


Artificial Intelligence

V07: Planning

Planning as search
Algorithms for classical planning
Next steps

Based on material by

- Stuart Russell, UC Berkeley
- Peter Ljunglöf, U Gothenburg / Chalmers
- Malte Helmert, U Basel



Educational objectives

- **Remember** PDDL semantics
- **Explain** in which regard and why **planning is the largest part in AI**
- **Comprehend** and **extend plans** given in PDDL
- **Know** the **road ahead** for more **complicated planning problems**

“In which we see how an agent can take advantage of the structure of a problem to construct complex plans of action.”

➔ Reading: AIMA, ch. 10 [+ ch. 11]

(ch. 10-11.2 covered here)





1. PLANNING AS SEARCH

Planning and AI

Classical planning

- «*Planning is the **art and practice of thinking before acting***» Patrik Haslum
«*Devising a plan of action to achieve one's goal*» AIMA p. 366
- Planning agents **seen so far**:
 - **Problem solving** agent (V03/V04): atomic representation → needs domain-specific heuristics
 - Hybrid propositional **logic** agent (V06a): **ground** (i.e., variable-free) sentences → may get swamped
- ➔ The part of AI being **conducted by most** researchers today calling themselves «**AI guys**»

Why is planning so big?

- Solved **applications**: Large logistics problems, operational planning, robotics, scheduling, ...
- Community: **Search** is its basis; **logic & knowledge** representation is **part of it** → treated at specialized (ICAPS) and major AI (IJCAI, AAI, ECCAI) international conferences
- AI's tendency of **spawning new disciplines**:
 - Many now autonomous disciplines started as a field of study within AI
 - Examples: Computer vision, robotics, information retrieval, automatic speech recognition
 - Currently machine learning seems to take this path
- Other universalist tendencies: “**Everything is search**”, “**everything is optimization**”



One of planning's big shots: Malte Helmert of University of Basel (→ see http://ai.cs.unibas.ch/misc/tutorial_aaai2015/)

Automated planning

Setting

- a **single agent** in a (→ multi-agent / game-playing possible)
- **fully observable**, (→ conformant planning possible)
- sequential and discrete, (→ temporal and real-time planning possible)
- **deterministic** and (→ probabilistic planning possible)
- **static** (offline) environment (→ online possible)

Tool: **Planning Domain Definition Language (PDDL)**

- A **subset of FOL**, more expressive than propositional logic
- Used to **define the planning task as a search** problem:
 - Initial states and goal states
 - A set of *Action(s)* in terms of **preconditions** and **effects** → *Result(s, a)*
 - **Closed world assumption**: Unmentioned state variables are assumed false
- It allows for **factored** representation (collection of variables)
- Derived from the **STRIPS** planning language

“Tower of Hanoi” is a classic planning example



PDDL / STRIPS operators

Tidily arranged action descriptions, restricted language

From action **schema** to ground action

- Action schema (variables are **universally quantified** $[\forall]$):

Action(*Fly*(*p*, *from*, *to*),

Precondition: $At(p, from) \wedge Plane(p) \wedge Airport(from) \wedge Airport(to)$

Effect: $\neg At(p, from) \wedge At(p, to)$

- Ground action (all variables have been substituted with values):

Action(*Buy*(*MarshallGuitarBox*, *StringsMusicStore*),

Precondition: $At(StringsMusicStore) \wedge Sells(StringsMusicStore, MarshallGuitarBox)$

Effect: *Have*(*MarshallGuitarBox*)

→Note: this abstracts away many important details of buying!

Note that capitalization of atoms (predicates & terms) is different here as compared to Datalog (V06b), to be consistent with AIMA.

Upper-case constants



Restricted language allows for **efficient** algorithms

- Action precondition: conjunction of positive literals
- Action effect: conjunction of literals
- Applicability of action *a* in state *s*: *iff* $s \models Precondition(a)$
→ E.g., $\forall p, from, to (Fly(p, from, to) \in Actions(s)) \Leftrightarrow s \models (At(p, from) \wedge Plane(p) \wedge Airport(from) \wedge Airport(to))$
- Computing the result: $Result(s, a) = (s - Del(a)) \cup Add(a)$ without explicit reference to time!
(**delete list** contains all negative literals in *Effects*(*a*), **add list** all positives)

Example: air cargo transport

- A classical transportation problem: Loading / unloading cargo, flying between different airports
- Actions: $Load(cargo, plane, airport)$, $Unload(cargo, plane, airport)$, $Fly(plane, airport, airport)$
- Predicates: $In(cargo, plane)$, $At(cargo \vee plane, airport)$
- **Complete PDDL planning problem description** (with all variables **existentially quantified** \exists):

Initial & goal state are given; $Action(s)$ and $Result(s, a)$ follow from action schemas.


```

Init( $At(C_1, SFO) \wedge At(C_2, JFK) \wedge At(P_1, SFO) \wedge At(P_2, JFK)$ 
 $\wedge Cargo(C_1) \wedge Cargo(C_2) \wedge Plane(P_1) \wedge Plane(P_2)$ 
 $\wedge Airport(JFK) \wedge Airport(SFO)$ )
Goal( $At(C_1, JFK) \wedge At(C_2, SFO)$ )
Action( $Load(c, p, a)$ ,
  PRECOND:  $At(c, a) \wedge At(p, a) \wedge Cargo(c) \wedge Plane(p) \wedge Airport(a)$ 
  EFFECT:  $\neg At(c, a) \wedge In(c, p)$ )
Action( $Unload(c, p, a)$ ,
  PRECOND:  $In(c, p) \wedge At(p, a) \wedge Cargo(c) \wedge Plane(p) \wedge Airport(a)$ 
  EFFECT:  $At(c, a) \wedge \neg In(c, p)$ )
Action( $Fly(p, from, to)$ ,
  PRECOND:  $At(p, from) \wedge Plane(p) \wedge Airport(from) \wedge Airport(to)$ 
  EFFECT:  $\neg At(p, from) \wedge At(p, to)$ )

```

- **Plan:**

[$Load(C_1, P_1, SFO), Fly(P_1, SFO, JFK), Unload(C_1, P_1, JFK),$
 $Load(C_2, P_2, JFK), Fly(P_2, JFK, SFO), Unload(C_2, P_2, SFO).$]



Example: blocks world


The blocks world

- A block is either on the table or on another block
- Blocks can be stacked (only if one fits directly on another)
- Goal: produce a given configuration of blocks on the table (specified as which is on top of what)
- Challenge: **No explicit quantifiers in PDDL** → need to introduce artificial predicates
 - Example: *Precondition*: $\neg\exists x On(x, B)$ not directly expressible → introduce predicate *Clear*(B)

Example

```

Init(On(A, Table) ∧ On(B, Table) ∧ On(C, A)
    ∧ Block(A) ∧ Block(B) ∧ Block(C) ∧ Clear(B) ∧ Clear(C))
Goal(On(A, B) ∧ On(B, C))
Action(Move(b, x, y),
    PRECOND: On(b, x) ∧ Clear(b) ∧ Clear(y) ∧ Block(b) ∧ Block(y) ∧
              (b ≠ x) ∧ (b ≠ y) ∧ (x ≠ y),
    EFFECT: On(b, y) ∧ Clear(x) ∧ ¬On(b, x) ∧ ¬Clear(y))
Action(MoveToTable(b, x),
    PRECOND: On(b, x) ∧ Clear(b) ∧ Block(b) ∧ (b ≠ x),
    EFFECT: On(b, Table) ∧ Clear(x) ∧ ¬On(b, x))
  
```



→ A possible solution sequence: $[MoveToTable(C, A), Move(B, Table, C), Move(A, Table, B)]$

How difficult is planning?

Computational complexity of classical planning

Problem definition (see V06a appendix for SAT)

- The **PlanSAT** problem: Does there exist a plan that achieves the goal?
- The **bounded PlanSAT** problem: Does there exist a solution of length at most k ?
→ useful for optimal (i.e., shortest plan) planning

The PSPACE class contains problems **solvable** by a deterministic algorithm **with its memory constrained to be polynomial in the input** length
→ **larger & more difficult than NP** (but no constraint on time)

Complexity

- PlanSAT and bounded PlanSAT are **PSPACE-complete**
→ i.e., **difficult** (assumed to be not even in NP)!
- PlanSAT **without negative** preconditions and without negative effects is in P
→ i.e., **solvable**

Practice

- **Sub-optimal planning** is sometimes **easy**
- **PDDL has facilitated** the development of **very accurate** domain-independent **heuristics** making planning feasible (formalisms based on FOL have had less success)




2. ALGORITHMS FOR CLASSICAL PLANNING

Planning as state-space search

...approachable with any algorithm from V03 or local search

Two formulations

- **Forward** (progression): search considers actions that are **applicable**
- **Backward** (regression): search considers actions that are **relevant**
- **Neither** of them is **efficient without good heuristics!**



Futility of uninformed forward search


- Example 1: Buying a copy of AIMA
 - Tool: Action schema $Buy(isbn)$ with effect $Own(isbn)$
 - 10-digit ISBN leads to $10^{10} = 10 \text{ billion ground actions}$ to be enumerated
- Example 2: Moving all cargo from airport A to airport B
 - Setting: 10 airports with 5 planes and 20 pieces of cargo at each
 - Obvious solution: load all cargo at A in one of the planes, fly to B , unload everything (41 actions)
 - search graph has 2000^{41} nodes up to this depth (assuming $\sim 2'000$ actions per state on average)

Heuristics for forward state-space search

Enabled by factored representations for states & actions

Possible domain-independent heuristics

- **Relaxing actions** (i.e., adding new links to the graph to ease the problem)
 - **Ignore-preconditions heuristic**: All actions are applicable anytime
→ leads e.g. easily to the 2 different heuristics for the n -puzzle of V03
 - **Ignore-delete-lists heuristic**: Removing all negative literals from effects
→ enables making monotonic progress towards goal, achievable e.g. with hill climbing
- **State abstractions** (i.e., collapsing multiple states into a single one to shrink the graph)
 - Reduce the state space by e.g. **ignoring some** fluents



Winners of the bi-annual ICAPS planning competition often used

- Heuristic search (→ see FastDownward system: Helmert et al. 2004, <http://www.fast-downward.org/>)
- Planning graphs (→ see next slides)
- SAT solvers (→ see V06a and below)

Planning graphs

An alternative to basic state-space search

Challenges so far

- **Exponential** size of the search trees
- **Not all heuristics** are **admissible** in general



ELSEVIER

Artificial Intelligence 90 (1997) 281–300

Artificial
Intelligence

Fast planning through planning graph analysis*

Avrim L. Blum*, Merrick L. Furst¹School of Computer Science, Carnegie Mellon University, 5000 Forbes Avenue,
Pittsburgh, PA 15213-3891, USA

Received December 1995; revised September 1996

Abstract

We introduce a new approach to planning in STRIPS-like domains based on constructing and analyzing a compact structure we call a planning graph. We describe a new planner, Graphplan, that uses this paradigm. Graphplan always returns a shortest possible partial-order plan, or states that no valid plan exists.

We provide empirical evidence in favor of this approach, showing that Graphplan outperforms the total-order planner, Prodigy, and the partial order planner, UCPOP, on a variety of interesting natural and artificial planning problems. We also give empirical evidence that the plans produced by Graphplan are quite sensible. Since searches made by this approach are fundamentally different from the searches of other common planning methods, they provide a new perspective on the planning problem. © 1997 Elsevier Science B.V.

Keywords: General purpose planning; STRIPS planning; Graph algorithms; Planning graph analysis

Solution: the planning graph

1. **Propositionalize** the search tree:

replace all action schemas by sets of ground actions (to remove variables etc.)

2. **Approximate** the complete propositionalized tree

- **Polynomial size:** $O(n(a + l)^2)$ for a actions, l literals and n levels
- Useful to **create admissible heuristics** like **set-level** heuristic (→ see appendix):
cost of achieving \wedge of goals = $\sum g_i$ (level cost) of goals in first level without mutual exclusivity

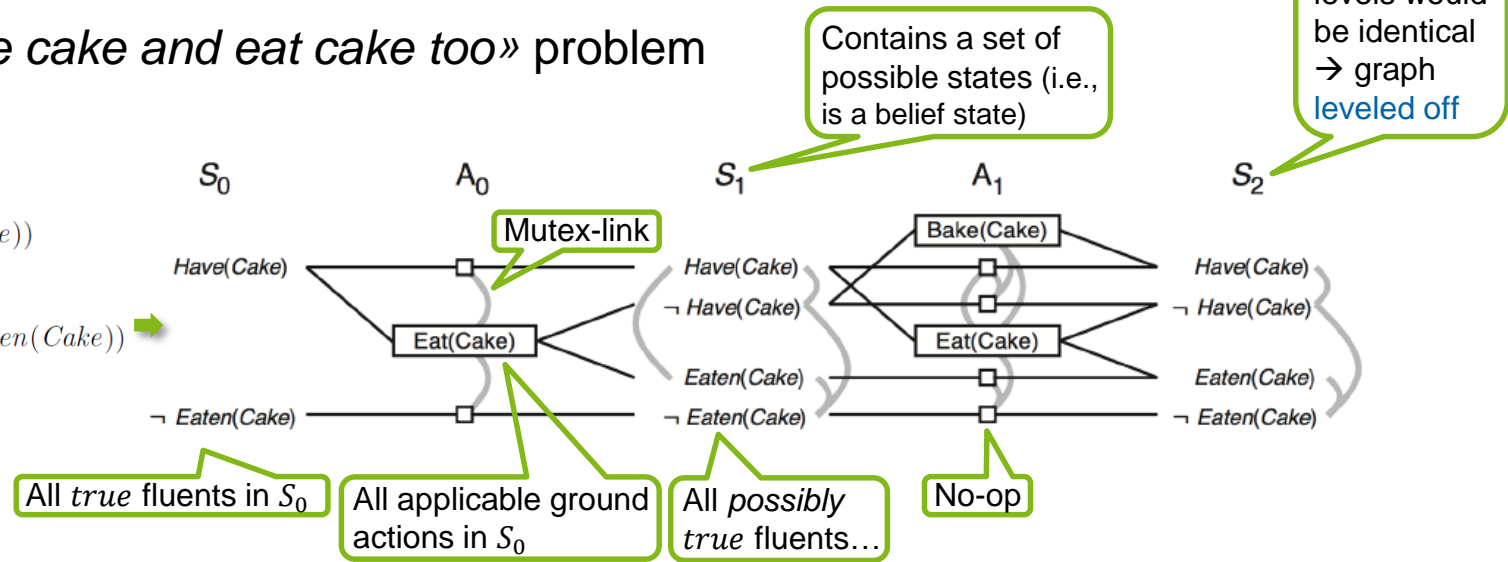
The planning graph

Organized in **alternating levels** of possible states S_i and applicable actions A_i

- S_i holds all **fluents that could be true** at that point
- A_i holds all **actions** that could have their preconditions satisfied
- Links *between* levels represent **preconditions** and **effects**
- Links *within* the levels express **conflicts** (“mutex”-links)

Example: the «*have cake and eat cake too*» problem

Init(*Have*(*Cake*))
Goal(*Have*(*Cake*) \wedge *Eaten*(*Cake*))
Action(*Eat*(*Cake*))
 PRECOND: *Have*(*Cake*)
 EFFECT: \neg *Have*(*Cake*) \wedge *Eaten*(*Cake*)
Action(*Bake*(*Cake*))
 PRECOND: \neg *Have*(*Cake*)
 EFFECT: *Have*(*Cake*)

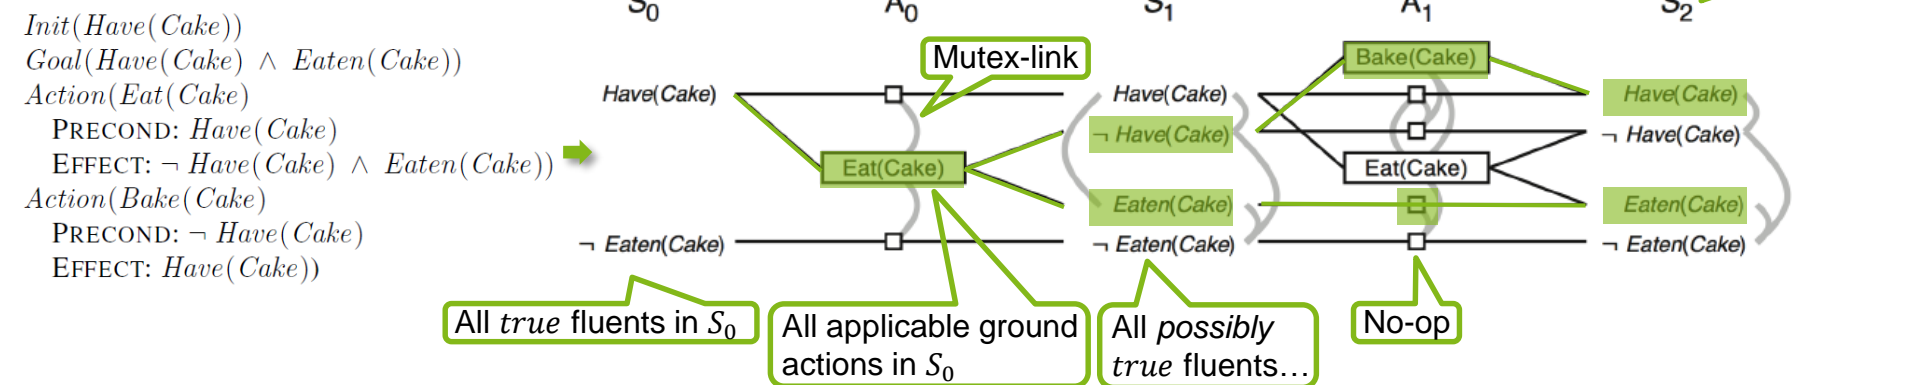


The planning graph

Organized in **alternating levels** of possible states S_i and applicable actions A_i

- S_i holds all **fluents that could be true** at that point
- A_i holds all **actions** that could have their preconditions satisfied
- Links *between* levels represent **preconditions** and **effects**
- Links *within* the levels express **conflicts** (“mutex”-links)

Example: the «*have cake and eat cake too*» problem



The GraphPlan algorithm

Plan directly using the planning graph

```

function GraphPlan(problem) returns solution or failure
  graph ← Initial-Planning-Graph(problem) #i.e., create  $S_0$ 
  goals ← Conjunctions(problem.GOAL) #AND of all goal literals
  nogoods ← an empty hash table #used for the same purpose as in constraint learning (→ see V05)
  for t = 0 to  $\infty$  do
    if goals all non-mutex in  $S_t$  of graph then
      solution ← Extract-Solution(graph, goals, Numlevels(graph), nogoods) #e.g. CSP or backward search
      if solution ≠ failure then return solution
    if graph and nogoods have both leveled off then return failure
  graph ← Expand-Graph(graph, problem)

```

Description

- GraphPlan **expands** the graph with new levels A_i & S_{i+1} **until** \nexists **mutex links between goals**
 - The `nogoods` list records `(level, goal)` pairs that couldn't be satisfied at that level
→ prevents `ExtractSolution` from searching again if called with the same arguments
 - **To extract** the actual **plan**, the algorithm **searches backwards** in the graph
 - Initial state is the last level S_n of the planning graph
 - Available actions at level S_i are conflict-free subsets of actions at A_{i-1} with effects realizing S_i 's goals
 - Goal is to reach a state at S_0 such that all goals are satisfied
- The plan **extraction** is the difficult part and is usually done with **greedy-like heuristics**

SATplan and CSP solvers

More alternatives to planning

Translate PDDL description into a SAT problem or a CSP

- Goal state and all actions have to be propositionalized
 - Action schemas have to be replaced by a set of ground actions (variables to be replaced by constants)
 - Fluents need to be introduced for each time step
 - ...
- combinatorial explosion

Cost – benefit

- Remove a part of the benefits of the expressiveness of PDDL to...
- ...gain access to efficient solution methods of SAT and CSP solvers






3. NEXT STEPS

Planning in the real world

Use cases

- **Spacecraft operation in real time (1998)**
NASA's Deep Space 1 was controlled by a planning & scheduling system devising and carrying out plans like «*During the next week take pictures of the following asteroids and thrust 90% of the time*» (→ see [Nilsson, 2010] ch. 32.2.1)
- **Factory scheduling (1985)**
4 week (3 shifts a day) production plan at Hitachi for assembly line of 350 products with 35 machines and >2000 different operations (→ see AIMA ch. 11.2.2, **HTN** on next slide)
- **Military operation planning (1990)**
A scheduling program helped with the logistics of 1st gulf war and is said to have «*paid back all of DARPA's 30 years of investment in AI in a matter of a few months*» (→ see [Nilsson, 2010] ch. 23.3.3)



Challenges

- Taking **resources** (incl. time) into account → **scheduling**
- Being **overwhelmed by** state space **size** → **hierarchical planning**
- Needing to **incorporate human** wisdom → hierarchical planning
- Coping with **uncertainty** → **conformant** / **contingency** / **online planning** (analog to AIMA ch. 4)
- Planning with **multiple agents** → planning with **cooperative and adversarial** multiagents is **unsolved**



Outlook: Hierarchical planning

The need for **abstraction**

- Atomic actions for humans: ca. 10^3 muscles, max. 10 mindful activations per second
- Plan for a lifetime: ca. 10^9 wake seconds → ca. 10^{13} possible actions per life
- Plan for 2 weeks vacation: ca. $(60 * 60 * 24 * 14) \times 10^3 \times 10 \approx 10^{10}$ actions
- Methods seen so far work only for thousands (i.e., $\ll 10^{10}$) of actions
→ hierarchical decomposition
(e.g., “go to ZRH” → “get to train station, take train to Zurich airport, ascend to departure hall” → ...)

Technical solution sketch

- **Hierarchical task networks** (HTN): more **factored** representations for **actions** (besides states)
- Two kinds of actions:
 - **Primitive actions**: standard precondition-effect schemas
 - **High level actions** (HLA): e.g. “go to ZRH” → have one or more possible **refinements**
 - Refinement: a **sequence of HLAs or primitive actions**, maybe recursive
- Key benefits: Possibly huge **speed** improvements, possibility for **humans to define HLAs**
- Implementation: E.g. by **forward breadth-first search** (but can be done much better)

Where's the intelligence?

Man vs. machine

- Planning is foremost an **exercise in controlling combinatorial** explosion
- It does so by **combining** efficient **search & logical** reasoning
 - necessary speedups are achieved by **domain-independent heuristics** that exploit structure in the representation
 - this is **really smart**
- But: There is **no clear understanding** yet of **which** methods **work best on what** problems
- In contrast to popular opinion, AI planning is **widely applied in practice** today
 - Also, research is not “dead”, but **less hyped** at the moment
 - Probably planning is the **best** that **symbolic AI** currently offers



Automating university timetabling by planning?

A search exercise

Automatic **scheduling** is a relevant subfield of AI planning. Likewise, automated timetable generation (often focused on university teaching timetables) is a vibrant field of study.

- Conduct a **quick literature research** on automated timetabling (e.g. <https://scholar.google.ch/scholar?q=automated+timetabling>)
- What **kind of approaches** are proposed? How do they **relate to AI planning** as you have heard of here?
- With your current understanding of AI – **how would you approach the problem?** What are your **options?**

Zeit	Montag 20.02.	Dienstag 21.02.	Mittwoch 22.02.	Donnerstag 23.02.	Freitag 24.02.
08:00 - 08:45	T_JT14a... Stdm TH 561 T_JT13l... Stdm TH 561 T_JT13l... Stdm TH 561 T_JT14a... Stdm TH 561				
08:50 - 09:35	T_VSRep.BA Stdm TH 561 T_JT14l... Stdm TH 561 T_JT14l... Stdm TH 561				
10:00 - 10:45	T_JT14l... Tugg TH 547 T_JT14l... Tugg TH 547 T_VSRep.BA Tugg TH 547 T_JT13l... Tugg TH 547				
10:50 - 11:35	T_JT13l... Tugg TH 547 T_JT14a... Tugg TH 547 T_JT14a... Tugg TH 547				
12:00 - 12:45					
12:50 - 13:35					T_JT14l... Stdm ZL 05.05 T_JTRep... Stdm ZL 05.05 T_JT14a... Stdm ZL 05.05 T_JT13l... Stdm ZL 05.05
14:00 - 14:45		T_JT14l... Tugg TH 544 T_JT14l... Tugg TH 544 T_VSRep.BA Tugg TH 544 T_JT13l... Tugg TH 544 T_JT13l... Tugg TH 544			T_JT13l... Stdm ZL 05.05 T_JT13l... Stdm ZL 05.05 T_JT14a... Stdm ZL 05.05
14:50 - 15:35		T_JT14a... Tugg TH 544 T_JT14a... Tugg TH 544			T_JT14a... Tugg ZL 05.16 T_JT13l... Tugg ZL 05.16 T_JT13l... Tugg ZL 05.16
16:00 - 16:45					T_JTRep... Tugg ZL 05.16 T_JT14l... Tugg ZL 05.16 T_JT14a... Tugg ZL 05.16



Review

- **Planning is AI's main field**, due to success stories like remotely controlling a NASA spacecraft in real time
 - Planning refers to problem solving techniques (**i.e., search**) on factored (**i.e., logic-based**) representations of states and actions, allowing for **fast algorithms**
 - **PDDL** describes the initial and goal **states as conjunctions** of literals; **actions** in terms of their **preconditions and effects**
- Effective domain-independent **heuristics** are derived **by subgoal independence** or **problem relaxation**
- A **planning graph** is constructed **incrementally**
 - Each **layer containing** the **superset of** all actions/literals that could occur in this time step, including **mutex relationships**
 - Can be used to **derive** useful **heuristics**; or directly for **planning via GraphPlan**
- Other approaches are using **SAT or CSP solvers**
 - **FOL-based** planning has much-needed expressiveness for larger real-world problems, but yet **no efficient algorithms** (missing heuristics)
 - Workarounds include **hierarchical planning** through HTNs
- It is yet **unknown which approach is best**





APPENDIX

Using planning graphs to devise heuristics


I.e., only one action per level / time step

A (**serial**) planning graph facilitates domain-independent heuristics

- Problem is unsolvable if any goal literal fails to appear in the final level
- **Level cost** of goal g_i : Level in planning graph at which g_i first appeared
- Heuristics for conjunctions of goals:
 - **Max-level heuristic**: **maximum** of the **level costs** of all subgoals
 - **Level sum heuristic**: **sum** of **level costs** of all subgoals (assuming independence → not necessarily admissible)
 - **Set-level heuristic**: **First level** at which all subgoals appear **without any pair being mutex**
- As with CSPs: **checking** for **pair-wise** consistency often **pays off**; higher order often doesn't

Heuristic values for
 $Have(Cake) \wedge Eaten(Cake)$:

- Max-level: $\max(1,0) = 1$
- Level sum: $1 + 0 = 1$
- Set-level: 2 (**accurate!**)




Historical remark: Linear planning

Planners in the early 1970s considered **totally ordered** action **sequences**

- Problems were decomposed in subgoals
- Resulting subplans were strung together in some order
 - This is called **linear planning**

But, linear planning **is incomplete!**

- There are some very simple problems it cannot handle
- E.g., the **Sussman anomaly**: Unsolvable by linear planner
 - A complete planner must be able to interleave subplans



Enter **partial-order planning**, state-of-the-art during the 1980s and 90s

- Today mostly **used for specific tasks**, such as operations scheduling
- Also used **when it is important for humans to understand** the plans
 - E.g., operational plans for spacecraft and Mars rovers are checked by human operators before uploaded to the vehicles