AAAI 2018
Workshop on Planning and Inference

Organizers:
Roni Khardon, Akshat Kumar, Alex Ihler

- Thanks to the PC

Christopher Amato, David Barber, Vaishak Belle, Alan Fern, Kristian Kersting, Qiang Liu, Radu Marinescu, Denis Maua, Sriraam Natarajan, Gerhard Neumann, Pascal Poupart, David Poole, Regis Sabbadin, Scott Sanner, Prasad Tadepalli, Jan-Willem van de Meent.

- Thanks to AAAI that managed all the non-academic aspects of the workshop

Overview

- Why? Lots of pieces in the literature. Work is disjoint. Not unified. Thanks to recent progress time is ripe for impact.
  - facilitate synergy and exchange of ideas
  - identify opportunities and challenges

Overview

- What?
  - equivalence/reductions
  - approximations/scalability
  - applications

Stochastic Planning and Probabilistic Inference

- Simplest setting: finite horizon, goal based formulation and atomic states and actions

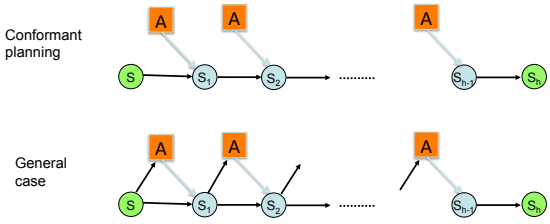

Stochastic Planning and Probabilistic Inference

- Simplest setting: finite horizon, goal based formulation and atomic states and actions

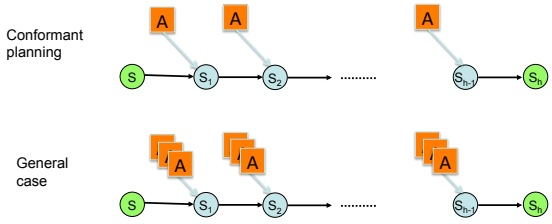

Stochastic Planning and Probabilistic Inference

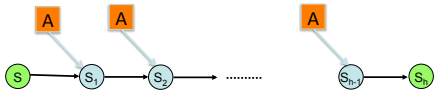

- P(goal|s_0, As) ~ P(As |s_0,goal) ~ P(As,goal |s_0)
- "Reward-weighted" distribution

$$\hat{p}(s_{1:t}, a_{1:t}, t|\pi) = \frac{r_t(s_t, a_t)(s_{1:t}, a_{1:t}|\pi)}{U(\pi)}$$

generalizes this to Infinite horizon, reward per stage, stationary/non-stationary policy

Stochastic Planning and Probabilistic Inference

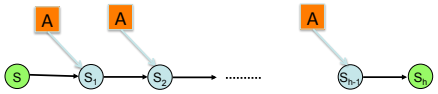

- Value iteration = alternating variable elimination

Stochastic Planning and Probabilistic Inference

- Factored states and actions

Conformant

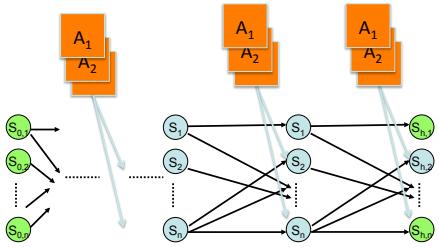

Stochastic Planning and Probabilistic Inference

- Factored states and actions

Non-stationary policy

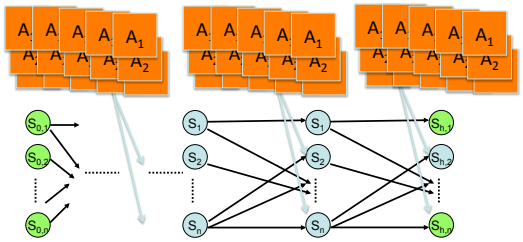

Stochastic Planning and Probabilistic Inference

- Factored states and actions

Stationary policy

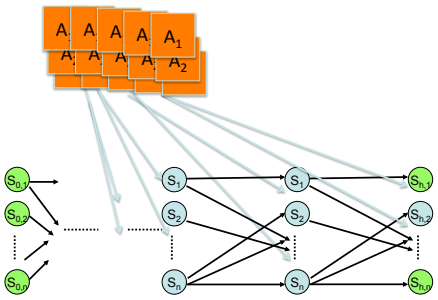

Approximation

- P(goal|s_0, As) ~ P(As |s_0,goal) ~ P(As,goal |s_0)
- "Reward-weighted" distribution
- Value iteration = alternating variable elimination
- …
- All of these are correct/exact
- Can this scale to large problems?
- How to best approximate?
- Interaction b/w reduction and approx

How to Solve?

- Symbolic (lifted) dynamic programming
- Weighted model counting
- Search: And/Or, Branch and bound
- Expectation Maximization
- Variational approximations
- Belief propagation
- Policy gradients
- Planning as Mixed Integer Linear Program
- Particle filters / Point based methods
- …

Problem Variants

- discrete vs. continuous
- single vs. multi-agent
- general vs. spatial problems
- propositional vs. relational
- model based planning vs. reinforcement learning
- exact/optimal vs. approximate vs. heuristic solutions

Applications

- Task planning in robotics
- Spatial management of invasive species
- Taxi fleet control
- Power flow in electric grids
- …

Overview

- Program:
  - keynote talks
  - invited presentations
  - contributed papers
- Why? Lots of pieces in the literature. Work is disjoint. Not unified. Thanks to recent progress time is ripe for impact.
  - facilitate synergy and exchange of ideas
  - identify opportunities and challenges

    → Please propose challenges/opportunities for all

Agenda

- Equivalence/reductions

  what works? when/why?
- Approximations/scalability

  what works? when/why?
- Applications

  what can we already do? what properties of problems/solutions are important in real-world large problems?