- Homework 2B due today
- Midterm Exam 1 week from Thursday
- Watch for Homework 3
  - Not collected or graded (but do it)
  - Watch for posted solutions

State Space Planner

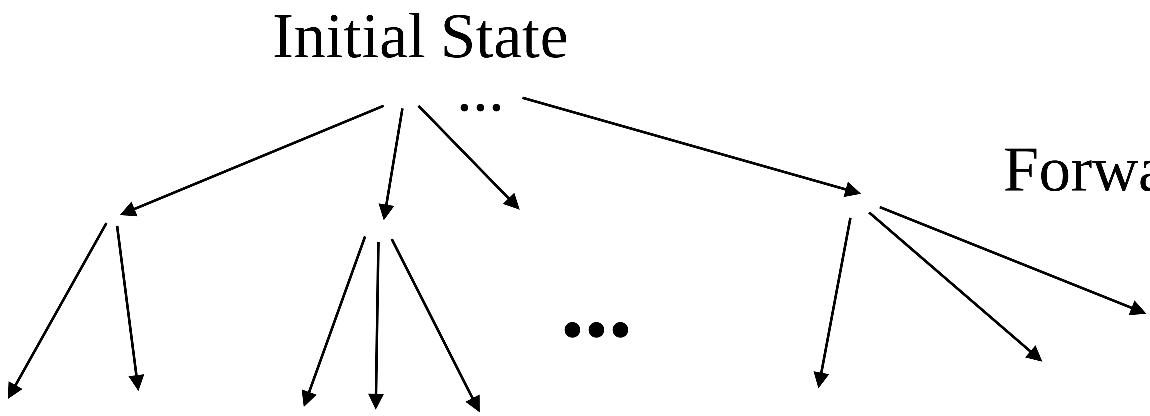

Forward Search for Goal

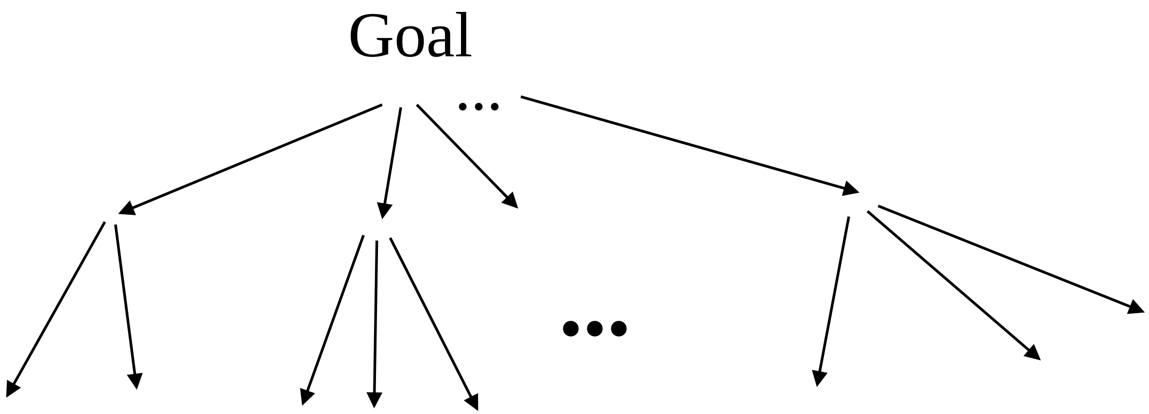

Backward Search for subset of Initial State fluents

One reason why planning beats searching

State Space Planners

- Intuitive
- Efficiency can be problematic
  - Interacting conjunctive subgoals
  - Often heuristics can help
- Linear / Total Order planners
- Usually incomplete
  - Sussman Anomaly
  - Seldom expose all planning decisions
  - Action (plan step) scheduling decisions

Plan Space Planner

partial-order planner; nonlinear planner;...

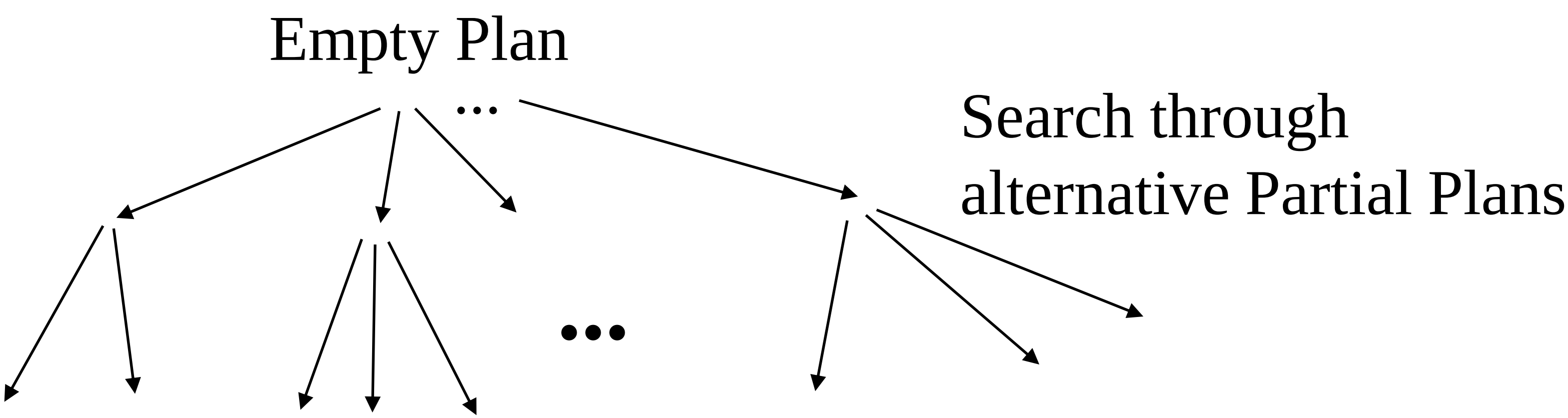

Partial plan ≡ set of constraints

Constraint set denotes all action sequences that satisfy its constraints

Empty plan ≡ all action sequences

Search through alternative constraints for a partial plan that achieves the goal

Alternative constraint vocabularies: Operator selection, Codesignation, Non-codesignation, Protection, Ordering…

Nonlinear Planners

- More elegant theoretically
- Can be complete
- Can also be inefficient
  - Redundancies
  - SNLP: Systematic Nonlinear Planner
  - Can cost as much to check for redundancy
- Heuristics can help

Planning Heuristics

- Domain independent heuristics are weak
  - Strong heuristics are often domain specific
  - Conflicts w/ domain independence
- Automatically learn heuristics
- Planning competitions at ICAPS

Learning

domain dependent planning knowledge
search control heuristics

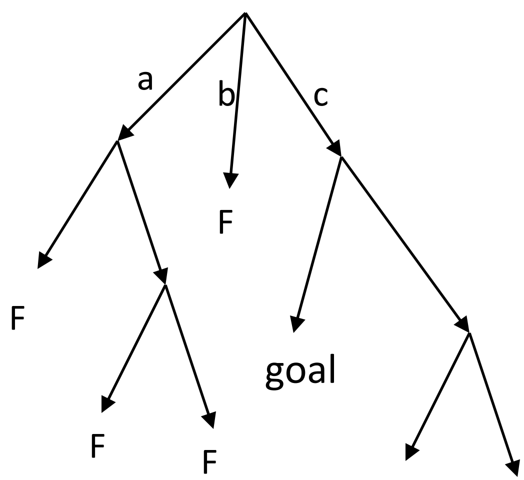

Imagine a preference function

Called with a set of alternatives

Returns a ranking or a permuted list or …

All failures under “a” and “b” $\rightarrow$ a training episode

Statistically characterize failures, next similar situation prefer “c” to “a” or “b” at top node

Learn to build towers from the bottom up

Learn to avoid using goal blocks as temporary supports

Desirable heuristic:
Unstacking all relevant blocks to the table realizes a factor of two from optimal (difficult / impossible to learn)

Propositionalization
(make planning propositional)

- Give up soundness
  - Gain (great) efficiency
  - Need to check solutions
- Strong results in Graph Theory and Satisfiability
- Cannot be applied at the first order model level
  - Infinite sets are problematic
  - Function symbols
  - "Everyone has a father" infinite number of elements
- But finite at the computation level
  - Herbrand Universe: all ground instances (can be infinite)
  - If F.O. Δ is unsatisfiable, there is a finite derivation of { } from its Herbrand universe

GraphPlan

- Planning Graph:
  - Series of “Levels” each w/
  - Reachable ground instance fluents
  - Possible actions (instantiated operators)
  - Constraints (as in nonlinear planning)
    - Mutual exclusion: mutex On(A,B) & Clr(B)
    - Identification of Effects to Preconditions
- Determine if
  - A solution exists
  - No solution is possible
  - Else add a level
- Heuristics are very important for efficiency

Paradigm Problems
w/ Classical Planning

Logic's semantic brittleness
- what if axioms are approximate?
- refutational inference
- others; paraconsistency

World Uncertainty
- initial state
- sensors
- effectors

Inferential Uncertainty
- Frame Problem
- Ramification Problem – light switch
- Qualification Problem – birds

Single Agent Assumption

Time
- temporal planning
- interacting overlapping actions

Continuous Change

Planning is undecidable
(not so bad - why?)

Planning is often inefficient
(worse - why?)

What’s Hard?

- Start your car
- Finding your new advisor’s office
- Hammering a nail into a board
- Landing an airplane
- Playing PacMan / the Sims / ...

Analytic Models

- Search, deduction, classical planning, …
- Uncertainty is a (perhaps ***the***) major problem in intelligent behavior
- An unfortunate combination:
  - Complex world
  - Analytic model
  - Semantics of logical inference (satisfiability, propagating constraints)

This results in brittleness

Hybrid Models

- Prior (*a priori*) commitments
  - Constraints
    - Always some hard constraints
    - Sometimes also soft constraints or preferences
  - Analytic
  - Determines / defines a family of models
  - Often parametrically related
- Access to Examples / Observations of the world
  - Training examples
    - Labeled by expert or the world itself
    - Supervised learning
  - Resolve remaining degrees of freedom to fit the examples
  - Calibrate parameters with the world
    - Select values
    - Estimate distributions
    - Posterior commitments

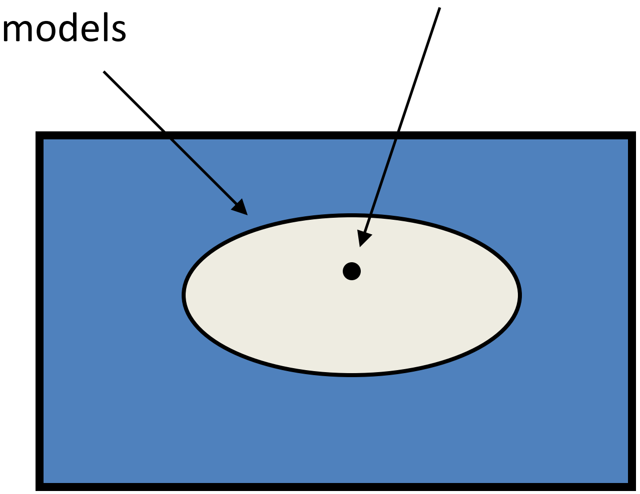

Inference mechanisms

Tolerant of small errors

Optimization

Reinforcement Learning

- Difficulty / impossibility of an adequate analytic world model
- Learning augments prior commitments
- Behaviorism (?)
- Control Theory
- Bellman’s Dynamic Programming
  - Plan = Policy (not action sequences)
  - Policy: State $\rightarrow$ Action
  - Find optimal policy (not a satisfiability problem)
  - Markov assumption
  - Markov Decision Process (MDP)

RL Model

- States are individuated by perception (i.e., intrinsic features) State = set of intrinsic features / Markov
- World is a finite exclusive and exhaustive set of states
  "Finite" is now required
- World changes are state transitions

  Same
- Actions may have probabilistic effects

  New action ontology
- Rewards (+ or -) occur probabilistically (for us, on state arrival)

  New goal ontology
- Learn how to act so as to maximize rewards

  New plan ontology (more a "policy")

What is new?

Grid World

ACTIONS

← Left
→ Right
↑ Up
↓ Down

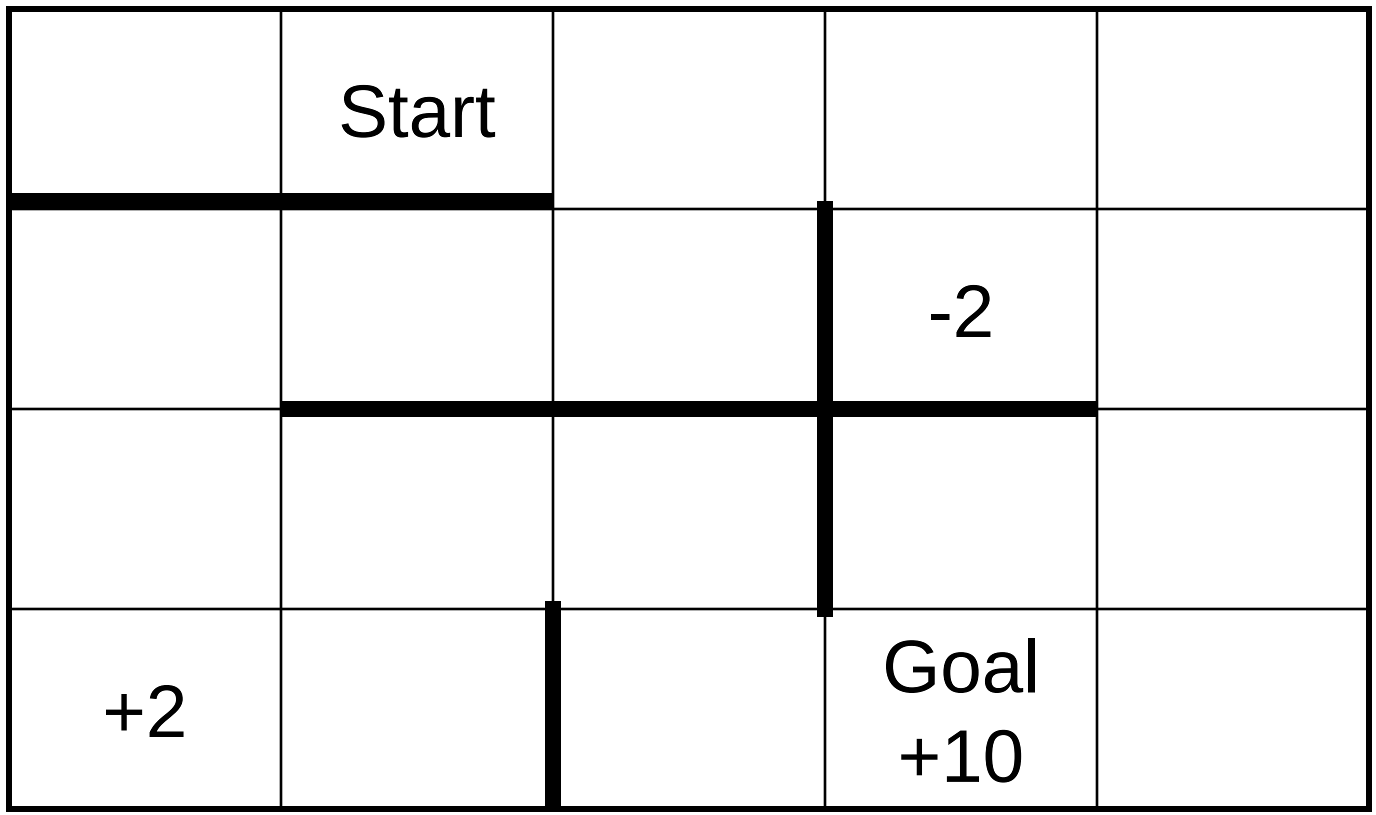

Grid World Policy

ACTIONS

← Left
→ Right
↑ Up
↓ Down