Observability

- Robot can see and hear only the immediate surroundings.
- Poker player cannot see the opponents' cards.
- Formalization: split state variables to those that can be, and to those that cannot be observed.
- More general formalization: set of atomic and compound formulas; truth-values of these formulas can be observed, but not necessarily the truth-values of their subformulas.

Problem definition

A 5-tuple $\langle P, I, O, G, B \rangle$ consisting of

- a set P of state variables,
- a propositional formula I over P,
- a set O of operators,
- a propositional formula G over P, and
- a set $B \subseteq P$ of observable state variables.

is *a problem instance in nondeterministic planning with partial observability*.

Conditional plans: definition

B = the observable state variables, O = the operators. A plan is $\langle N, b, l \rangle$ where

- N is a finite set of nodes,
- $b \in N$ is the initial node,
- $l : N \rightarrow (O \times N) \cup 2^{\mathcal{L} \times N}$ assigns each node
  - an operator and a successor node $\langle o, n \rangle \in O \times N$ or
  - a set of conditions and successor nodes $\langle \phi, n \rangle$, where $n \in N$ and ϕ is a formula over B.

Restrictions on observability

Let $\langle P, I, O, G, B \rangle$ be a problem instance in conditional planning.

If $P = B$, the problem instance is *fully observable*.

If $B = \emptyset$, the problem instance is *unobservable*.

If $B \neq \emptyset$ and $P \neq B$, the problem instance is *partially observable*. We also use this term for the general case that includes full observability and unobservability.

Observational classes

- When only variables in $B = \{p_1, \ldots, p_m\}$ are observable, then it is not possible to distinguish between states s and s' such that $s(p) = s'(p)$ for all $p \in B$.
- Observability partitions states to classes $S_1 \cup S_2 \cup \cdots \cup S_n$ of indistinguishable states. One class for every valuation of the observable state variables.

  $S = S_1 \cup S_2 \cup \cdots \cup S_n$,
  $S_i \cap S_j = \emptyset$ for any $\{i, j\} \subset \{1, \ldots, n\}$ such that $i \neq j$.

Observational classes: cardinality

Full observability: state space $S = \{s_1, \ldots, s_n\}$ is partitioned to singleton classes $S_1 = \{s_1\}, S_2 = \{s_2\}, \ldots, S_n = \{s_n\}$.

Unobservability: The partition has only one class $S_1 = S$ consisting of all the states.

Partial observability: the number and cardinality of S_i may be anywhere between 1 and n.

Belief states and the belief space

- Current state is not in general known during plan execution. Instead, the a *set of possible current states* is known.
- A set of possible current states is a *belief state*.
- The set of all belief states is the *belief space*.
- If there are n observationally indistinguishable states, then there are 2^n belief states.

Transition graph for the belief space

1. Let F be a belief state (a set of states).
2. Operator o is executable in F if it is executable in every $s \in F$.
3. When o is executed, possible next states are $E = \mathit{img}_o(F)$.
4. Observations select some observational class S_j.
5. New belief state is $F' = \mathit{img}_o(F) \cap S_j$.

   $\Longrightarrow$ There is an arc from F to F' in the transition graph for the belief space.

Belief states and the belief space

EXAMPLE (next slide):

Belief space generated by states over two Boolean state variables.

$n = 2$ state variables, $2^n = 4$ states, $2^{2^n} = 16$ belief states

red action: complement the value of the first state variable

blue action: assign a random value to the second state variable

Belief space: example, cont'd

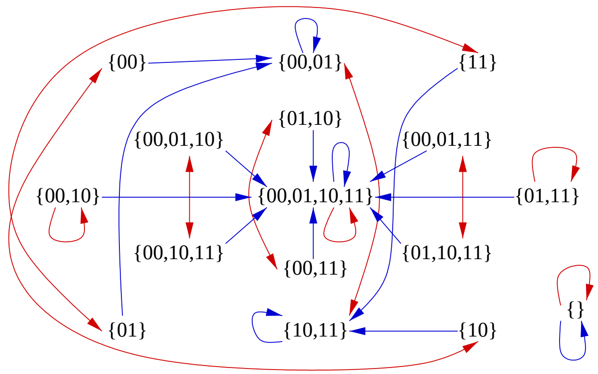

Belief space: example II

A robot without any sensors, anywhere in the classroom.

Actions: go North, South, East, West

Plan for getting out: 6 $\times$ West, 7 $\times$ North, 1 $\times$ East, 1 $\times$ North

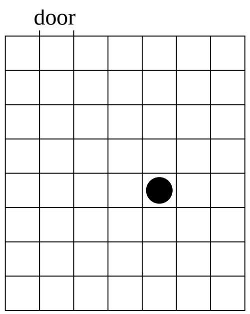

Belief space: example II, belief state initially

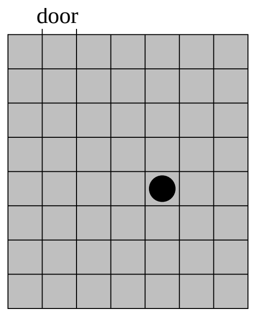

Belief space: belief state after W

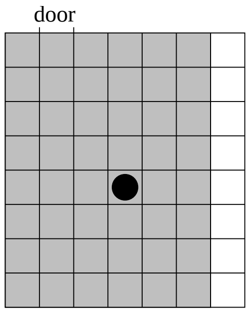

Belief space: after WW

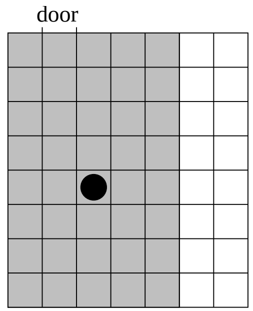

Belief space: after WWW

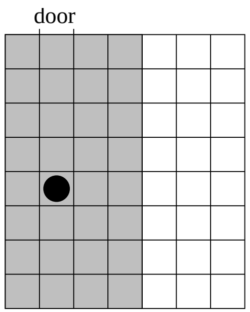

Belief space: after WWWW

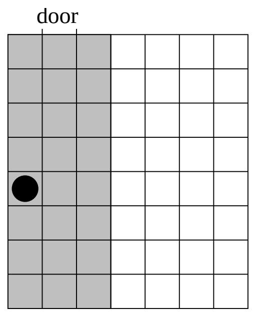

Belief space: after WWWWW

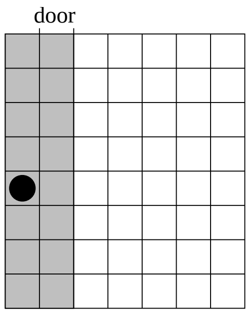

Belief space: after WWWWWW

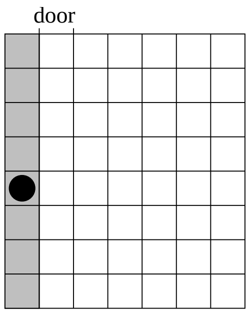

Belief space: after WWWWWWN

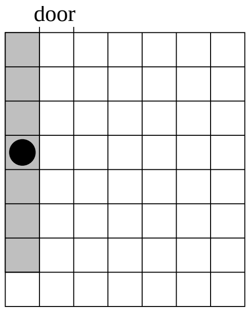

Belief space: after WWWWWWNN

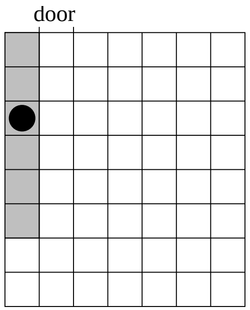

Belief space: after WWWWWWNNN

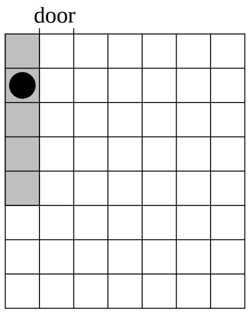

Belief space: after WWWWWWNNNN

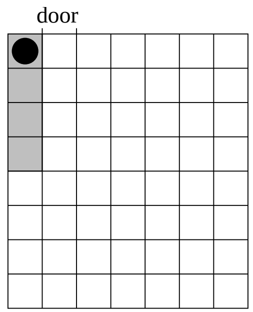

Belief space: after WWWWWWNNNNN

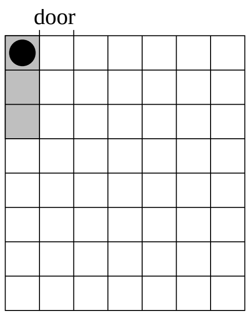

Belief space: after WWWWWWNNNNNN

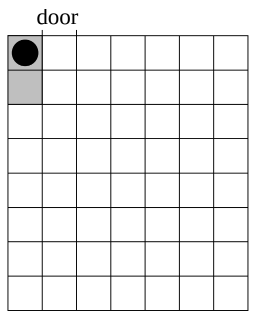

Belief space: after WWWWWWNNNNNNN

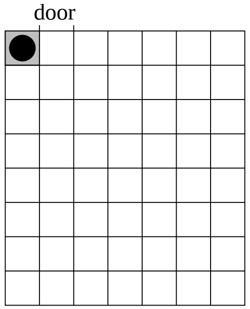

Belief space: after WWWWWWNNNNNNNE

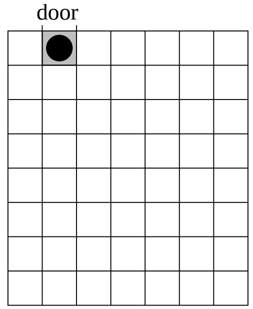

Belief space: after WWWWWWNNNNNNNEN

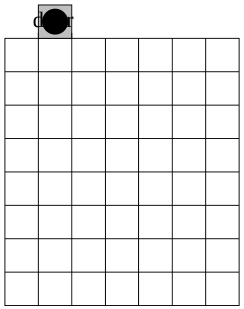

Planning under unobservability

1. Find a sequence of operators $o_1, \ldots, o_n$ that reaches a goal state in G starting from any initial state in I.
2. Find a sequence of operators $o_1, \ldots, o_n$ that reaches a belief state $S \subseteq G$ starting from the initial belief state I.

   Edge from belief state Z to Z' if o is applicable in every state in Z and *img*$_o(Z) = Z'$.

Regression for nondeterministic operators

1. Nondeterministic operator $o = \langle c, e\rangle$ with $e = p_1 e_1|\cdots|p_n e_n$ in *normal formal II* so that every e_i is deterministic.
2. $\text{regr}_{\langle c,e\rangle}(\phi)$ is now defined as $\text{regr}_{\langle c,e_1\rangle}(\phi) \wedge \cdots \wedge \text{regr}_{\langle c,e_n\rangle}(\phi)$.

THEOREM. If $S' = \{s'|s' \models \phi\}$, then *spreimg*$_o(S') = \{s|s \models \text{regr}_o(S')\}$.

PROOF IDEA: To guarantee reaching S' from a state s, S' must be reached from s with every effect $e_1, \ldots, e_n$.

Regression for nondet. ops: example

$o = \langle A, (0.5B|0.5\neg C)\rangle$

$$\begin{aligned}\text{regr}_o(B \leftrightarrow C) &= \text{regr}_{\langle A,B\rangle}(B \leftrightarrow C) \wedge \text{regr}_{\langle A,\neg C\rangle}(B \leftrightarrow C)\\ &= (A \wedge (\top \leftrightarrow C)) \wedge (A \wedge (B \leftrightarrow \bot))\\ &= (A \wedge C) \wedge (A \wedge \neg B)\\ &= A \wedge C \wedge \neg B\end{aligned}$$

Algorithms for UO planning: heuristic search

Regression + heuristic search (A$*$, best first, ...)

- new problem: Testing $I \models \text{regr}_o(\phi)$ is co-NP-hard.
- heuristic 1: strong distances as in FO algorithms

  If only one initial state, distance estimation as in deterministic planning can be used, but even accurate estimates for deterministic planning may be very misleading.
- heuristic 2: size of belief state

Algorithms for UO planning: heuristic search

Use strong distances of states as a heuristic:

$$\begin{aligned}D_0 &= G\\ D_{i+1} &= D_i \cup \textstyle\bigcup_{o \in O} \textit{spreimg}_o(D_i) \text{ for all } i \geq 1\end{aligned}$$

A lower bound on plan length for belief state Z is j if $Z \subseteq D_j$ and $Z \not\subseteq D_{j-1}$.

This is an **admissible heuristic** (does not overestimate the cost).

Algorithms for UO planning: heuristic search

- Use the cardinality of belief states as a heuristic.
- Backward search: Prefer operators that *increase the size of the belief state*, i.e. find a plan suffix that reaches a goal state from more starting states.
- Forward search: Prefer operators that *decrease the size of the belief state*, i.e. reduce the uncertainty about the current state and make reaching goals easier.
- This heuristic is **not admissible**.