## Principles of AI Planning

February 7th, 2007 - Nondeterministic planning with partial observability Introduction
Reduction to fully observable case
Idea
Basic translation
Caveat
Observations
Discussion
Forward search
Idea
Algorithm
Backward search
Idea
Observations
Algorithm

## Example

Summary
M. Helmert, B. Nebel (Universität Freiburg)

## Nondeterministic planning with partial observability

Planning with partial observability is harder than both the fully observable and unobservable cases:

- Memoryless plans (where the next action to take only depends on the current situation) as in the fully observable case are not sufficient.
- Of course, we cannot define a memoryless plan based on individual states because limited observability makes some states indistinguishable.
- It is also not sufficient to consider memoryless plans where the action to take is based on the current observation class.
- Conformant (i.e., non-branching) plans as in the unobservable case are also clearly not powerful enough.


## Principles of AI Planning

Nondeterministic planning with partial observability

## Malte Helmert Bernhard Nebel

Albert-Ludwigs-Universität Freiburg
February 7th, 2007

[^0]Al Planning

## Strong planning

- We will (mostly) consider the strong planning problem.
- Generalizations to the strong cyclic planning are similar to the fully observable case.


## Algorithms

Similar to other variants of the planning problem, there are three major approaches to nondeterministic planning with partial observability:

- Reduction to another problem
- Forward search
- Backward search

We will consider one example for each of these.

Reduction to fully observable case

- Memoryless plans are not sufficient for the partially observable case because a plan must take into account the knowledge collected in previous observations etc.
- During plan execution, this knowledge is represented in the current belief state.
- One idea for solving a partially observable task $\mathcal{T}$ is to map it to a fully observable task $\mathcal{T}^{\prime}$ where each belief state of $\mathcal{T}$ corresponds to a state of $\mathcal{T}^{\prime}$.


## Algorithms

Three approaches
Reduction to another problem:

- Reduce to planning with full observability.

Forward search (progression):

- Define the search space as an AND/OR tree.
- Define a heuristic function for such trees.
- Use a tree search algorithm such as AO* or Proof Number Search.

Backward search (regression):

- Start from the set of goal states.
- Find state sets from which already generated state sets can be reached by applying operators and making observations.
M. Helmert, B. Nebel (Universität Freiburg)

Al Planning

Reduction to fully observable case
State variables

Let $\mathcal{T}=\langle A, I, O, G, V\rangle$ be the input task with state set $S$. We define the fully observable task $\mathcal{T}^{\prime}=\left\langle A^{\prime}, I^{\prime}, O^{\prime}, G^{\prime}, A^{\prime}\right\rangle$.

State variables

- For each state $s \in S$, there is one state variable $v_{s} \in A^{\prime}$.
- Intuition: $v_{s}$ is true in a state of $\mathcal{T}^{\prime}$ iff it is possible that we are currently in $s$.
- Formally: $A^{\prime}:=\left\{v_{s} \mid s \in S\right\}$


## Reduction to fully observable case

Initial state formula

Let $\mathcal{T}=\langle A, I, O, G, V\rangle$ be the input task with state set $S$. We define the fully observable task $\mathcal{T}^{\prime}=\left\langle A^{\prime}, I^{\prime}, O^{\prime}, G^{\prime}, A^{\prime}\right\rangle$.
Initial state formula

- The initial state of $\mathcal{T}^{\prime}$ is fully deterministic (in terms of $A^{\prime}$ ), as there is only one possible initial belief state in $\mathcal{T}$
- For all states $s$ in the initial belief state of $\mathcal{T}$, variable $v_{s}$ is initially true. Other variables are initially false.
- Formally: $I^{\prime}:=\bigwedge_{s \in S, s \models I} v_{s} \wedge \bigwedge_{s \in S, s \mid \neq I} \neg v_{s}$.

Reduction to fully observable case
Initial state formula

Let $\mathcal{T}=\langle A, I, O, G, V\rangle$ be the input task with state set $S$. We define the fully observable task $\mathcal{T}^{\prime}=\left\langle A^{\prime}, I^{\prime}, O^{\prime}, G^{\prime}, A^{\prime}\right\rangle$.
Operators (preconditions)

- Each operator $o=\langle c, e\rangle \in O$ is translated to an operator $o^{\prime}=\left\langle c^{\prime}, e^{\prime}\right\rangle \in O^{\prime}$.
- To test whether operator $o$ is applicable, we must verify that all states in the current belief state satisfy $c$.
- Again, this is equivalent to saying that no state in the current belief state violates $c$.
- Formally: $c^{\prime}:=\bigwedge_{s \in S, s \neq c} \neg v_{s}$.


## Reduction to fully observable case

Initial state formula

Let $\mathcal{T}=\langle A, I, O, G, V\rangle$ be the input task with state set $S$.
We define the fully observable task $\mathcal{T}^{\prime}=\left\langle A^{\prime}, I^{\prime}, O^{\prime}, G^{\prime}, A^{\prime}\right\rangle$.
Goal formula

- A goal belief state of $\mathcal{T}$ is one where all possible states satisfy $G$.
- This is equivalent to saying that no state in the current belief state violates $G$
- We can express that by saying that none of the variables $v_{s}$ for states $s$ violating $G$ are true.
- Formally: $G^{\prime}:=\bigwedge_{s \in S, s \nmid G G} \neg v_{s}$.


## Reduction Basic translation

## Reduction to fully observable case

Initial state formula
Let $\mathcal{T}=\langle A, I, O, G, V\rangle$ be the input task with state set $S$.
We define the fully observable task $\mathcal{T}^{\prime}=\left\langle A^{\prime}, I^{\prime}, O^{\prime}, G^{\prime}, A^{\prime}\right\rangle$.
Operators (effects)

- Each operator $o=\langle c, e\rangle \in O$ is translated to an operator $o^{\prime}=\left\langle c^{\prime}, e^{\prime}\right\rangle \in O^{\prime}$.
- After applying operator $o$, we can possibly be in state $s \in S$ iff we were previously in some state in which $o$ is applicable and from which applying $o$ can lead to $s$.
- This is modeled by an effect
$\left(\left(\bigvee_{t \in \text { preimg }_{o}(s)} v_{t}\right) \triangleright v_{s}\right) \wedge\left(\neg\left(\bigvee_{t \in \text { preimgo }_{o}(s)} v_{t}\right) \triangleright \neg v_{s}\right)$.
- Formally: $e^{\prime}:=\bigwedge_{s \in S}\left(\left(\left(\bigvee_{t \in \text { preimgo }_{o}(s)} v_{t}\right) \triangleright v_{s}\right) \wedge\right.$

$$
\left.\left(\neg\left(\bigvee_{t \in \text { preimgo }^{(s)}} v_{t}\right) \triangleright \neg v_{s}\right)\right)
$$

## Reduction to fully observable case

Done?

- We have translated state variables, initial state formula, goal formula and operators.

Is that it?

- So far, our translation is independent of the set of observable variables $V$ !
- Moreover, the resulting planning task is deterministic!

Is there an error in our modeling?
M. Helmert, B. Nebel (Universität Freiburg)

## Reduction to fully observable case

Observations
Let $\mathcal{T}=\langle A, I, O, G, V\rangle$ be the input task with state set $S$.
We define the fully observable task $\mathcal{T}^{\prime}=\left\langle A^{\prime}, I^{\prime}, O^{\prime}, G^{\prime}, A^{\prime}\right\rangle$.
Observations

- In general, our formalism allows observations to be general formulas over $V$. However, it is sufficient to only consider atomic observations $u \in V$.
- If we observe $u$ in a belief state $b$, we can end up in two different belief states: one containing exactly the states of $b$ where $u$ is true, and one containing exactly the states of $b$ where $u$ is false.
- In other words, either the belief states where $u$ is false or the belief states where $u$ is true are ruled out
- Formally: Translate observation of $u \in V$ into an operator $\left\langle T, e_{u}^{\prime}\right\rangle \in O^{\prime}$ with $e_{u}^{\prime}:=\left(\bigwedge_{s \in S, s \neq u} \neg v_{s}\right) \mid\left(\bigwedge_{s \in S, s \models u} \neg v_{s}\right)$.


## Reduction to fully observable case

Not done

Is there an error in our modeling?

- No, but it is not complete yet: There are solvable partially observable tasks $\mathcal{T}$ for which $\mathcal{T}^{\prime}$ (as defined so far) is unsolvable.
- The reason for this is that he have not yet modeled the possibility of observing state variables.
Modeling observations requires introducing nondeterminism in $\mathcal{T}^{\prime}$.

[^1]AI Planning

## Reduction to fully observable case

Discussion

- Note that the reduction works both for strong and for strong cyclic planning.
- The reduction has a significant drawback: Since it introduces as many state variables as there are states in the original task, the resulting problem is exponentially larger than the original one.
- This will usually not be practical.
- On the other hand, there does not really exist any truly "practical" algorithm for nondeterministic planning with partial observability.


## Reduction to fully observable case

Complexity result

- Using an exponential-time planning algorithm for fully observable planning, $\mathcal{T}^{\prime}$ can be solved in time $O\left(c^{\left\|\mathcal{T}^{\prime}\right\|}\right)$, and $\left\|\mathcal{T}^{\prime}\right\|=O\left(c^{\|\mathcal{T}\|}\right)$.
- Thus, we have a double-exponential $\left(O\left(c^{d\|\mathcal{T}\|}\right)\right)$ algorithm for nondeterministic planning for partial observability.
- We will later prove that this is worst-case optimal.

Search in AND/OR trees

In forward search, plans are represented as trees whose nodes represent the situations arising during plan execution.

- The root node represents the initial situation.
- OR nodes correspond to choosing and applying operators.
- Note how these relate to operators in $\mathcal{T}^{\prime}$ in the earlier reduction.
- AND nodes correspond to making observations.
- Note how these relate to nondeterminism in $\mathcal{T}^{\prime}$ in the earlier reduction.


## Search in AND/OR trees

Example


Forward search Algorithm

## AND/OR trees

Formal definition

## Definition

An AND/OR tree is a labeled rooted tree where

- internal nodes are labeled with $(\wedge)$ or $(\vee)$
(AND nodes/OR nodes), and
- leaves are labeled with $(\top)$ or $(\perp)$ (true leaves/false leaves).


## AND/OR trees

Truth value

## Definition

An AND/OR tree evaluates to true iff

- it is a true leaf,
- it is an OR node with a child that evaluates to true, or
- it is an AND node whose children all evaluate to true.
M. Helmert, B. Nebel (Universität Freiburg)


## Partial plan trees

Definition (ctd.)
A partial plan tree for a nondeterministic planning task $\langle A, I, O, G, V\rangle$ with state set $S$ is an AND/OR tree with the following properties:

- ...
- An OR node $n$ (also called an operator node) has one child $n_{0}$ for each operator $o \in O$ applicable in $b(n)$, with associated belief state $b\left(n_{o}\right)=a p p_{o}(b(n))$.
- An AND node $n$ (also called an observation node) has an associated formula $\varphi(n)$ over $V$. It has two children:
- $n^{\top}$ with $b\left(n^{\top}\right)=\{s \in b(n) \mid s \vDash \varphi\}$
- $n^{\perp}$ with $b\left(n^{\perp}\right)=\{s \in b(n) \mid s \notin \varphi\}$.


## Partial plan trees

## Definition

A partial plan tree for a nondeterministic planning task $\langle A, I, O, G, V\rangle$ with state set $S$ is an AND/OR tree with the following properties:

- Each node $n$ has an associated belief state $b(n)$.
- If $n$ is the root node, then $b(n)=\{s \in S|s|=I\}$.
- A leaf node $n$ is labeled with $(T)$ iff $b(n) \models G$. In this case it is called a goal node, otherwise an open node.
- ...

[^2]Forward planning as search in partial plan trees

- Clearly, a partial plan tree represents a strategy.
- This strategy is a strong plan iff the tree evaluates to true.

We thus obtain a (nondeterministic) forward search algorithm:
Forward search in partial plan trees
def expand-tree $(\mathcal{T})$ :
Set $T$ to the partial plan tree for $\mathcal{T}$ that consists
of a single leaf, labeled with the initial belief state.
while $T$ evaluates to false:
Choose some open leaf $n$ in $T$.
Replace $n$ by an operator or observation node, adding the necessary children to $T$.

## Backward search algorithms

- There is a conflict between plan size and observing:
- With many observations, plans become very big
- With few observations, it may be impossible to find a plan.

Trying out all possible ways to branch is not feasible. No good general solutions to this problem exist.

- AND-OR search algorithms use heuristics for making branching decisions.
- But they do not really work well...
- Let $B_{1}, \ldots, B_{n}$ be belief states with $B_{i} \subseteq C_{i}$ for all $i=1, \ldots, n$ for which we have a solution plan.
- Then we can find a plan for $B=B_{1} \cup \cdots \cup B_{n}$ by first observing in which class $C_{i}$ we are and then applying the corresponding plan for $B_{i}$.
- Backward search algorithms are similar in flavour to the ones for fully observable problems.
- Backward steps with operator application:
- Compute strong preimages.
- Backward steps with observations:
- Compute union of belief states from disjoint observational classes.
- Note: Can always take subsets of solved belief states to make them disjoint.


Observations in backward search

Backward search Observations
Observations in backward search
Example: Combining two belief states

M. Helmert, B. Nebel (Universität Freiburg)

Observations in backward search
Example: Combining two belief states, option 2


Observations in backward search
Example: Combining two belief states, option 4


No observability $\Rightarrow$ no branching
No choice between subplans during execution: option 1
o1


No observability $\Rightarrow$ no branching
Only one observational class: no choice between subplans
o1

M. Helmert, B. Nebel (Universität Freiburg)

AI Planning

Backward search Observations
No observability $\Rightarrow$ no branching
No choice between subplans during execution: option 2
o1


## Backward search Observations

Full observability $\Rightarrow$ arbitrary branching
A different plan can be used for every state


## Backward search

Example

- Blocks world with three blocks
- Goal: all blocks are on the table
- Only the variables clear $(X)$ are observable.
- A block can be moved onto the table if the block is clear.
- 8 observational classes corresponding to the 8 valuations of \{clear(A), clear(B), clear(C)\} (one of the valuations does not correspond to a blocks world state).


## A systematic backward algorithm

Idea: always split belief states into all observational classes
Initially, the set of solved belief states includes the set $b_{G} \cap C_{i}$ for each observational class $C_{i}$, where $b_{G}$ is the belief state containing all states satisfying the goal.

Then iterate the following steps:

1. Pick one belief state $b_{i}$ for each observational class and compute their union $b$.
2. If $b$ includes all initial states $\rightsquigarrow$ solution.
3. Otherwise, compute the strong preimage of $b$ with respect to some operator 0 .
4. Split the resulting set of states to belief states for different observational classes and add them to the set of solved belief states.

Plan construction by backward search
Example: goal belief state


Plan construction by backward search
Example: backward step with red-block-onto-table


Plan construction by backward search
Example: backward step with green-block-onto-table


Plan construction by backward search
Example: backward step with blue-block-onto-table


Plan construction by backward search
Example: backward step with red-block-onto-table


M. Helmert, B. Nebel (Universität Freiburg)

Backward search Example
Plan construction by backward search
Example: backward step with green-block-onto-table

M. Helmert, B. Nebel (Universität Freiburg)

Al Planning

Plan construction by backward search
Example: backward step with red-block-onto-table


Plan construction by backward search
Example: backward step with blue-block-onto-table



[^0]:    M. Helmert, B. Nebel (Universität Freiburg)

[^1]:    M. Helmert, B. Nebel (Universität Freiburg)

[^2]:    M. Helmert, B. Nebel (Universität Freiburg)

