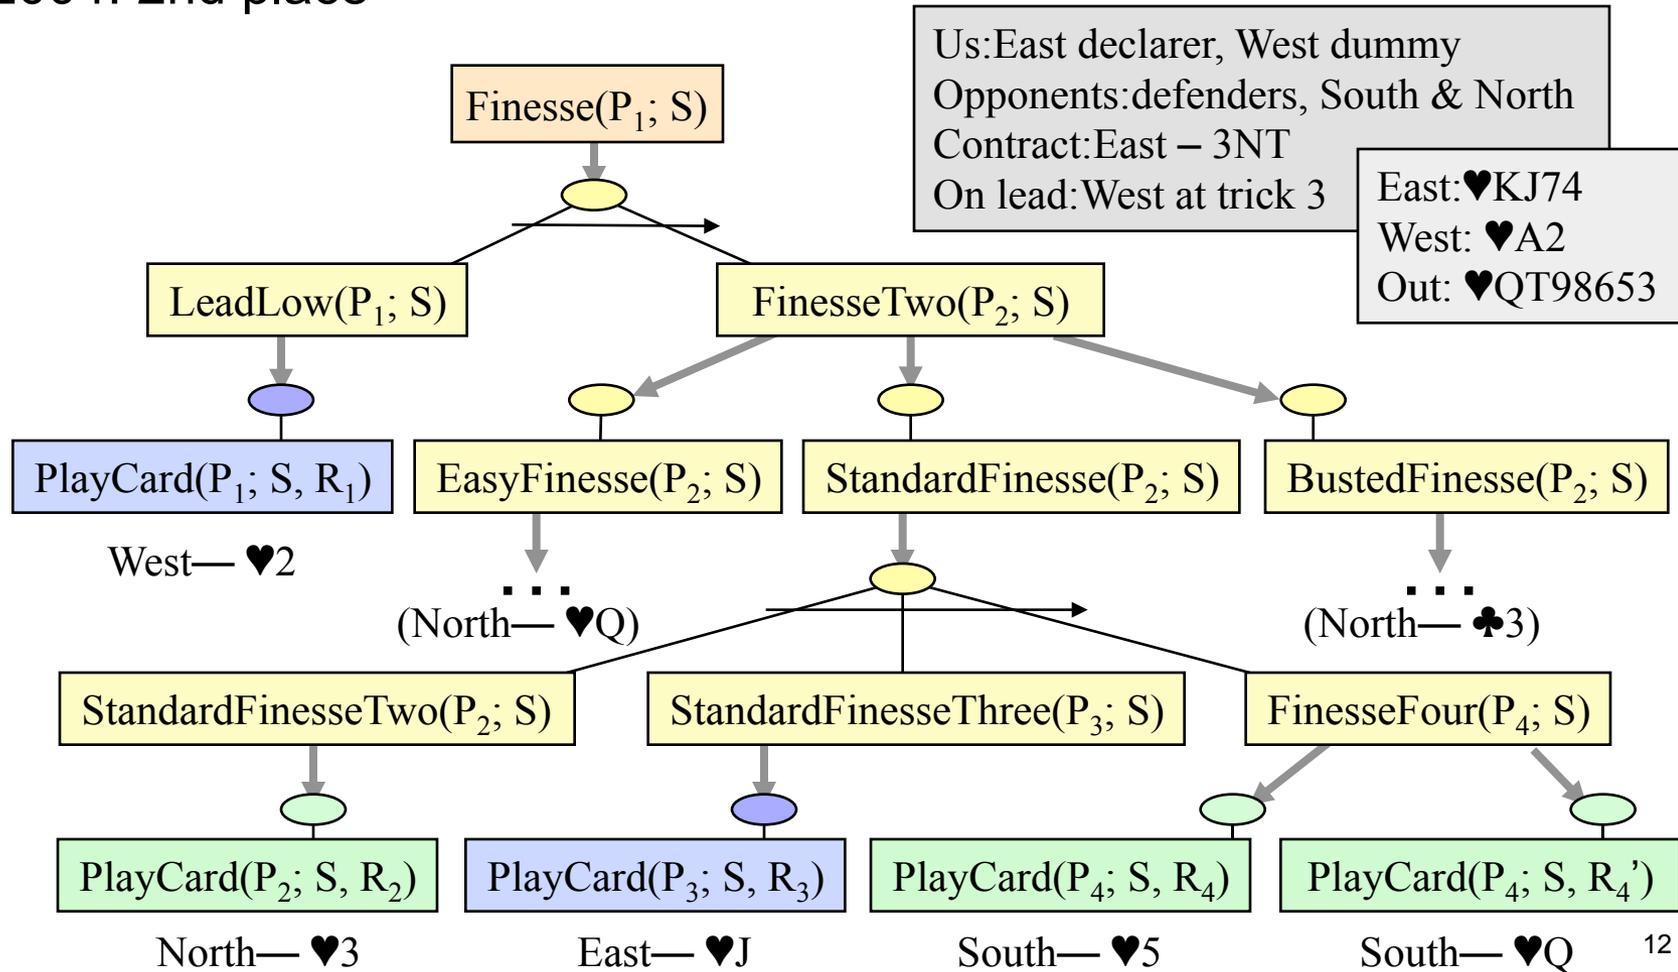
- Sheet-metal bending machines - Amada Corporation
 - Software to plan the sequence of bends
[Gupta and Bourne, *J. Manufacturing Sci. and Engr.*, 1999]



Games

E.g., *Bridge Baron* - Great Game Products

- 1997 world champion of computer bridge
[Smith, Nau, and Throop, *AI Magazine*, 1998]
- 2004: 2nd place



Other Applications

- Scheduling with Action Choices & Resource Requirements
 - Problems in supply chain management
 - HSTS (Hubble Space Telescope scheduler)
 - Workflow management
- Air Traffic Control
 - Route aircraft between runways and terminals. Crafts must be kept safely separated. Safe distance depends on craft and mode of transport. Minimize taxi and wait time.
- Character Animation
 - Generate step-by-step character behaviour from high-level spec
- Plan-based Interfaces
 - Dialogue management (plan a dialogue to convey something)
 - Plan recognition

Other Applications (cont.)

- **Web Service Composition**
 - Compose web services, and monitor their execution
 - Many of the web standards have a lot of connections to action representation languages
 - BPEL; BPEL-4WS allow workflow specifications
 - DAML-S allows process specifications
- **Business Process Composition /Workflow Management**
 - Including Grid Services/Scientific Workflow Management
- **Genome Rearrangement**
 - The relationship between different organisms can be measured by the number of “evolution events” (rearrangements) that separate their genomes
 - Find shortest (or most likely) sequence of rearrangements between a pair of genomes

Other Applications (cont.)

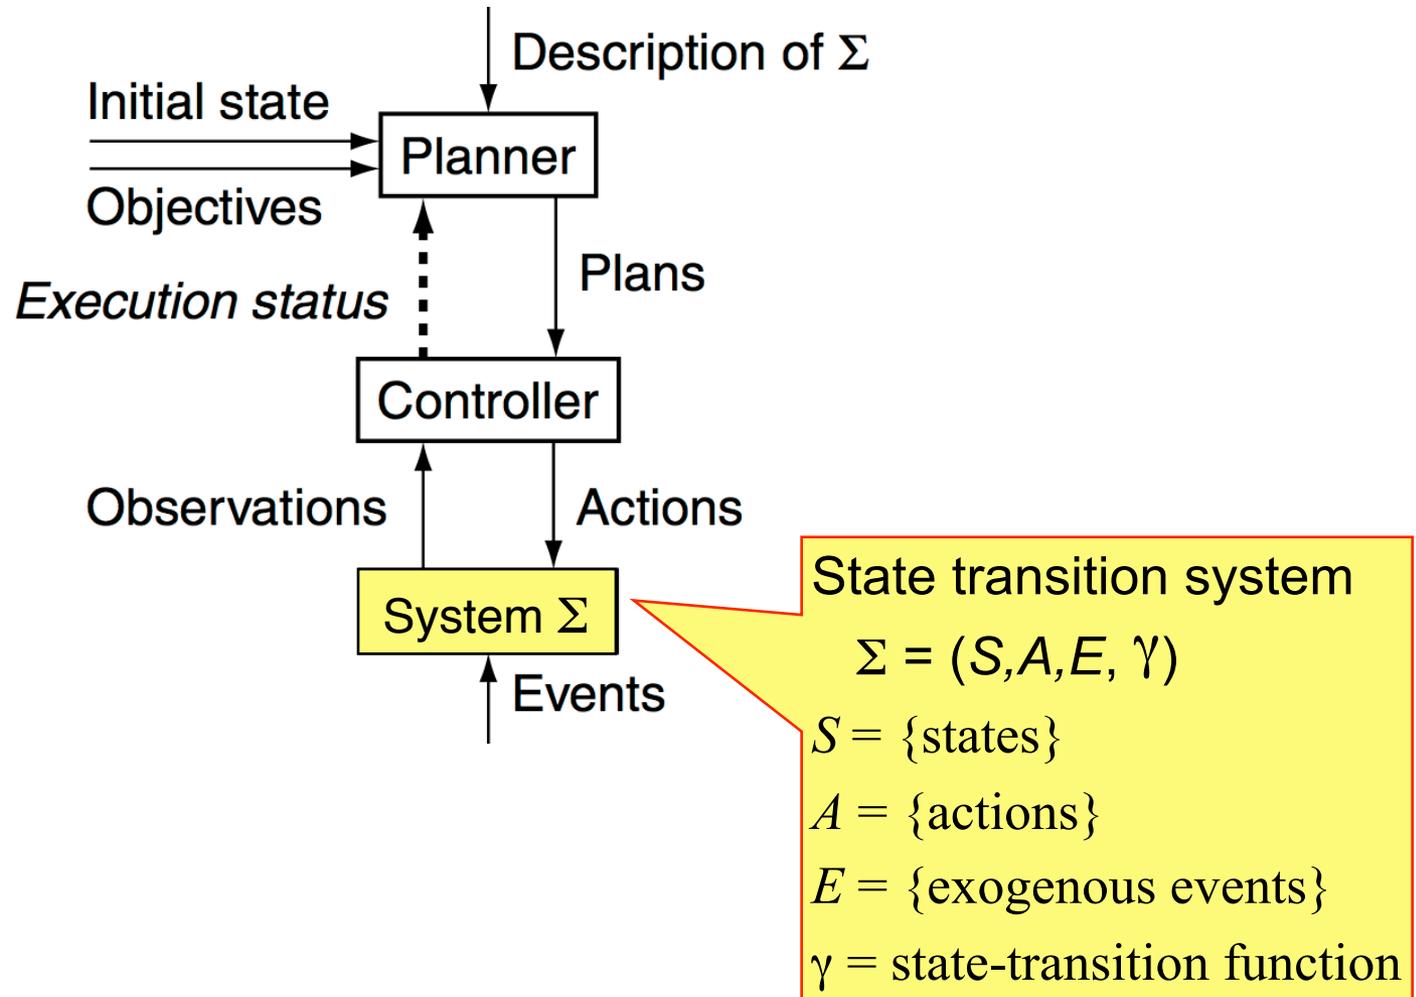
- Narrative generation
- Narrative understanding
- Software/Program synthesis
- Automated diagnosis
- Intelligent tutoring systems
- ...

Outline

- Conceptual model for planning
- Classes of planning problems
- Classes of planners and example instances
- Beyond planning
- Planning research – the big picture
- Some of what I hope you'll get from the course

Conceptual Model

1. Environment



State Transition System

$$\Sigma = (S, A, E, \gamma)$$

- $S = \{\text{states}\}$
- $A = \{\text{actions}\}$
- $E = \{\text{exogenous events}\}$
- State-transition function

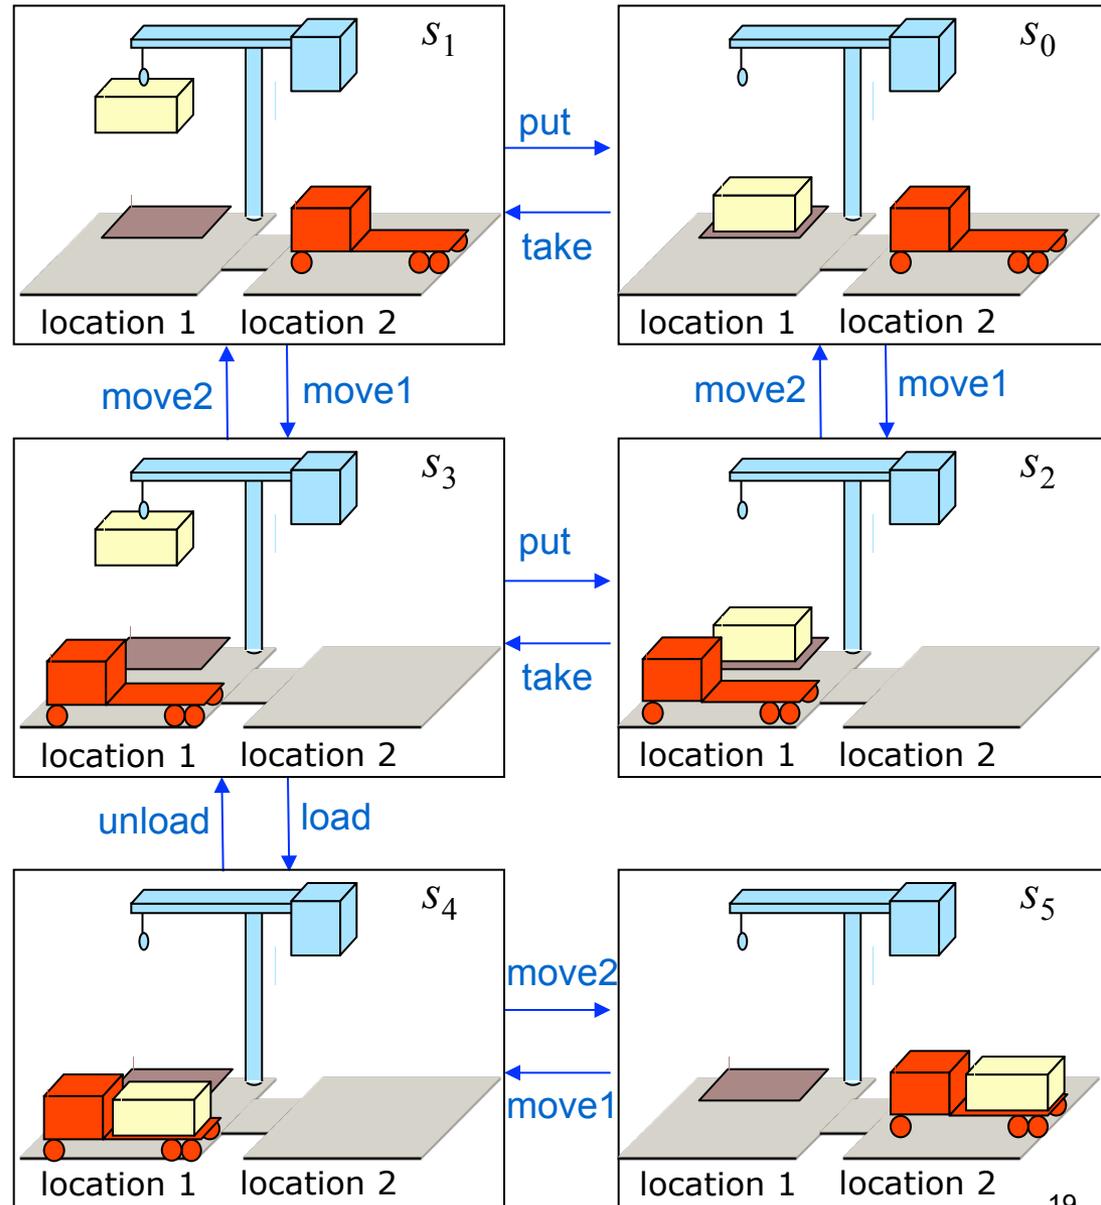
$$\gamma : S \times (A \cup E) \rightarrow 2^S$$

- $S = \{s_0, \dots, s_5\}$
- $A = \{\text{move1, move2, put, take, load, unload}\}$
- $E = \{\}$
- γ : see the arrows

State Transition System

$$\Sigma = (S, A, E, \gamma)$$

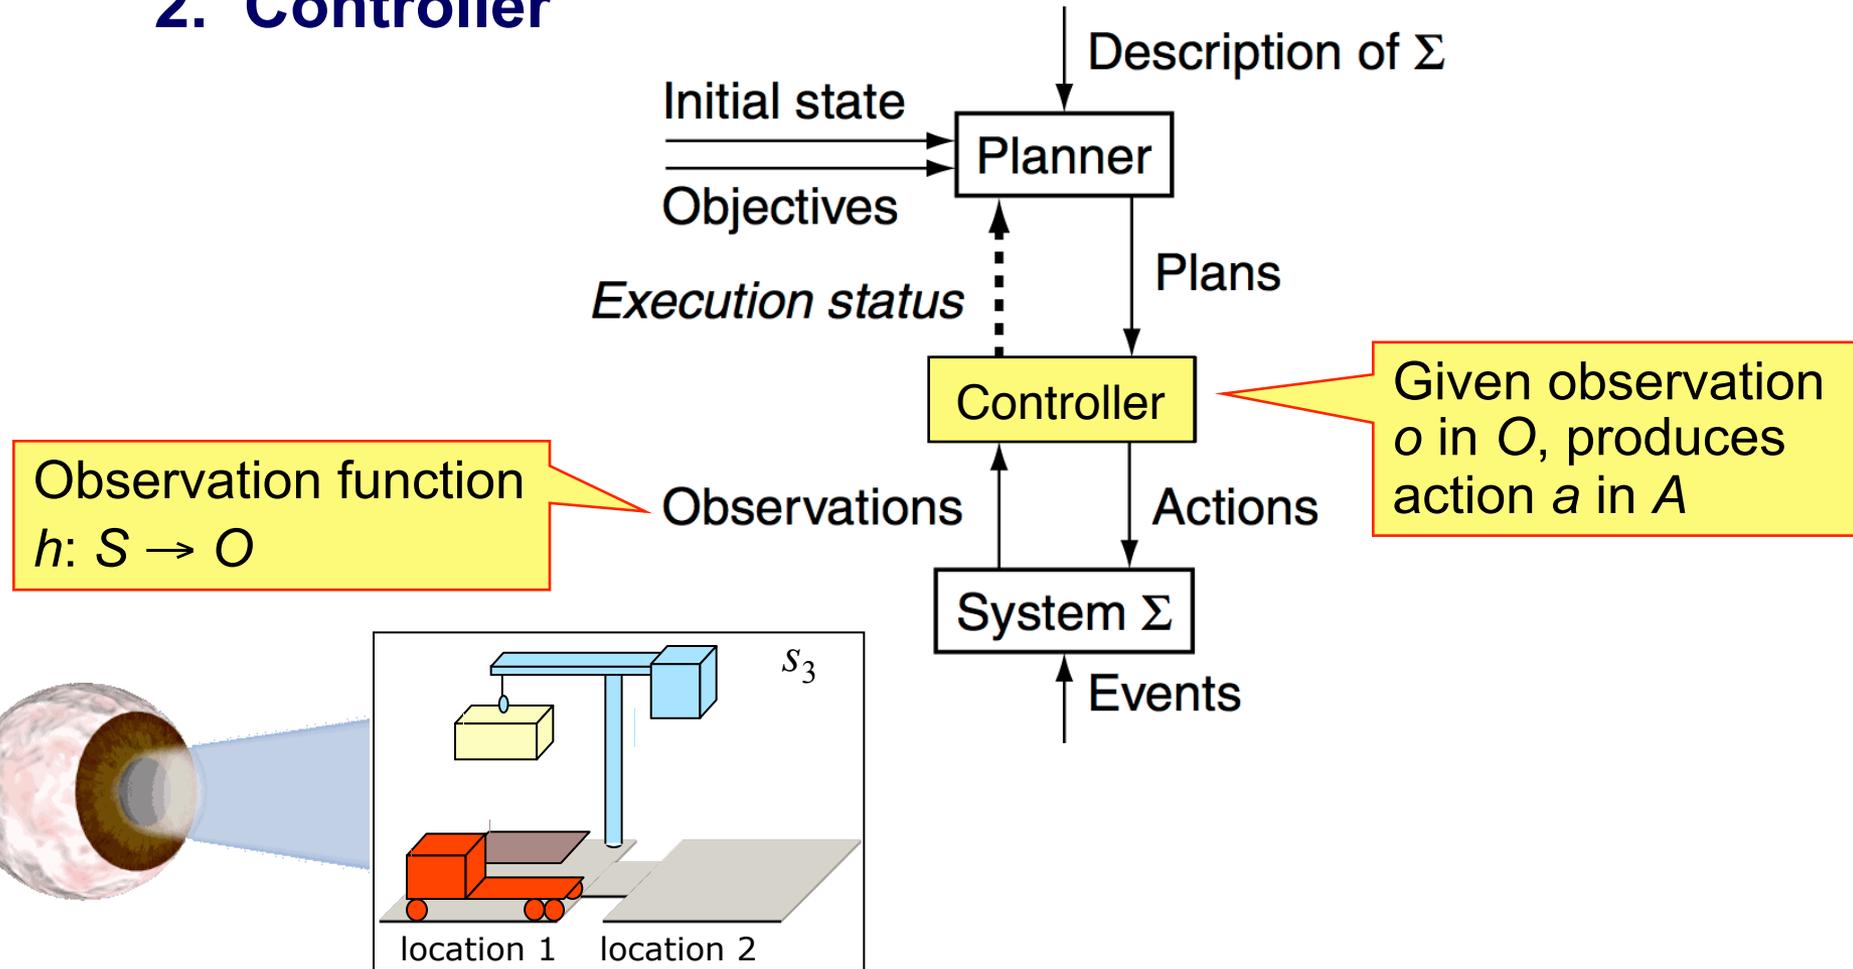
- $S = \{\text{states}\}$
- $A = \{\text{actions}\}$
- $E = \{\text{exogenous events}\}$
- State-transition function $\gamma : S \times (A \cup E) \rightarrow 2^S$
 - $S = \{s_0, \dots, s_5\}$
 - $A = \{\text{move1, move2, put, take, load, unload}\}$
 - $E = \{\}$
 - γ : see the arrows



Dock Worker Robots (DWR):

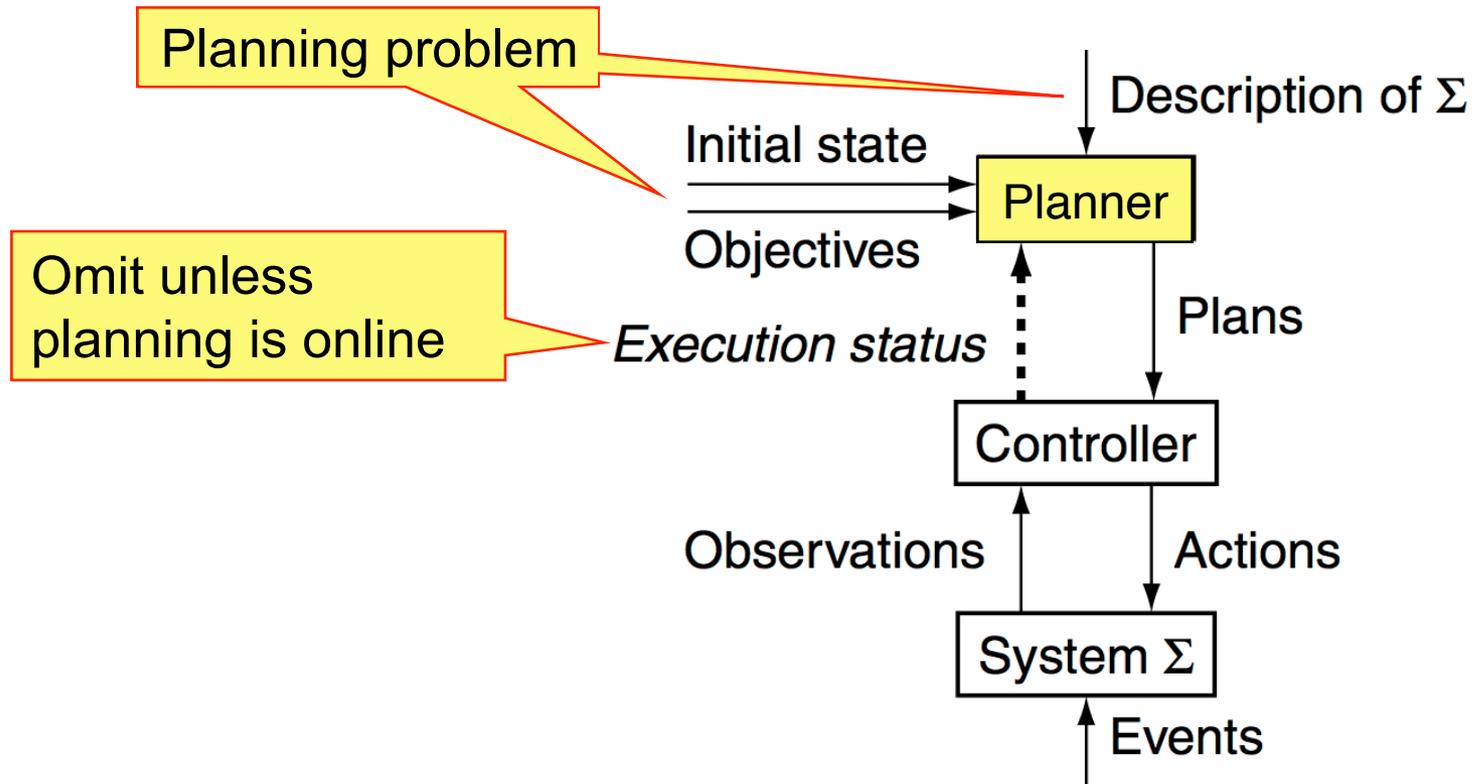
Conceptual Model

2. Controller



Conceptual Model

3. Planner's Input



Planning Problem

$$P = (\Sigma, s_0, G)$$

Σ : System Description

s_0 : Initial state(s)

E.g., Initial state = s_0

G : Objective

Goal state,

Set of goal states,

Set of tasks,

“trajectory” of states,

Objective function, ...

E.g., Goal state = s_5

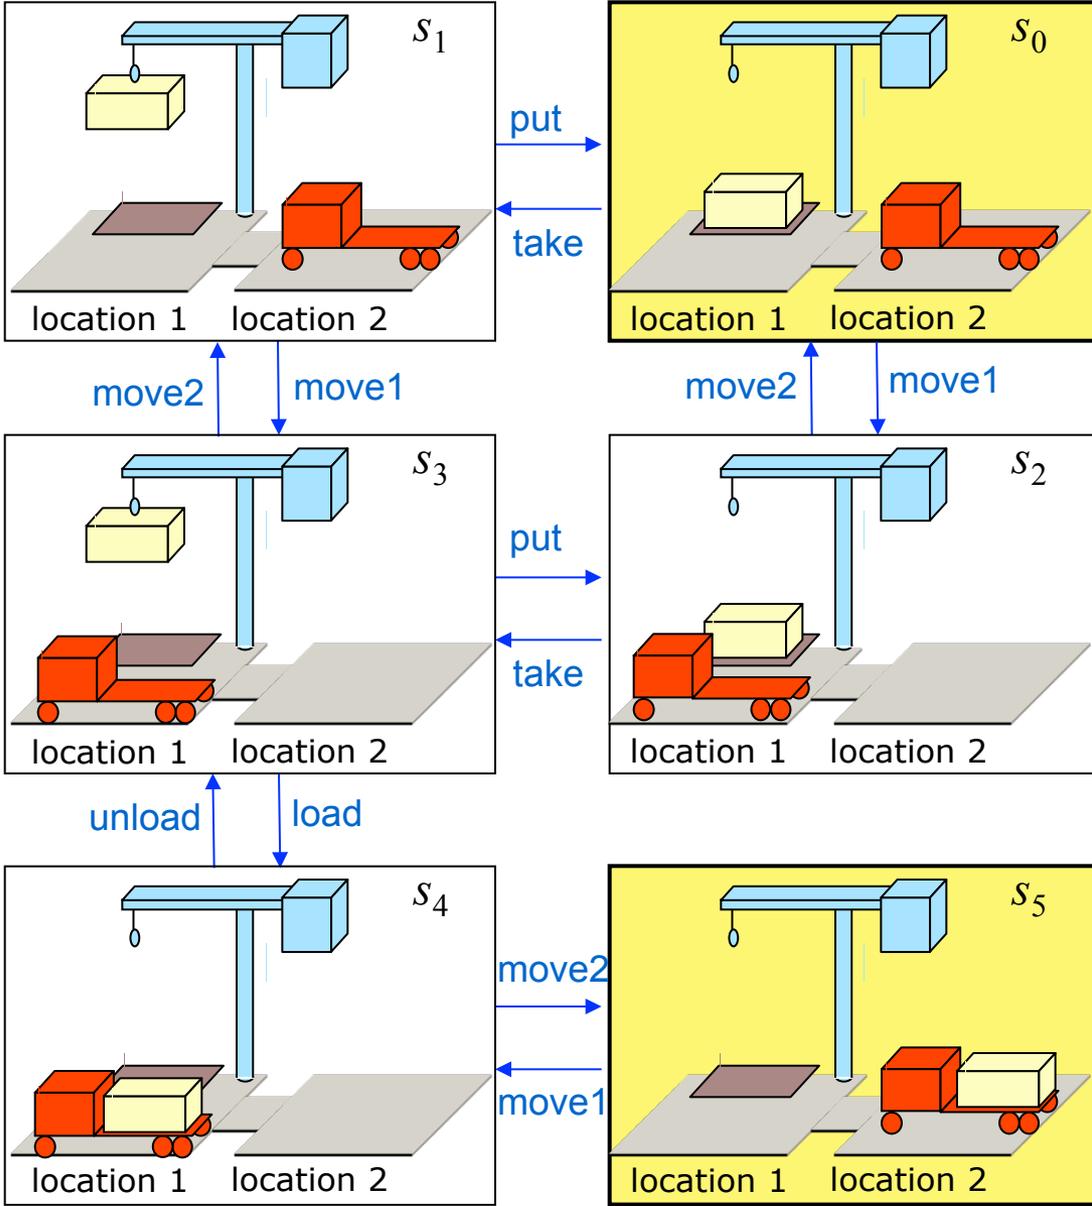
Planning Problem

$$P = (\Sigma, s_0, G)$$

Σ : System Description

s_0 : Initial state(s)
 E.g., Initial state = s_0

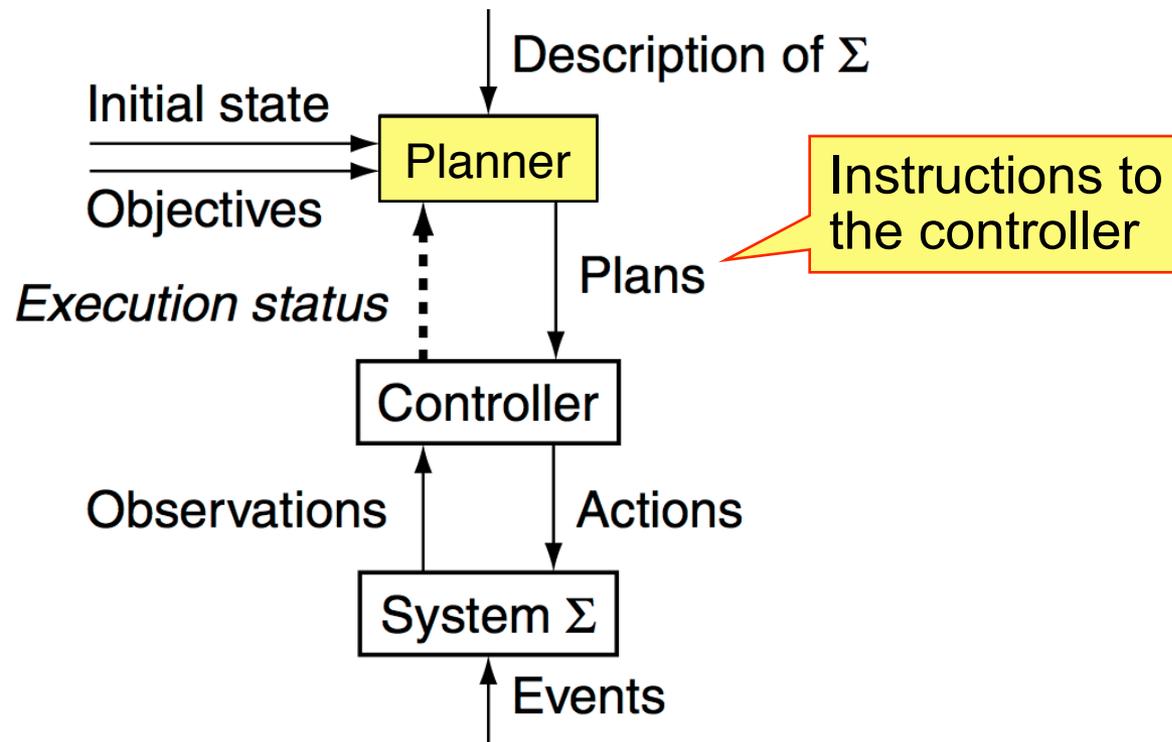
G : Objective
 Goal state,
 Set of goal states,
 Set of tasks,
 “trajectory” of states,
 Objective function, ...
 E.g., Goal state = s_5



The Dock Worker Robots (DWR) domain

Conceptual Model

4. Planner's Output



Plans

Classical plan:

a sequence of actions

E.g.,

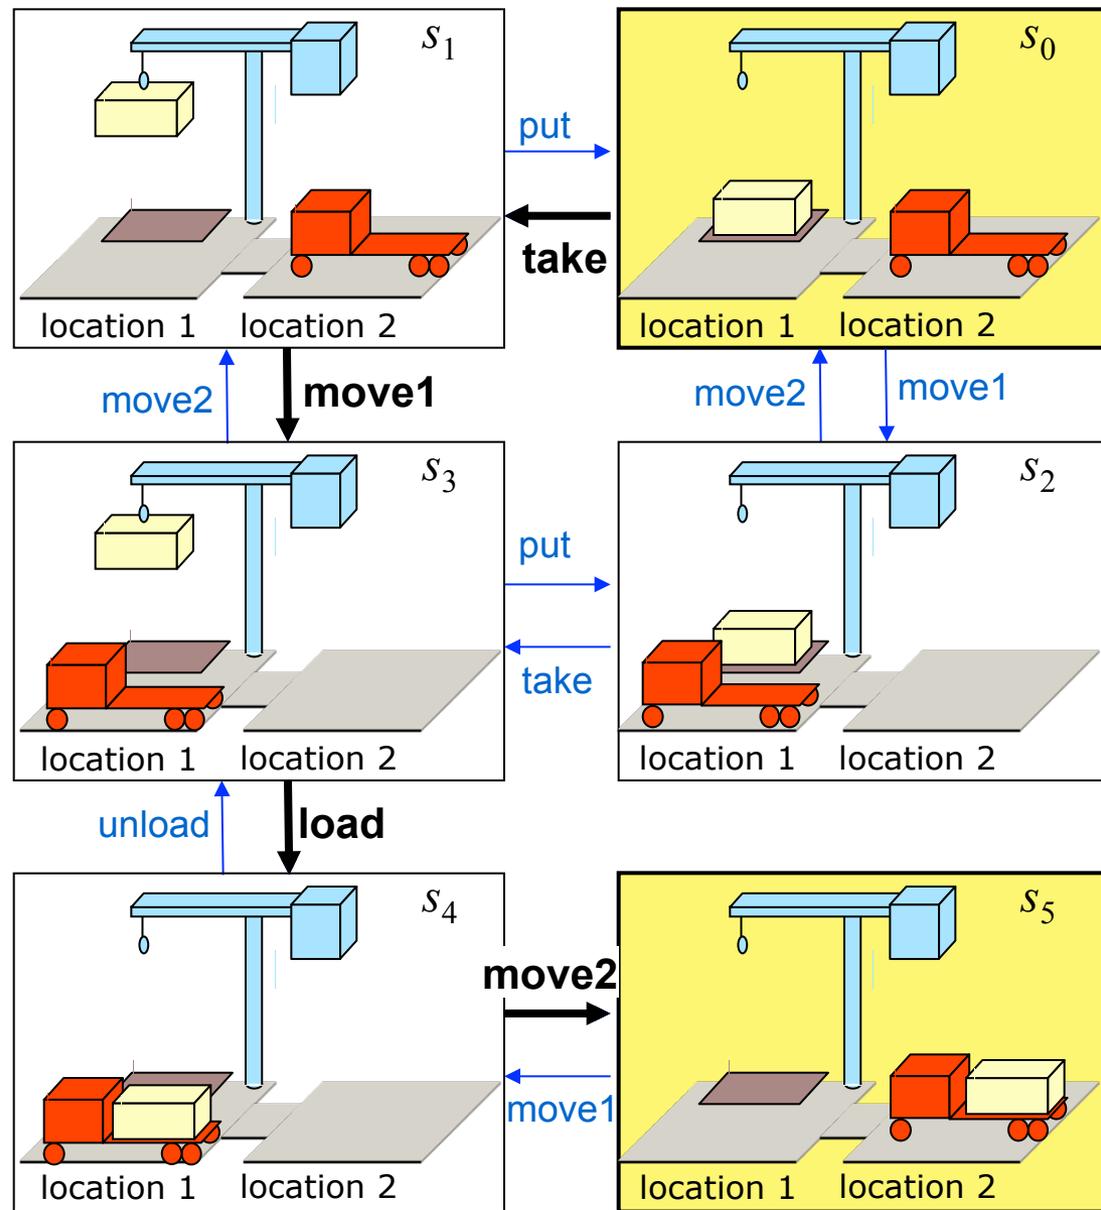
$\langle \text{take}, \text{move1}, \text{load}, \text{move2} \rangle$

Policy:

partial function from S into A

E.g.,

$\{(s_0, \text{take}),$
 $(s_1, \text{move1}),$
 $(s_3, \text{load}),$
 $(s_4, \text{move2})\}$



The Dock Worker Robots (DWR) domain

Outline

- Conceptual model for planning
- ➔ Classes of planning problems
- Classes of planners and example instances
- Beyond planning
- Planning research – the big picture
- Some of what I hope you'll get from the course

Planning

Agent: single agent or multi-agent

State:

complete or incomplete (logical/probabilistic)

state of the world and/or agent's state of knowledge

Actions:

world-altering and/or knowledge-altering (e.g. sensing)

deterministic or non-deterministic (logical/stochastic)

Goal Condition:

satisficing or optimizing

final-state or temporally extended/control knowledge/script

optimizing: preferences or cost or utility or ...

Reasoning:

offline or online (fully observable, partially observable)

Plans:

sequential, partial order, conformant, contingent, conditional
(controller or policy)

Different Classes Planning Problems

Varying components of the planning problem specification yields different classes of problems. E.g.,

dynamics: deterministic, nondeterministic, stochastic

observability: full, partial, none

horizon: finite, infinite

objective requirement: satisfying, optimizing

...

Different Classes Planning Problems

dynamics: **deterministic**, nondeterministic, stochastic

observability: full, partial, **none**

horizon: **finite**, infinite

objective requirement: **satisfying**, optimizing

...

- **classical planning**
- conditional planning with full observability (FOND)
- conditional planning with partial observability (POND)
- conformant planning
- markov decision processes (MDP)
- partial observable MDP (POMDP)
- preference-based/over-subscription planning

Different Classes Planning Problems

dynamics: deterministic, **nondeterministic**, stochastic

observability: **full**, partial, none

horizon: **finite**, **infinite**

objective requirement: **satisfying**, optimizing

...

- classical planning
- **conditional planning with full observability (FOND)**
- conditional planning with partial observability (POND)
- conformant planning
- markov decision processes (MDP)
- partial observable MDP (POMDP)
- preference-based/over-subscription planning

Different Classes Planning Problems

dynamics: deterministic, **nondeterministic**, stochastic

observability: full, **partial**, none

horizon: **finite**, **infinite**

objective requirement: **satisfying**, optimizing

...

- classical planning
- conditional planning with full observability
- **conditional planning with partial observability**
- conformant planning
- markov decision processes (MDP)
- partial observable MDP (POMDP)
- preference-based/over-subscription planning

Different Classes Planning Problems

dynamics: deterministic, **nondeterministic**, stochastic

observability: full, partial, **none**

horizon: **finite**, infinite

objective requirement: **satisfying**, optimizing

...

- classical planning
- conditional planning with full observability
- conditional planning with partial observability
- **conformant planning**
- markov decision processes (MDP)
- partial observable MDP (POMDP)
- preference-based/over-subscription planning

Different Classes Planning Problems

dynamics: deterministic, nondeterministic, **stochastic**

observability: **full**, partial, none

horizon: **finite**, **infinite**

objective requirement: satisfying, **optimizing**

...

- classical planning
- conditional planning with full observability
- conditional planning with partial observability
- conformant planning
- **markov decision processes (MDP)**
- partial observable MDP (POMDP)
- preference-based/over-subscription planning

Different Classes Planning Problems

dynamics: deterministic, nondeterministic, **stochastic**

observability: full, **partial**, none

horizon: **finite**, **infinite**

objective requirement: satisfying, **optimizing**

...

- classical planning
- conditional planning with full observability
- conditional planning with partial observability
- conformant planning
- markov decision processes (MDP)
- **partial observable MDP (POMDP)**
- preference-based/over-subscription planning

Different Classes Planning Problems

dynamics: **deterministic**, nondeterministic, stochastic

observability: full, partial, **none**

horizon: **finite**, **infinite**

objective requirement: satisfying, **optimizing**

...

- classical planning
- conditional planning with full observability
- conditional planning with partial observability
- conformant planning
- markov decision processes (MDP)
- partial observable MDP (POMDP)
- **preference-based/over-subscription planning**

Other Dimensions

dynamics: deterministic, nondeterministic, stochastic

: **explicit time, implicit time**

: **instantaneous, durative**

: **continuous, discrete, hybrid**

agents: multi-agent

perception: perfect, noisy

horizon: finite, infinite

objective requirement: satisfying, optimizing

objective form: final-state goal, temporally-extended goal, control knowledge, hierarchical task network (HTN), script/program (Golog)

plan form: sequential plan, partial order plan, controller, policy, generalized plan, program...

...

- classical planning
- conditional planning with full observability
- conditional planning with partial observability
- conformant planning
- markov decision processes (MDP)
- partial observable MDP (POMDP)
- preference-based/over-subscription planning

Other Dimensions

...and what do we do if we don't have a model?
We don't know the transition functions.
We don't know the reward we get from different actions.

Model-Free Planning

Reinforcement learning

Different Classes Planning Problems

Components of a Planning Problem: (Σ, s_0, G)

Varying these components defines different classes of planning problems and different planning output.

Dynamics: deterministic, nondeterministic, probabilistic

Observability: full, partial, none

Horizon: finite, infinite

Why is Planning Difficult?

- solutions to classical planning problems are **paths from an initial state to a goal state** in the transition graph
 - Efficiently solvable by Dijkstra's algorithm in $O(|V| \log |V| + |E|)$ time
 - Why don't we solve all planning problems this way?
- state space may be **huge**: 10^9 , 10^{12} , 10^{15} , ...states
- constructing the transition graph is infeasible!
- planning algorithms try to avoid constructing whole graph
- planning algorithms often are – but not guaranteed to be more efficient than obvious solution methods
constructing the transition graph and using e.g., Dijkstra's algorithm

Outline

- Conceptual model for planning
- Classes of planning problems
- ➔ Classes of planners and example instances
- Beyond planning
- Planning research – the big picture
- Some of what I hope you'll get from the course

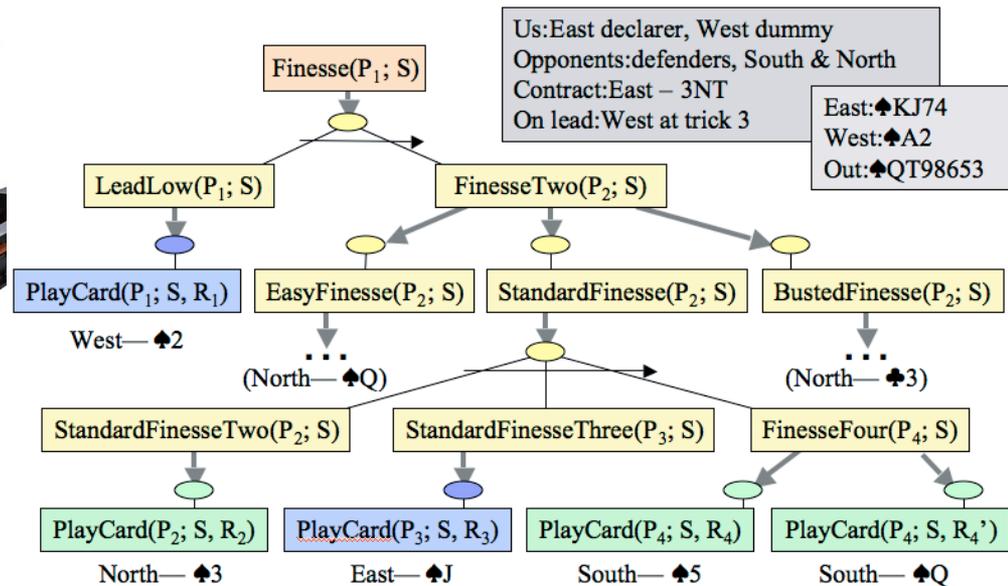
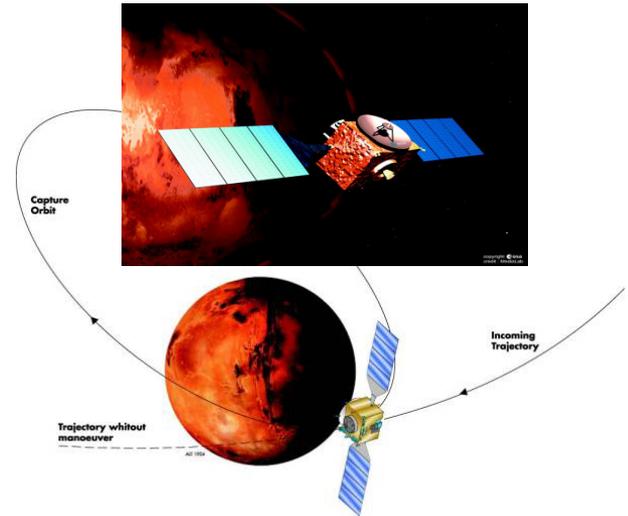
Three Main Classes of Planners

1. Domain-specific
2. Domain-independent
3. Domain-customizable

* *Ghallab, Nau, and Traverso's use "configurable" (which I don't like)
Also called "Domain-specific" or "Knowledge-Based"*

1. Domain-Specific Planners

- Made or tuned for specific domain
- Won't work well (if at all) in any other domain
- Many successful real-world planning systems work this way



2. Domain-Independent Planners

- In principle, a domain-independent planner works in any planning domain
- Uses no domain-specific knowledge except the definitions of the basic actions

2. Domain-Independent Planners

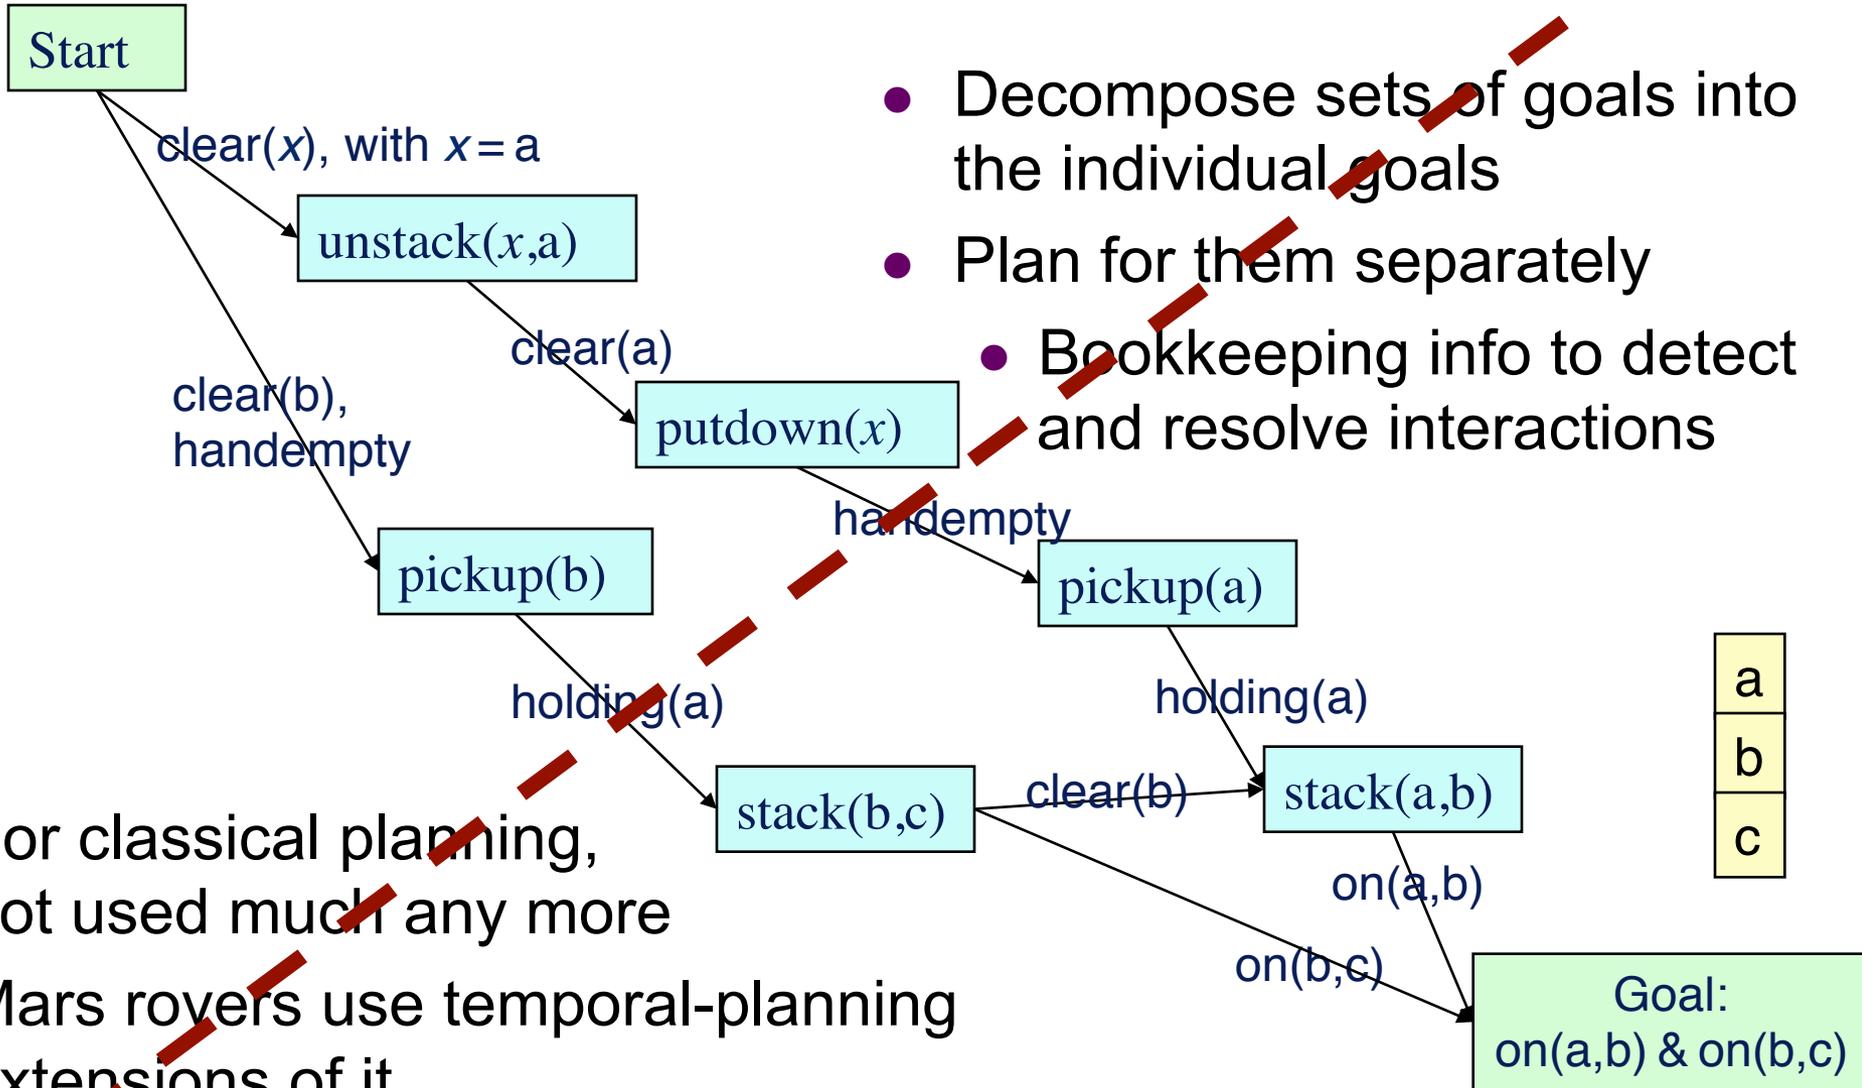
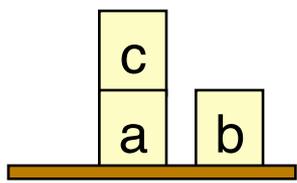
- In practice,
 - Not feasible to develop domain-independent planners that work in *every* possible domain
- Make simplifying assumptions to restrict the set of domains
 - *Classical planning*
 - Historical focus of most automated-planning research

Very active area of research. Many excellent planning systems.

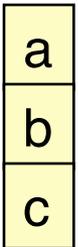
Restrictive Assumptions

- **A0: Finite system:**
 - finitely many states, actions, events
- **A1: Fully observable:**
 - the controller always knows the system's current state
- **A2: Deterministic:**
 - each action has only one outcome
- **A3: Static** (no exogenous events):
 - changes only occur as the result of the controller's actions
- **A4: Attainment goals:**
 - a set of goal states S_g
- **A5: Sequential plans:**
 - a plan is a linearly ordered sequence of actions $(a_1, a_2, \dots a_n)$
- **A6: Implicit time:**
 - Actions are instantaneous (have no duration)
- **A7: Off-line planning:**
 - planner doesn't know the execution status

Approach: Plan-Space Planning (Partial-Order Planning)



- Decompose sets of goals into the individual goals
- Plan for them separately
- Bookkeeping info to detect and resolve interactions



For classical planning,
not used much any more

Mars rovers use temporal-planning
extensions of it

Approach: Heuristic Search

The fastest state-of-the-art automated planning systems exploit Heuristic Search

Rick will be talking about this in the next hour

Approach: Translation to General Problem Solver

IDEA: Translate the planning problem into another kind of problem for which there are efficient solvers

- Find a solution to that problem
- Extract the plan from the solution, (sometimes w/ optimality guarantees)

Example “General Problem Solvers”

- SAT solvers
 - M [Rintanen]**, SATplan and Blackbox [Kautz & Selman]
- CSP solver
- Answer Set Programming (ASP) solvers
 - [Son *et al.*], [Lifschitz *et al.*], etc.
- Integer programming solvers such as Cplex
 - [Vossen *et al.*]
- MaxSAT

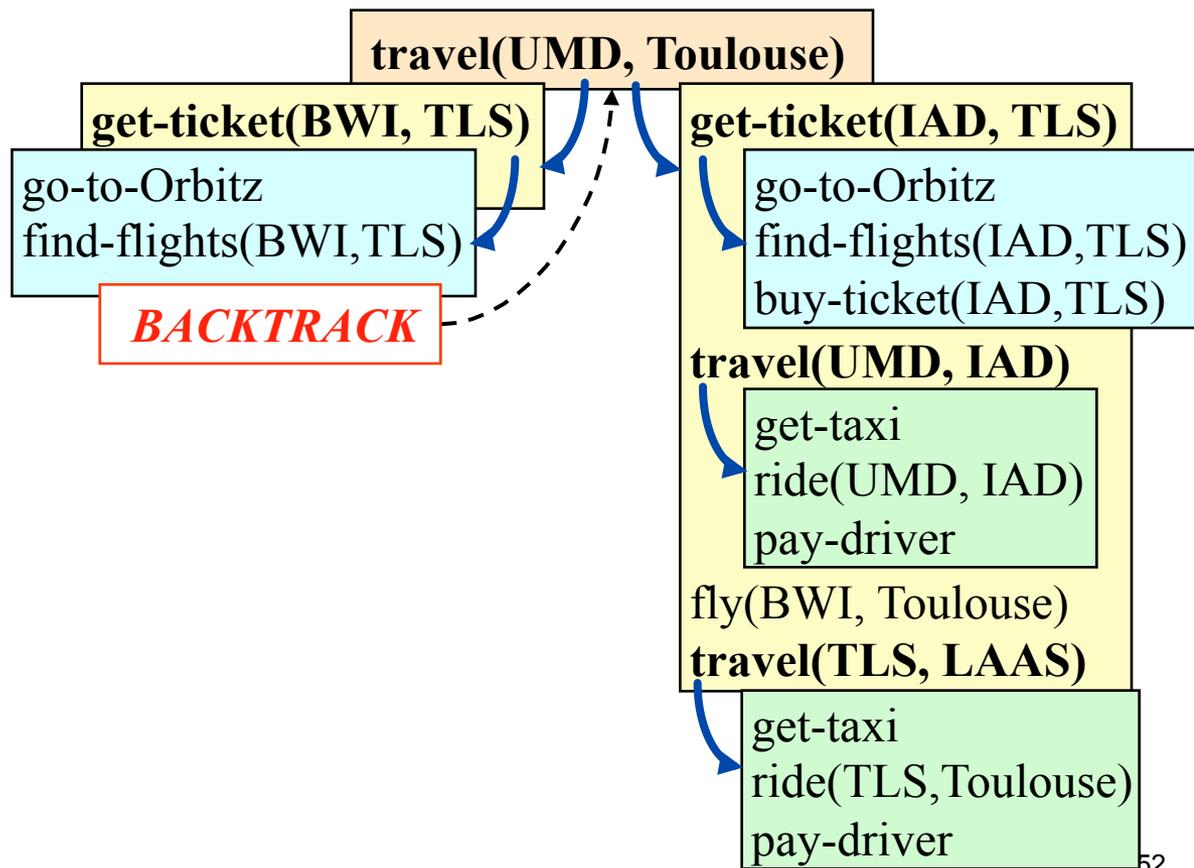
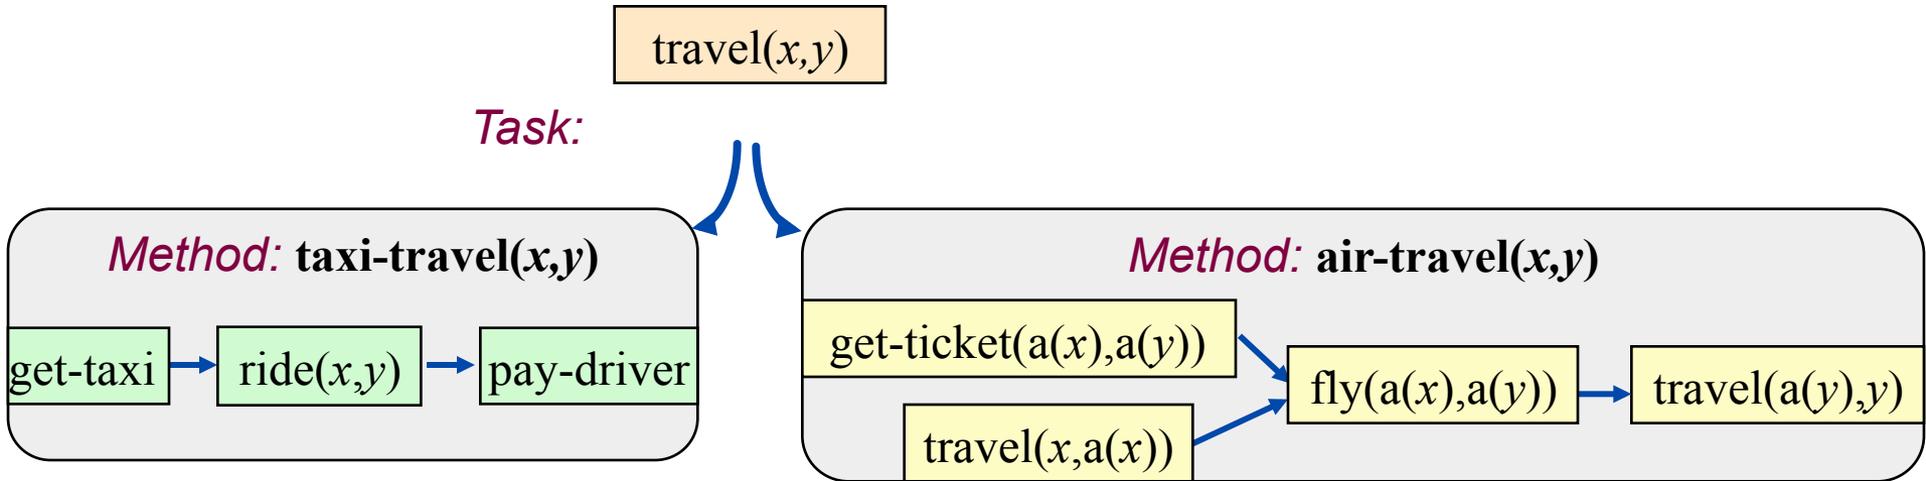
** competitive with state-of-the art heuristic search.

3. Domain-customizable

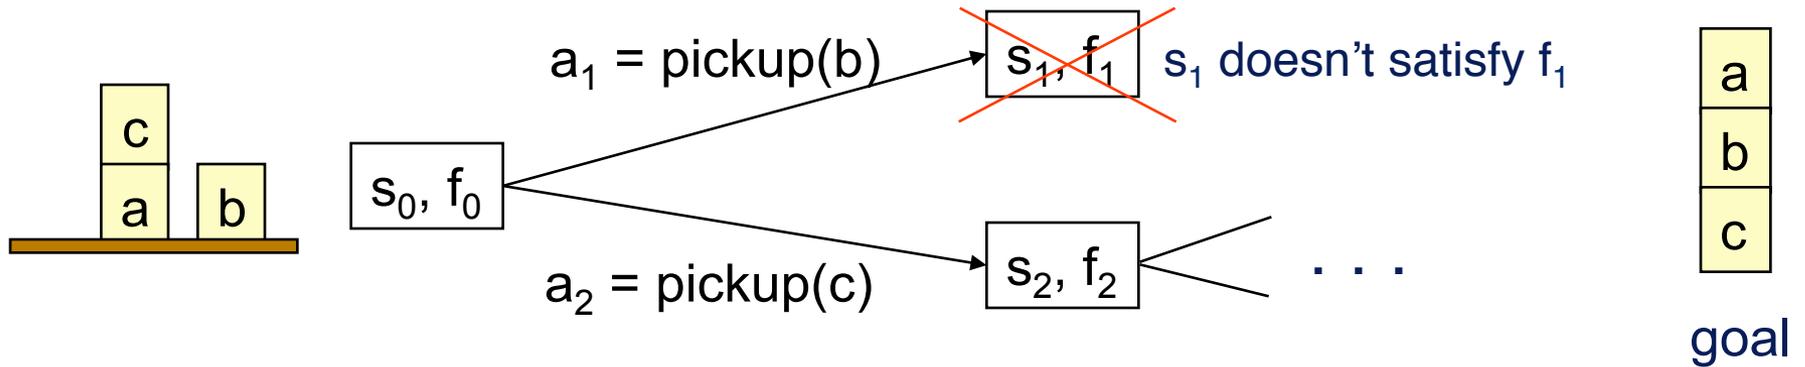
- Domain-independent planners are quite slow compared with domain-specific planners
 - Blocks world in linear time [Slaney and Thiébaux, *A.I.*, 2001]
 - Can get analogous results in many other domains
- But don't want to write a new planner for every domain!
- **Domain-customizable planners**
 - Domain-independent planning engine
 - Input (the “objective”) includes info about how to solve problems in the domain.
 - Hierarchical Task Network (HTN) planning
 - Planning with control formulas
 - Planning with a plan script or agent program

Approach: HTN Planning

- Problem reduction
 - *Tasks* (activities) rather than goals
 - *Methods* to decompose tasks into subtasks
 - Enforce constraints, backtrack if necessary
- Real-world applications
- Noah, Nonlin, O-Plan, SIPE, SIPE-2, SHOP, SHOP2



Approach: Planning with Control Formulas



- At each state s_i we have a *control formula* f_i in temporal logic

$$\text{ontable}(x) \wedge \neg \exists [y: \text{GOAL}(\text{on}(x, y))] \Rightarrow \text{O}(\neg \text{holding}(x))$$

“never pick up x from table unless x needs to be on another block”

- For each successor of s , derive a control formula using *logical progression*
- Prune any successor state in which the progressed formula is false
 - TLPlan [Bacchus & Kabanza]
 - TALplanner [Kvarnstrom & Doherty]

Approach: Planning w/ Program or Plan Script

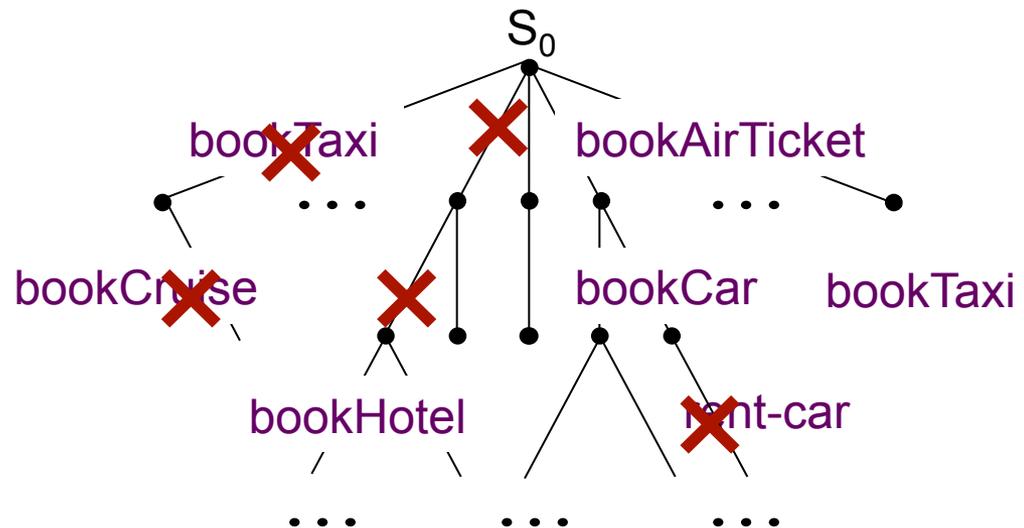
E.g., **Golog** [Levesque et al.]

Nondeterministic programs that act as procedural control knowledge, placing constraints on the valid action sequence/plans

E.g., *bookAirTicket(x); if far then bookCar(x) else bookTaxi(y)*

procedural constructs:

- sequence
- if-then-else
- nondeterministic choice
- while-do, etc.



Three Main Classes of Planners

1. Domain-specific
2. Domain-independent
E.g., Planning graph-based, SAT-based, heuristic search
3. Domain-customizable
E.g., HTN, domain control formula, agent programs/scripts

Outline

- Conceptual model for planning
- Classes of planning problems
- Classes of planners and example instances
- ➔ Beyond planning
- Planning research – the big picture
- Some of what I hope you'll get from the course

Broad Application of Planning Techniques

Planning algorithms are applicable to a broad range of applications that can roughly be viewed as reachability problems. E.g.,

- Software verification
- Diagnosis of dynamical systems
- Story understanding
- Situation assessment/Plan recognition
- Gene rearrangement
- ...

Outline

- Conceptual model for planning
- Classes of planning problems
- Classes of planners and example instances
- Beyond planning
- ➔ Planning research
- Some of what I hope you'll get from the course

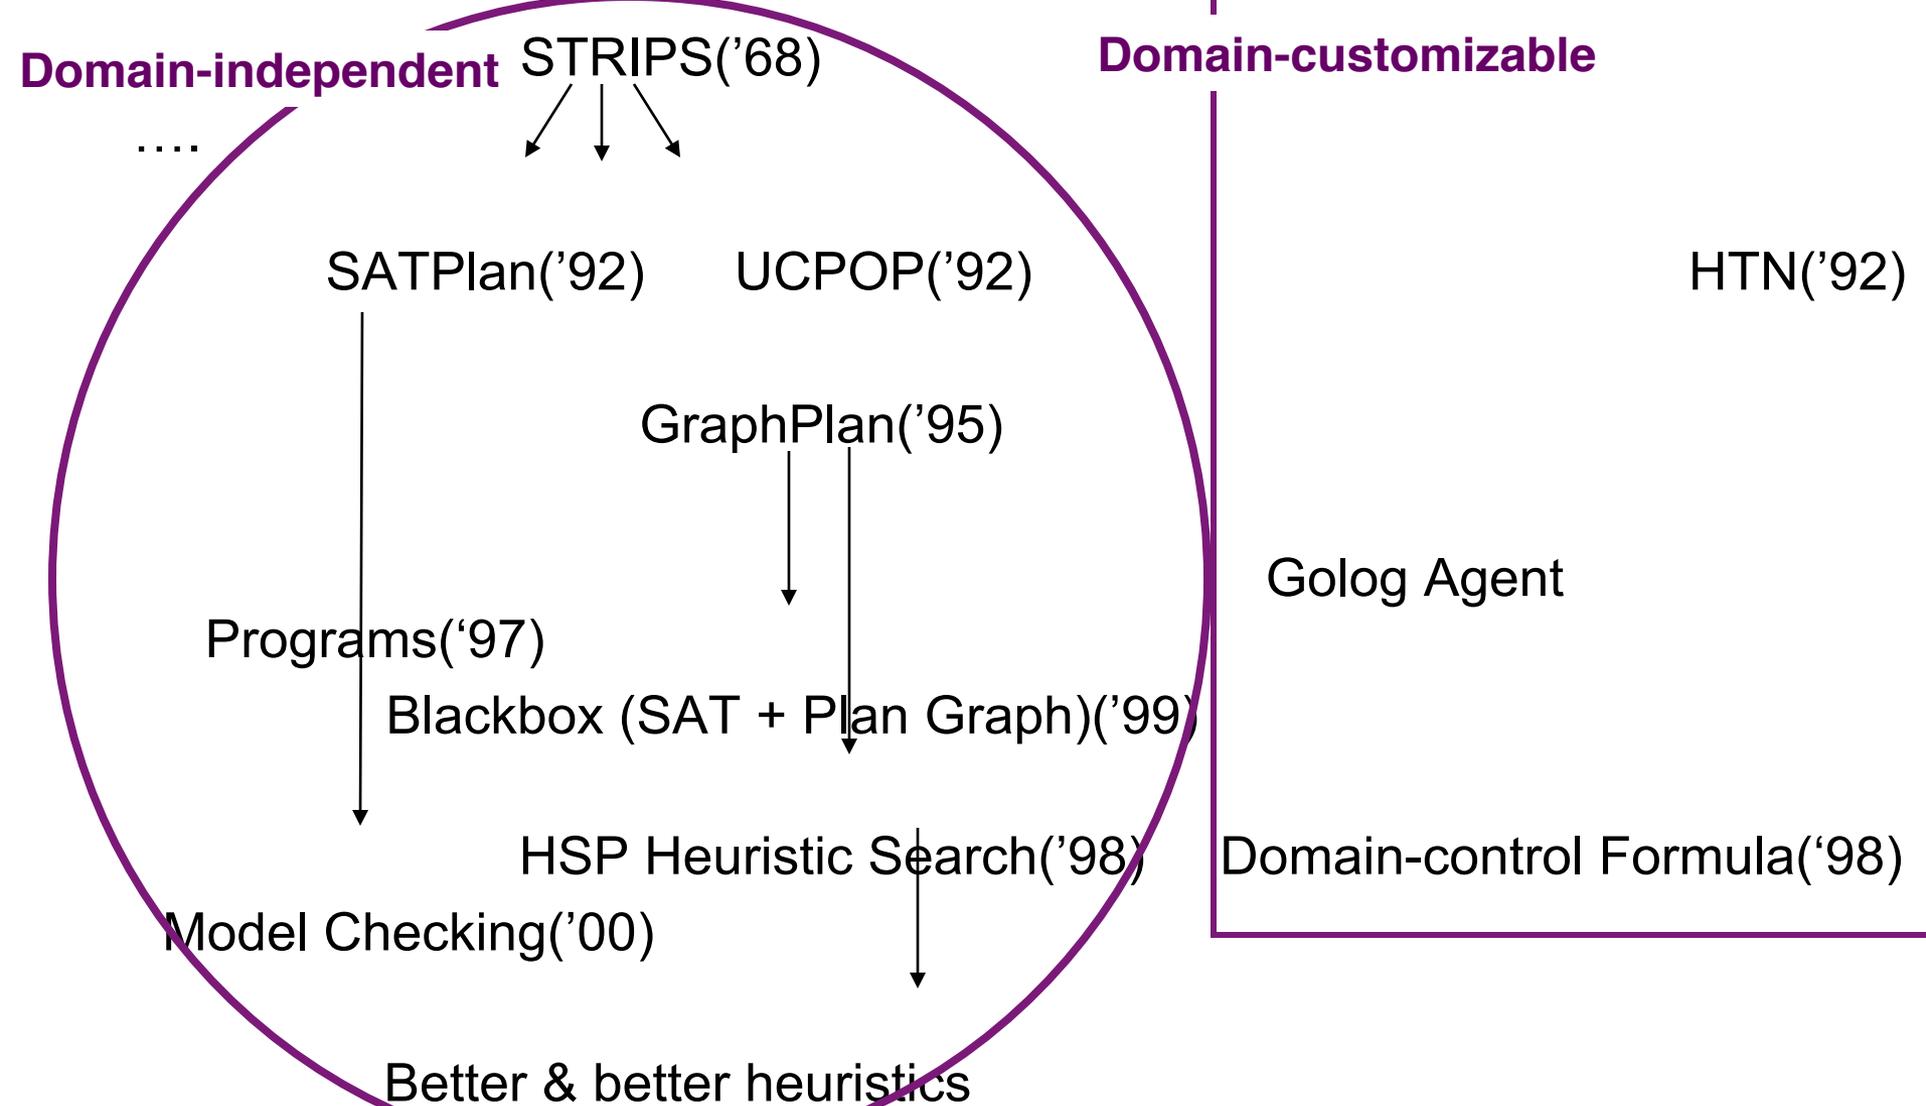
Planning Research – The big picture

At least four research communities that make fundamental contributions to research in planning:

- 1) Knowledge Representation and Reasoning Community
- 2) Automated Planning and Scheduling Community
- 3) Search Community
- 4) Machine Learning Community

Planning Algorithms – Recent Advances

Selected Historical perspective:



Planning Research – The big picture

The Landscape:

CONFERENCES

ICAPS* (Int. Conf on AI Planning and Scheduling)

*merging of AIPS and ECP

SoCS (Symposium on Combinatorial Search)

AAMAS (Int. Conf. on Autonomous Agents and Multiagent Systems)

KR

UAI (Uncertainty in AI) for MDP, POMDP, RL

ML, NIPS, ... for RL

IJCAI, AAI, ECAI

JOURNALS

JAIR, AIJ, ...

BIENNIAL COMPETITION and BENCHMARKING DOMAINS

IPC-*n* (International Planning Competition)

PDDL (Planning Domain Definition Language)

standard input language for most benchmark problem sets

Planning Research – The big picture

Recent Advances

Very “active” field -- lots of papers in top conferences

- Tremendous strides in deterministic plan synthesis
 - Biennial Intl. Planning Competitions
- Current interest is in exploiting the insights from deterministic planning techniques to other planning scenarios

Some topics of recent focus:

- Better heuristics
- Better search, real-time search, sampling,
- Richer domain customization (including preferences)
- From discrete to timed hybrid and/or continuous systems
- Planning and learning

Outline

- Conceptual model for planning
- Classes of planning problems
- Classes of planners and example instances
- Beyond planning
- Planning research – the big picture
- ➔ Some of what I hope you'll get from the course