New Results in Bounded-Suboptimal Search

Maximilian Fickert, Tianyi Gu, Wheeler Ruml

SAARLAND UNIVERSITY
SAARBRÜCKEN GRADUATE SCHOOL OF COMPUTER SCIENCE

University of New Hampshire

The Problem Setting

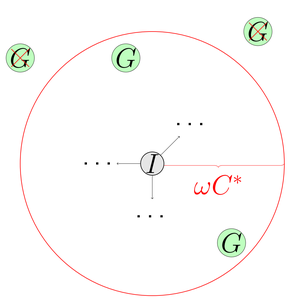

- Motivation: many real-world problems are too hard to solve optimally. Need bounded-suboptimal solution!
- Input: initial state (I), goal states (G), and a suboptimality bound $\omega \geq 1$.
- Objective: Find a solution with cost at most $\omega \cdot C^*$ as fast as possible.

Previous Approaches

- **Weighted A***
- **Explicit Estimation Search (EES)**[1]
  → focal search:
  - *cleanup* sorted by f, all generated nodes
  - *open* sorted by $\hat{f}$, only nodes with $g + \hat{h} \leq \omega \cdot f_{min}$
  - *focal* sorted by d, only nodes with $g + \hat{h} \leq \omega \cdot f_{min}$
    ($\hat{h}$ is corrected for the observed heuristic error[2])
- **Dynamic Potential Search (DPS)**[3]
  → best-first search on $\frac{\omega \cdot f_{min} - g(n)}{h(n)}$

[1] Thayer and Ruml, 2011.
[2] Thayer, Dionne, and Ruml, 2011.
[3] Gilon, Felner, and Stern, 2016.

New Algorithms: 1/2 Dynamic Expected Effort Search (DXES)

Best-first search on the expected effort:[4] $\frac{T}{p}$

- $T(n)$: search effort to find a solution under n
- $p(n)$: probability that n leads to a solution within estimated bound

n_1: $T = 10$, $p = 0.5$ $\rightsquigarrow 20$ $\qquad$ n_2: $T = 6$, $p = 0.25$ $\rightsquigarrow 24$

How to obtain T and p?

- Obtaining T use distance-to-go d
- Obtaining p from belief distributions

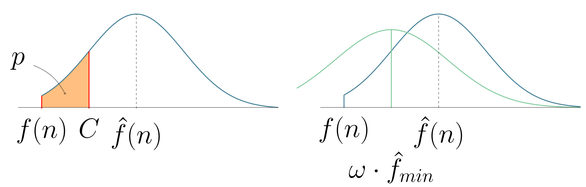

hard to estimate when raising the bound is useful!

New Algorithms: 2/2 A Round-Robin Scheme

Replace EES selection rule with alternating queues[5] :

- **open list:** sorted by $\hat{f}$
- **cleanup list:** sorted by f
- **focal list:** sorted by d(EES) or ud(DPS) or xe(DXES)

[4] First suggested by Dobson and Haslum (HSDIP'17)
[5] Helmert and Roger, 2010

Experiments

Planning Domains: IPC optimal tracks (48 domains)

Coverage	WA*	EES	DPS	DXES	RR-DPS	RR-d	RR-DXES
Sum (1652)	995	967	1012	894	982	1025	**1052**
Normalized(%)	58.7	57.0	60.0	51.5	57.9	60.7	**62.5**
Expansions	569	558	472	734	665	383	**371**

Search Domains:

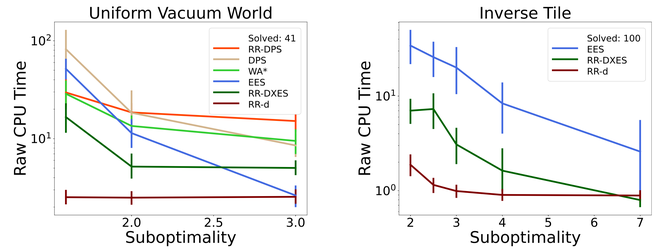

RR-d and RR-DXES perform best overall.

(see paper for more results)

Conclusion

- Weighted-A* is the first thing to try
- Round-Robin on d, $\hat{f}$, f is the next to try
- Round-Robin on xe, $\hat{f}$, f performs well in some domains

Still unresolved:

- When to raise bound, and when to pursue solution?
- How to best use belief distribution in bounded-suboptmal search?