Distributional Methods for Heuristic Search

University of New Hampshire

Tianyi Gu (advisor: Wheeler Ruml)

Motivation

- Distribution is more informative than a scalar value
- Bring the idea of explicitly estimating value uncertainty from RL to Heuristic Search

Contribution

AAAI-20: Data-Driven Nancy

- using data to estimate value uncertainty distribution
- using estimates to guide real-time heuristic search

IJCAI-21: Expected Effort Search (XES)

- using estimates to guide bounded-cost heuristic search

Improve Real-time Heuristic Search

How to Gather Information?

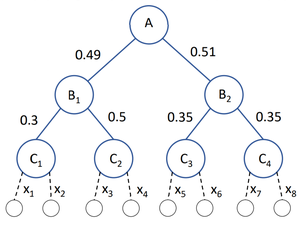

Given these search nodes, should agent at A move to B_1 or B_2

Which Node to Expand?

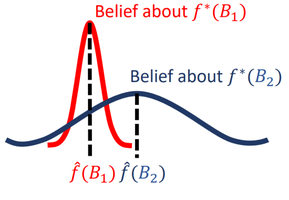

Should the agent expand nodes on the frontier under B_1 or B_2

Searching use beliefs

Risk-based Expansion: given beliefs about top-level action values, expand nodes on the frontier under top-level action that minimizes risk, the expected regret

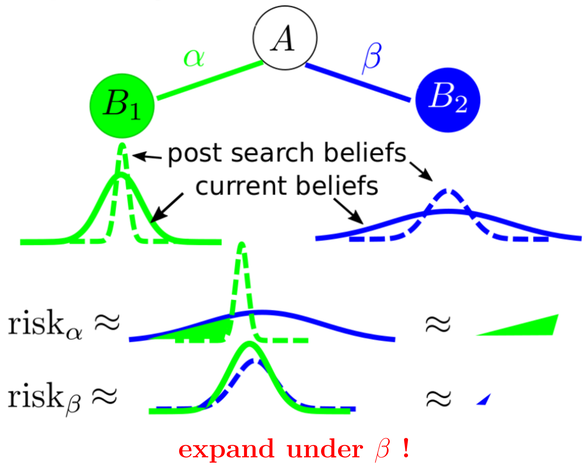

expand under β !

Where do beliefs come from?

Purpose of search is to gather information to inform decision-making process. Which information on the search frontier should be used to form beliefs about top-level actions?

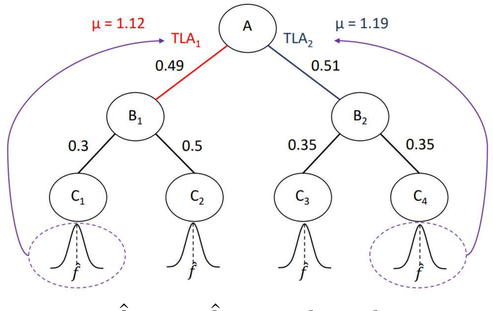

$$\hat{f} = g + \hat{h} = g + h + \epsilon d$$

- Assumption-based Nancy: Trucated Gaussian based on h and d
- Data-Driven Nancy: replace the assumptions with data

Improve Bounded-cost Search

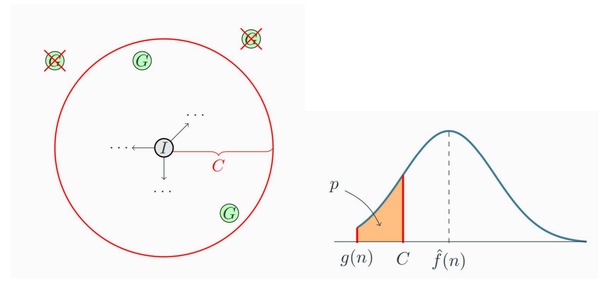

Left: the problem setting. Right: the probability of finding a solution within bound.

Previous Approaches:

- PTS: does not consider search effort
- BEES: does not consider the uncertainty of its estimates

Our approach: Expected Effort Search

1. estimate the probability of finding a solution within the bound $p(n)$
2. estimate the search effort $T(n)$
3. best-first search on the expected search effort $\frac{T(n)}{p(n)}$

Results:

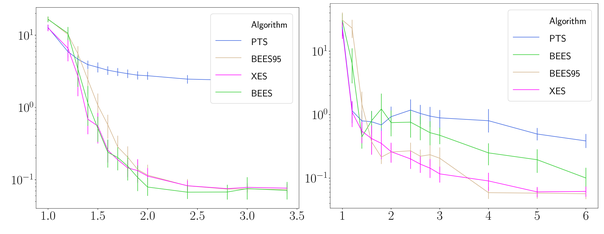

CPU time (in seconds) as a function of the cost bound (factor of optimal). Error bars show 95% confidence intervals on the mean across the commonly solved instances. Left: heavy vacuum world. Right: heavy sliding tile puzzle.