PRP Rebooted

mulab.ai/pr2

Advancing the State of the Art in FOND Planning

Christian Muise [1,3] *Sheila A. McIlraith* [2,3] *J. Christopher Beck* [2]

[1] Queen's University, Kingston, Canada
[2] University of Toronto, Toronto, Canada
[3] Vector Institute for AI, Toronto, Canada

Building on the rich history of planning technology for non-deterministic domains, we introduce a new planner, **PR2**, which significantly outperforms all previous planners across nearly every existing benchmark domain.

Fully **O**bservable **N**on-**D**eterministic

Application Areas

- Robotics with Uncertainty
- Probabilistic Planning
- Multi-agent Planning
- Contingent Planning
- Reactive Synthesis
- Dialogue Agents
- ...

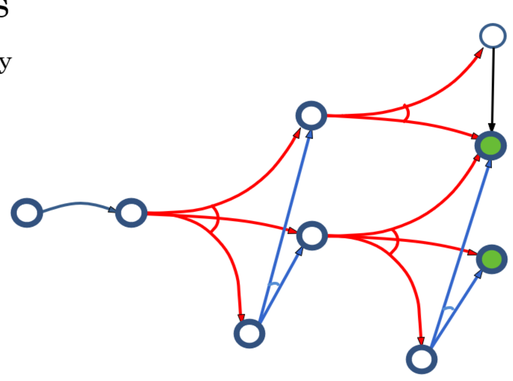

History of FOND Planners

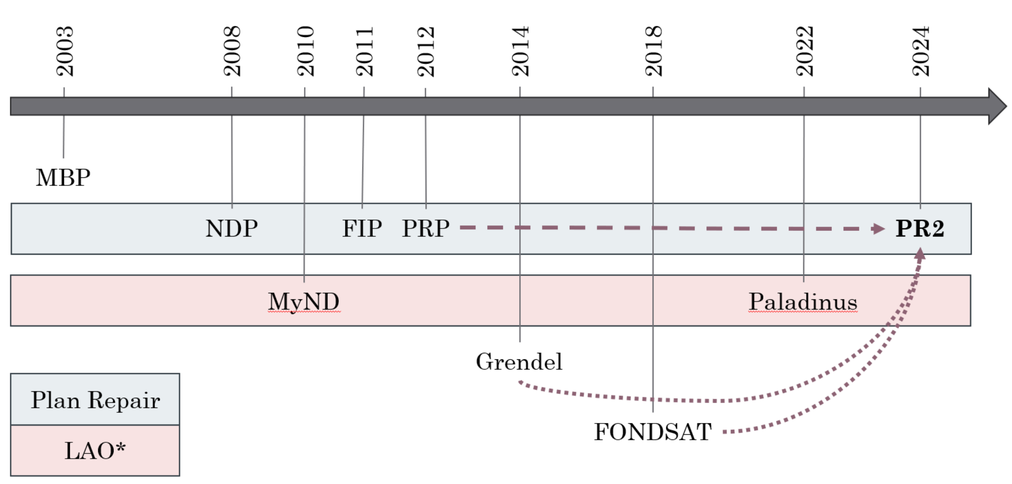

Algorithm 1: PR2 High-Level Planner

Input: FOND planning task, $\Pi = \langle \mathcal{V}, s_0, s_*, \mathcal{A} \rangle$
Output: Policy

```
incumbent = make_empty_solution();   FSAPS = ∅;
while !incumbent.is_strong_cyclic() do
    sol = make_empty_solution({s_0});
    while sol.REACHABLE contains unhandled nodes do
        n = sol.REACHABLE.pop_unhandled_node();
        switch analyze_node (n) do
            case 0: skip_if_strong_cyclic(n)
            case 1: skip_if_poisoned(n)
            case 2: match_complete_state(n)
            case 3: apply_predefined_path(n)
            case 4: match_complete_state(n)
            case 5: find_and_update_weak_plan(n)
            case default (case 6) do
                record_deadend(n);
                if n.state == s_0 then
                    return make_policy(incumbent.CONTROLLER);
    if sol.success_rate() ≥ incumbent.success_rate() then
        incumbent = sol;
return make_policy(incumbent.CONTROLLER);
```

How was it achieved?

Building off PRP

- Generalized state representation
- Logical regression
- Deadend detection and avoidance
- Rapid FOND plan repair/replanning

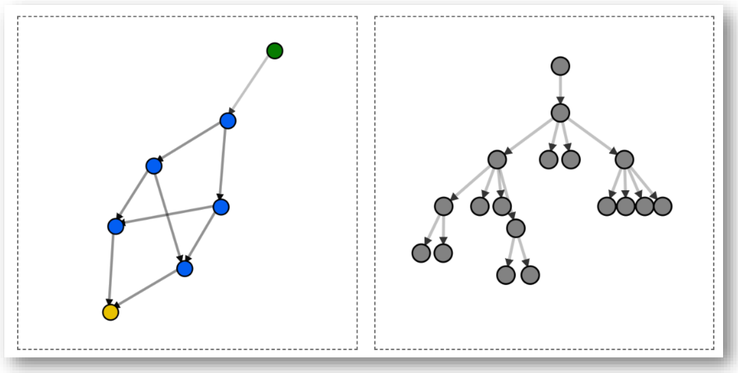

New Techniques

- Powerful solution representation
  - Suite of algorithms to process and leverage this representation
- "Poisoning" to identify less-promising areas of the search.
- FOND deadend-inspired heuristic
- Object sub-sampling to reduce search

New State of the Art

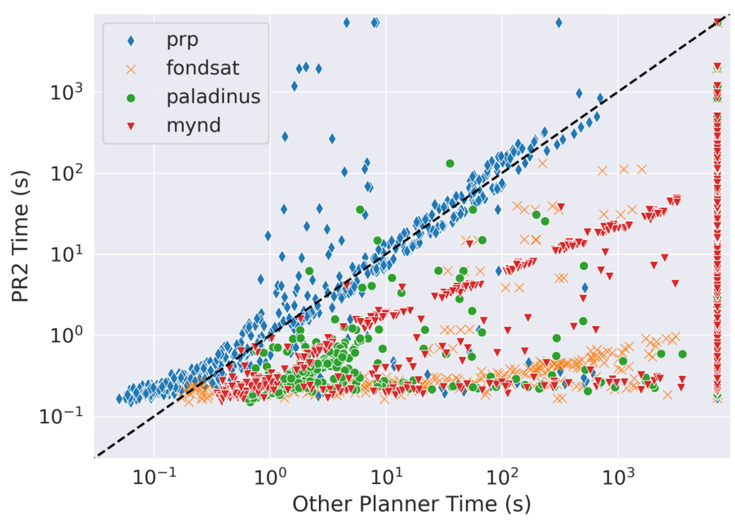

(a) Time comparison.

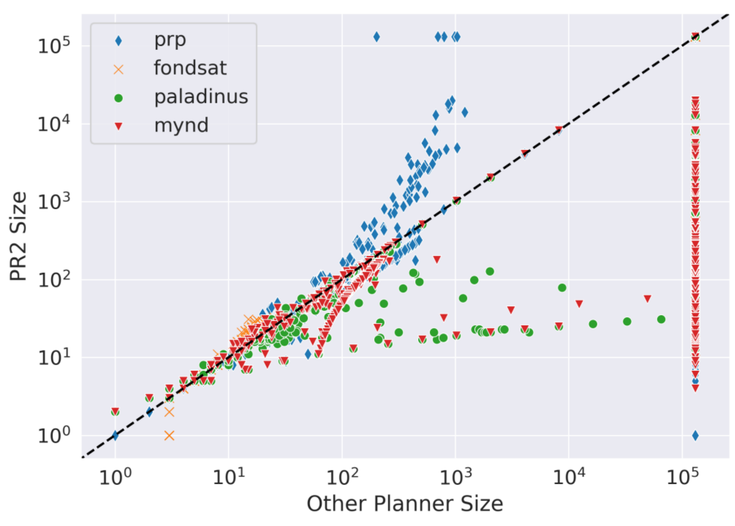

(b) Size comparison.

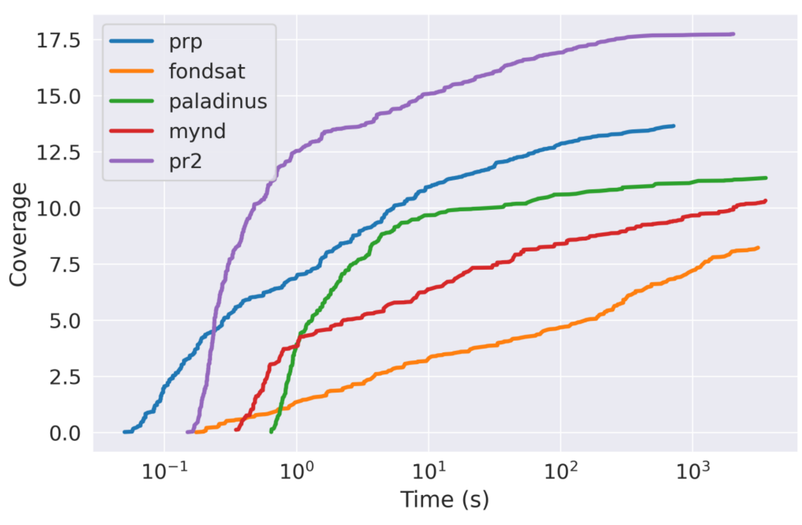

(c) Coverage over time.