FY24 Strategic University Research Partnership (SURP)

Fast planning under uncertainty with explicit operational and safety guarantees

Principal Investigator: Federico Rossi (347)

Co-Investigators: Martin Feather (512), Zachary Sunberg (University of Colorado at Boulder), Morteza Lahijanian (University of Colorado at Boulder)

Goal and Objective

Develop efficient algorithms for *planning under uncertainty* with explicit operational and *safety guarantees*

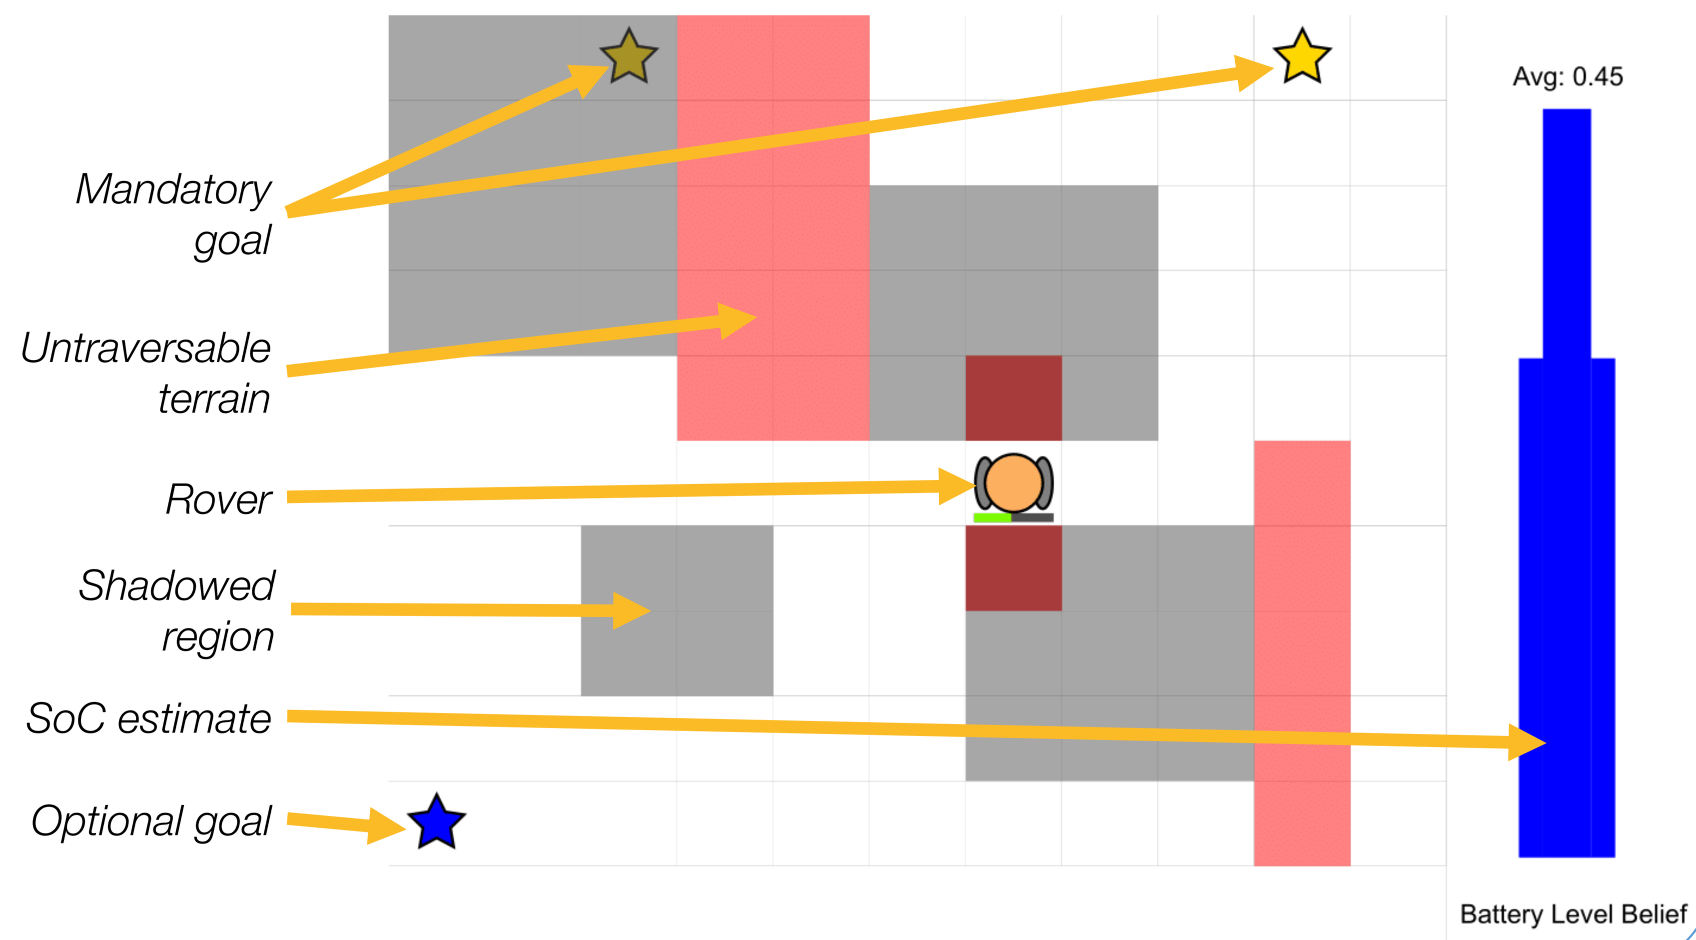

Background

Planning under uncertainty: POMDPs

- Efficient approximate algorithms
- No safety constraints

Safety-Aware POMDPs

- Assign extra penalty to "bad" end states
- Needs handcrafted penalties
- Asymptotic guarantees

Constrained POMDPs

- Provide rigorous way of describing complex constraints
- Algorithms are not scalable with many assumptions
- Violates Bellman's principle of optimality, resulting in pathological behavior

Our contribution: Recursively-Constrained POMDPs (RC-POMDPs) & reachability objectives

- Provide rigorous way of describing complex constraints
- Finite sample guarantees
- Scalable algorithms
- Good behaving algorithms

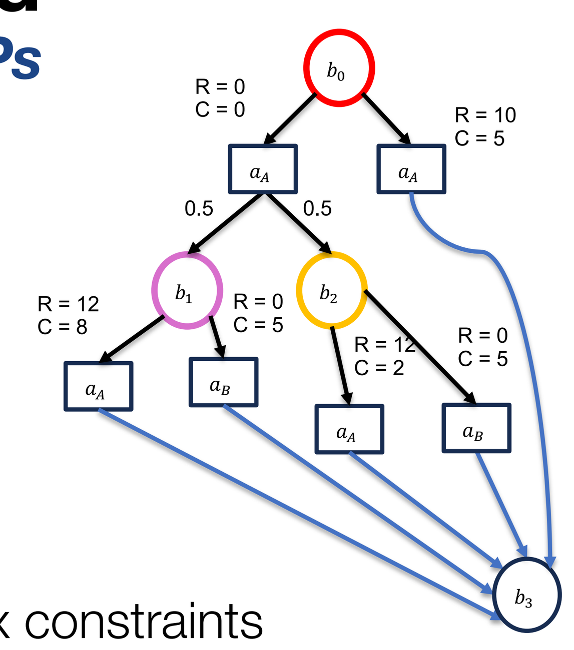

Approach: **POMDPs with Undiscounted Reachability Objectives**

Maximal Reachability Probabilities Problem (MRPP): Find a policy that maximizes probability of reaching set of target states

$$\pi^* = \arg\max P^{\pi}(\Diamond T)$$

Heuristic Search Value Iteration – Reachability Probabilities (HSVI-RP)

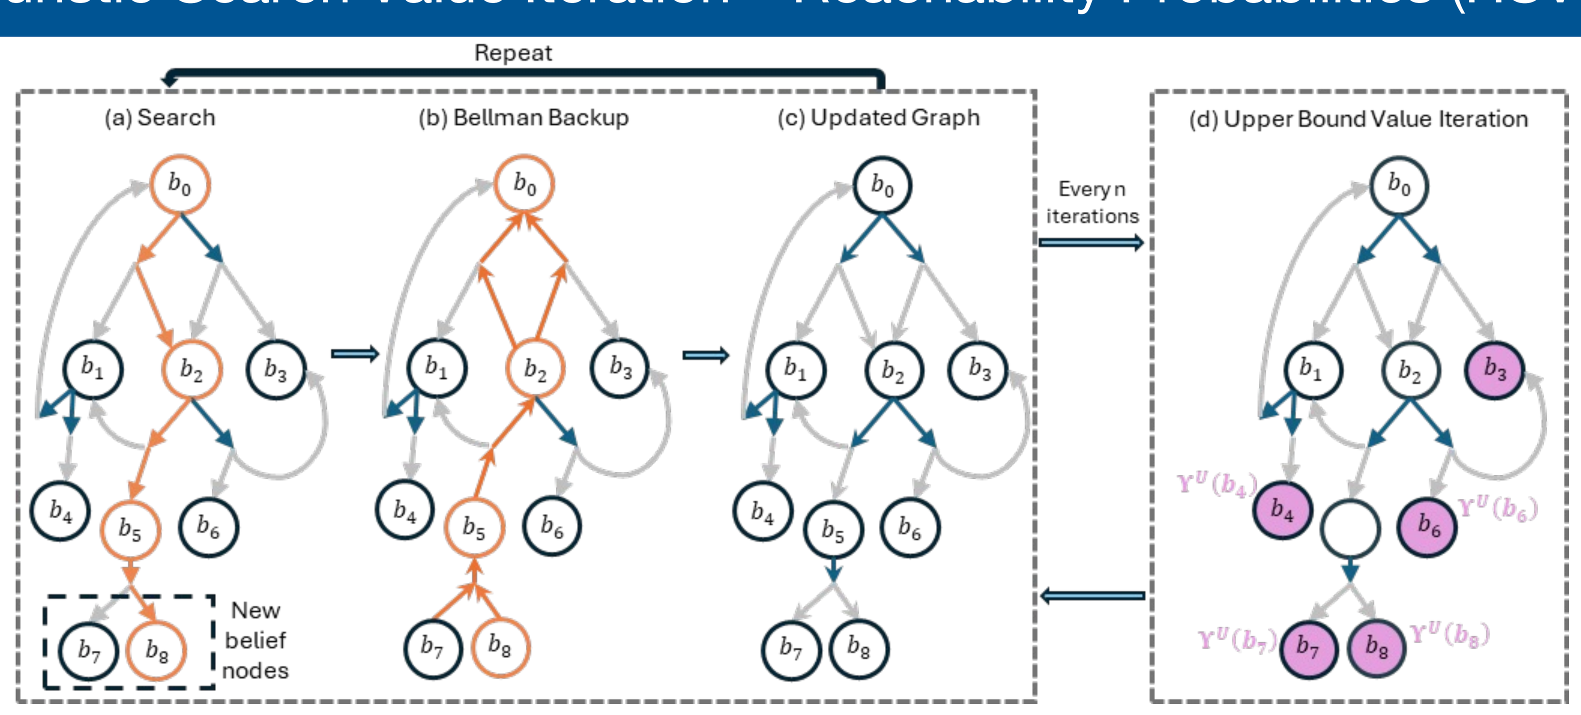

- Sound Upper and Lower Bounds: **Guarantees** on solution correctness
- Under-approximation (lower bound) **converges** to **optimal** solution

Probabilistic Temporal Logic Shields

- Compute upper and lower probability bounds for scLTL satisfaction
- Allow actions with lower bound above threshold
- Disallow actions with upper bound below threshold
- Enhance existing autonomy stack (e.g. online planning)

Safe Human-in-the-loop Control

- Safety advisor dissuades bad actions to human operator from probabilistic LTL constraints

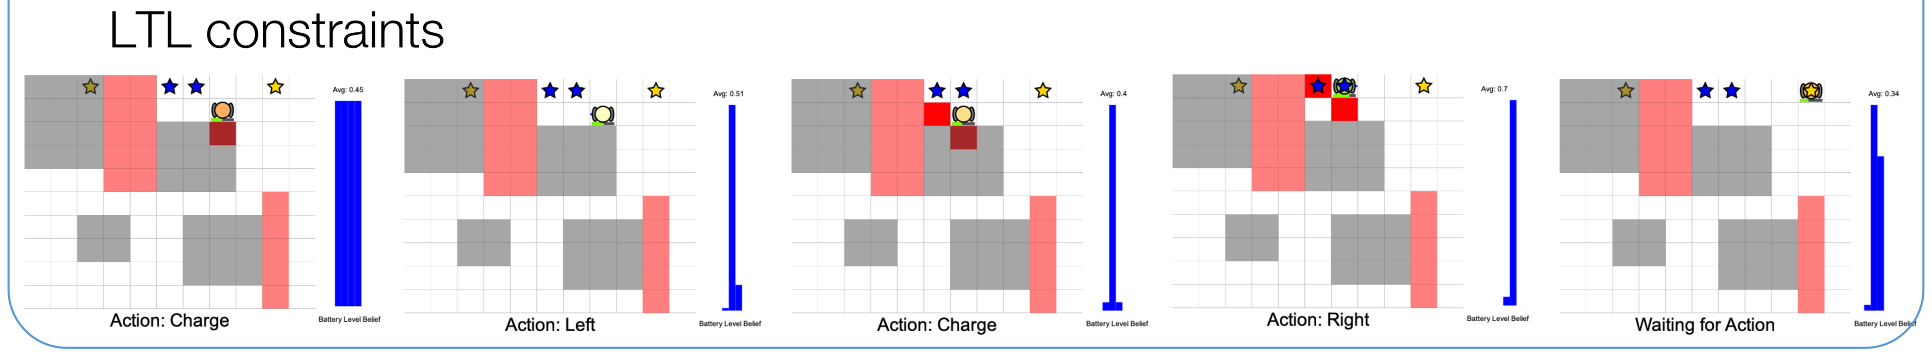

Approach: **Enhancing Online Planning and Human-in-the-Loop Control**

Co-safe Linear Temporal Logic (scLTL)

Complex task and operational specifications:

$$\phi := p \mid \neg\phi \mid \phi \vee \phi \mid X\,\phi \mid F\,\phi \mid \phi\, U\, \phi$$

$$\phi = \neg\,\text{deplete}\ U\ (\text{science} \wedge (\neg\,\text{depleted}\ U\ \text{unshadowed}))$$

Probabilistic Temporal Logic RC-POMDP

$$\pi^* = \operatorname{argmax}_{\pi} V_R^{\pi}(b_0)$$
$$\text{s.t. } P_{\phi}^{\pi}(b_0) \geq p$$

- Our proposed RC-POMDP algorithm can be used with HSVI-RP with small modifications for probabilistic scLTL constraints
- **Guarantees** on probabilistic scLTL constraint satisfaction

Benefits to JPL and NASA

Uncertainty is ubiquitous

- Environmental conditions (e.g., terrain-wheel interaction, air density)
- Actuator performance (e.g., thruster valve timing, heater performance)

Planning under uncertainty can increase science returns

From conservative, *worst-case* margins to explicit, *tight* representation of uncertainty

Plug-and-play with existing autonomy stacks

Planning with formal guarantees helps performance as well as V&V

- Formal guarantees span the *wide range of conditions* in which the autonomy is intended to operate
- Testing explores in detail *specific scenarios*
- Formal guarantees *complement* testing to help instill confidence that the autonomy will perform correctly and sufficiently rapidly

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

www.nasa.gov

Clearance Number: CL24-XXXX
Poster Number: RPC#182

Publications:

Qi Heng Ho, Tyler Becker, Ben Kraske, Zakariya Laouar, Martin Feather, Federico Rossi, Morteza Lahijanian, and Zachary N. Sunberg, "Recursively-Constrained Partially Observable Markov Decision Processes", Conf. Uncertainty in Artificial Intelligence, Barcelona (Spain), 2024.
Qi Heng Ho, Martin S. Feather, Federico Rossi, Zachary Sunberg, and Morteza Lahijanian, "Sound and Efficient Algorithms for POMDPs with Reachability Objectives via Heuristic Search", Conf. on Uncertainty in Artificial Intelligence (UAI), Barcelona (Spain), 2024.

PI/Task Mgr. Contact Information:

Federico Rossi, federico.rossi@jpl.nasa.gov