Rover Science Autonomy

Probabilistic Planning for Science-Aware Exploration

Thesis Proposal

Trey Smith, Robotics Institute, Carnegie Mellon University

14 June 2004

Robotic Planetary Science

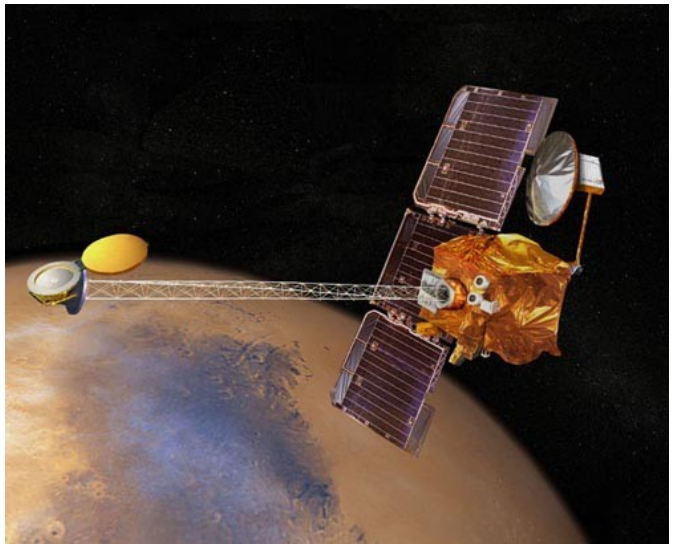

Mars Rover Operations Today

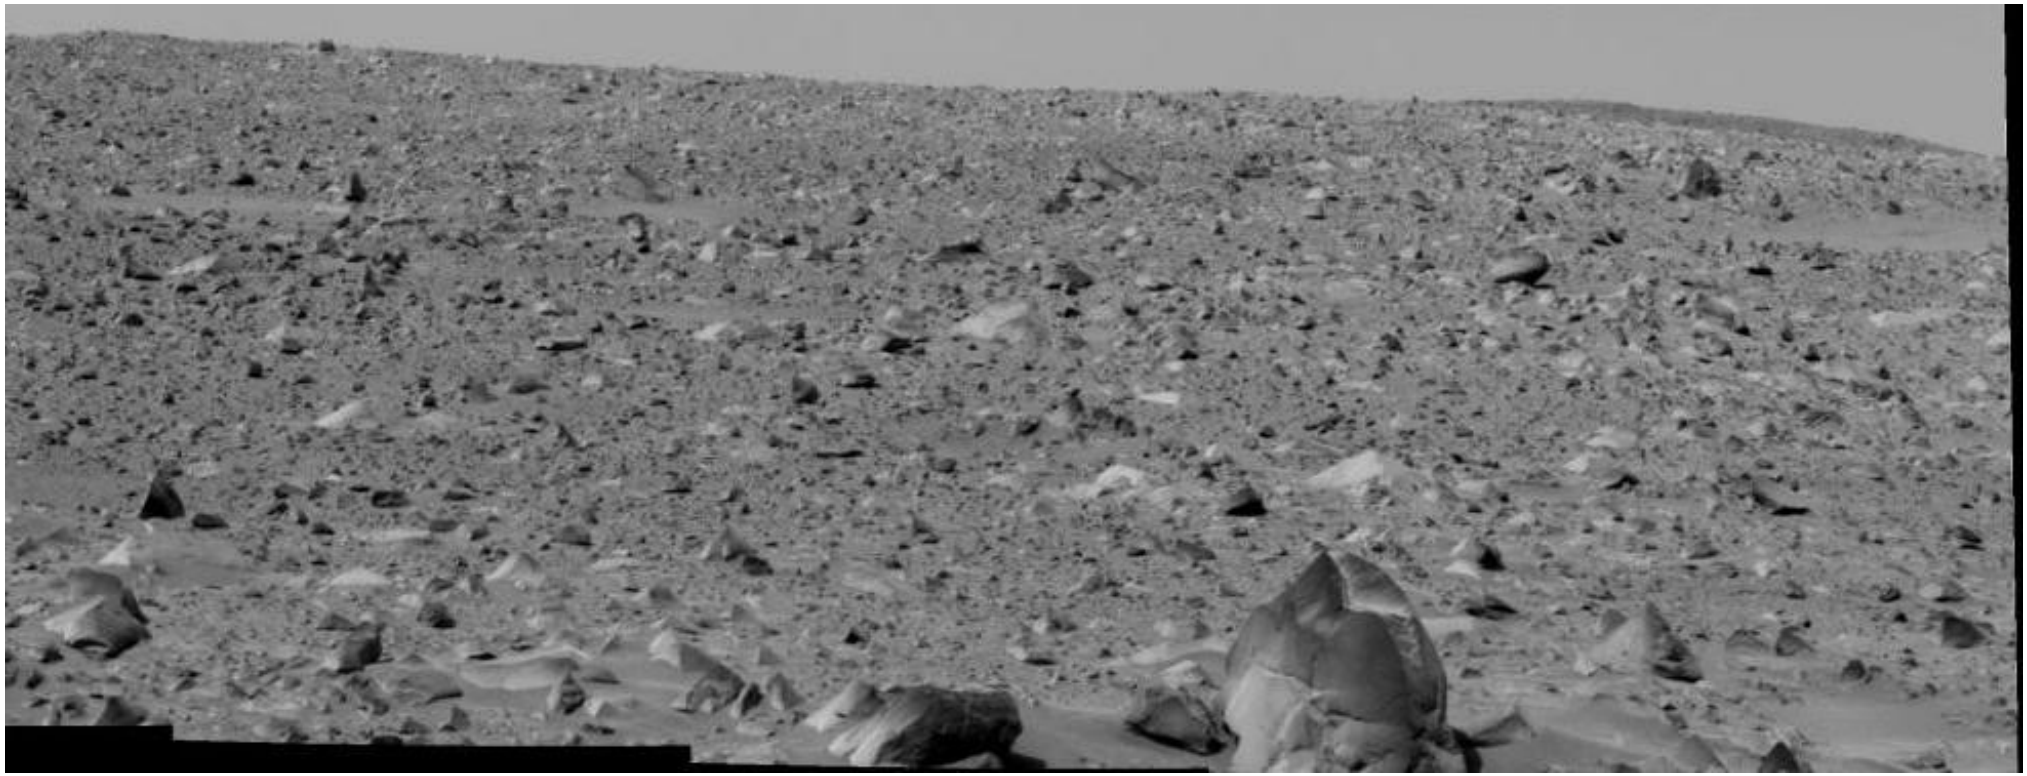

Mars Rover Operations Today

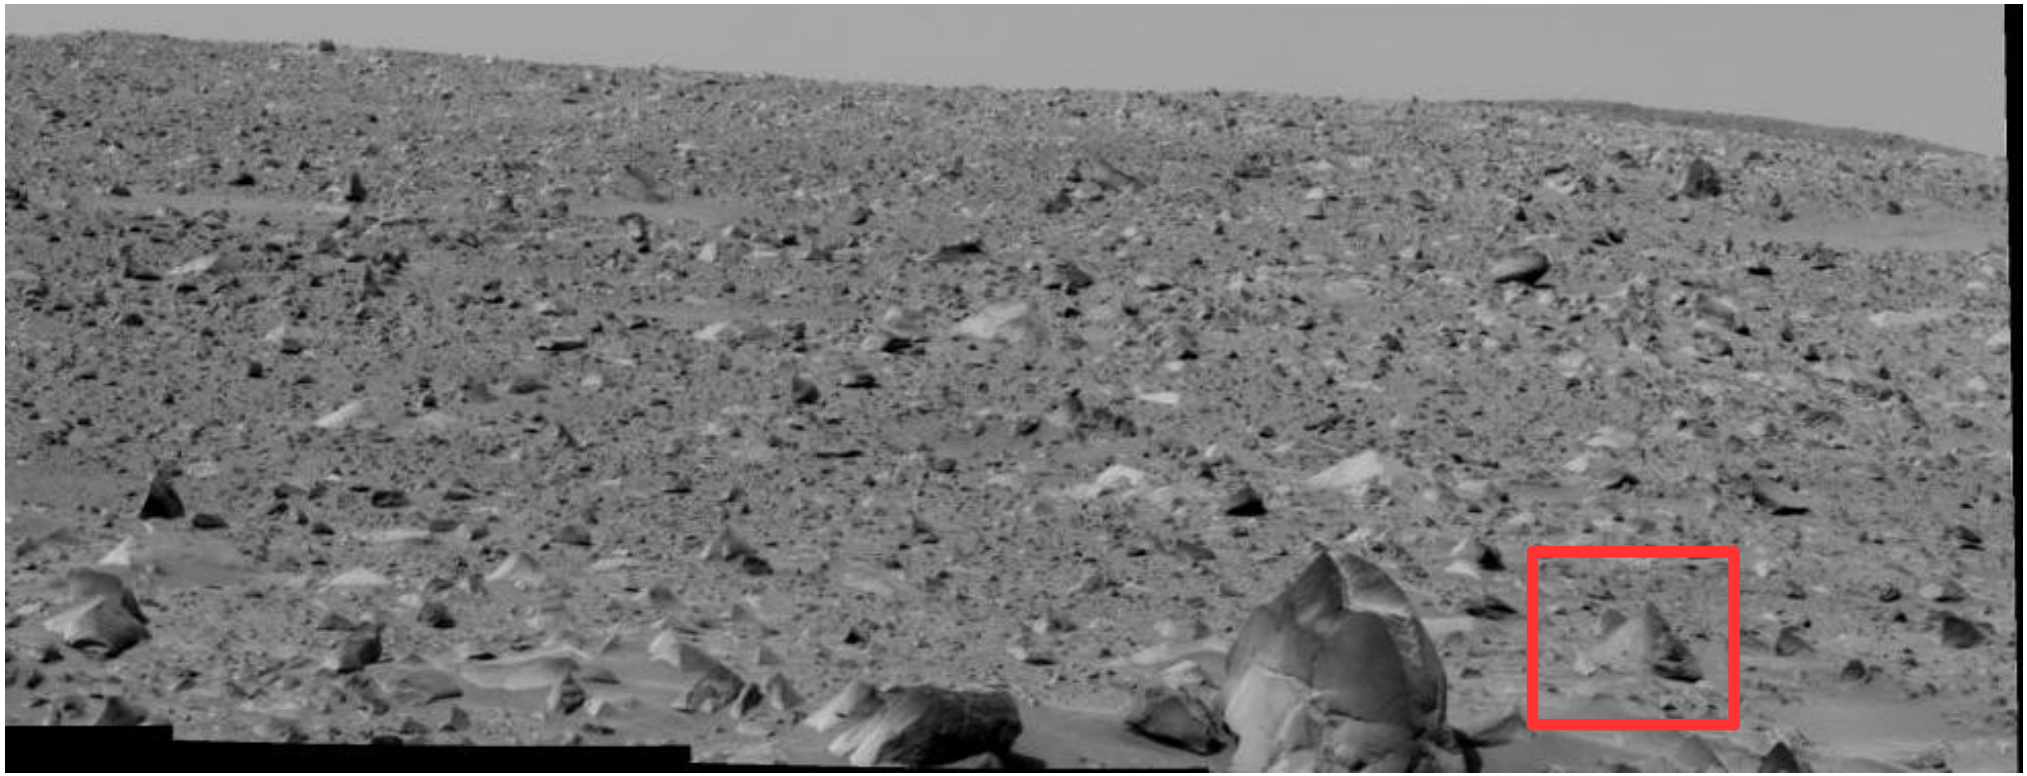

Mars Rover Operations Today

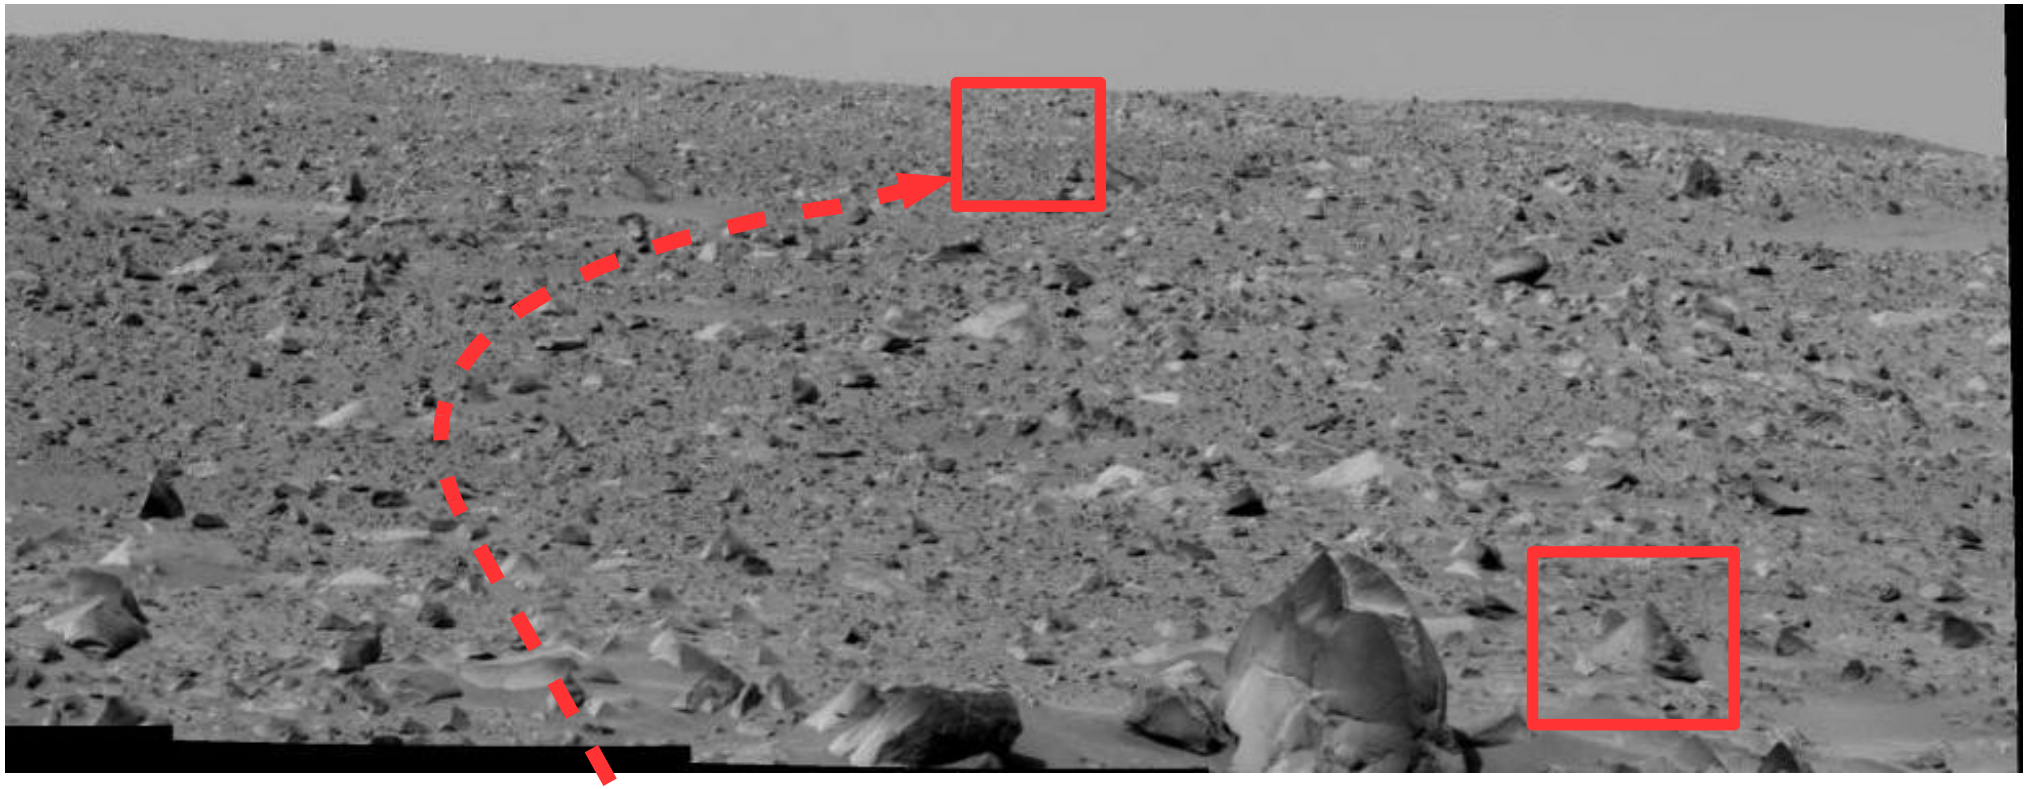

Over-the-Horizon Science

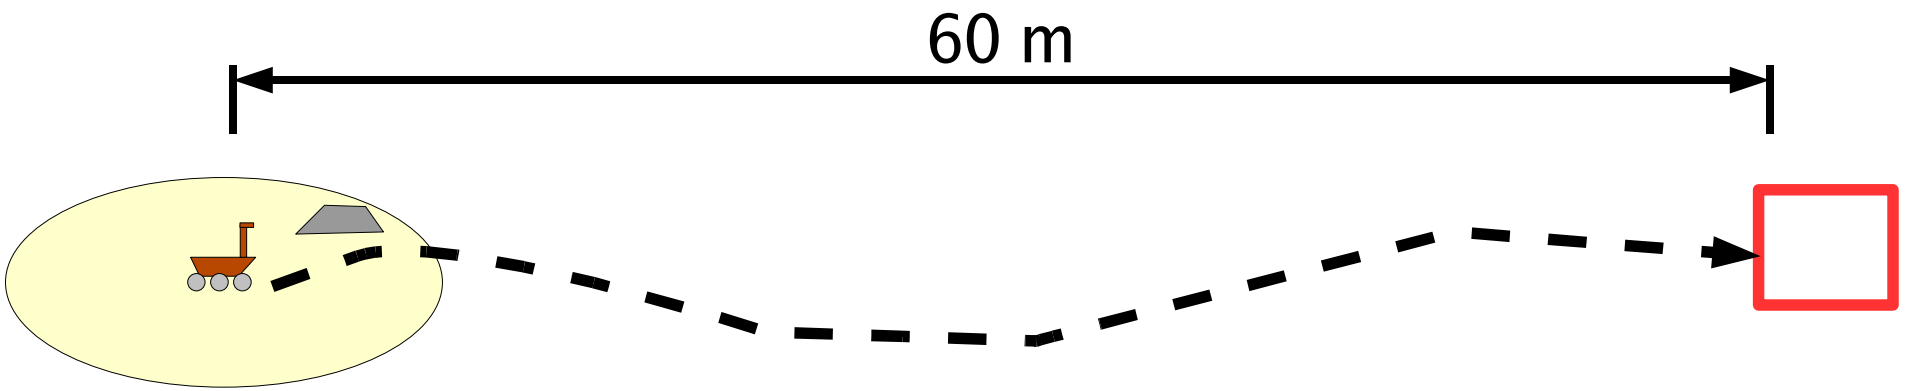

Over-the-Horizon Science

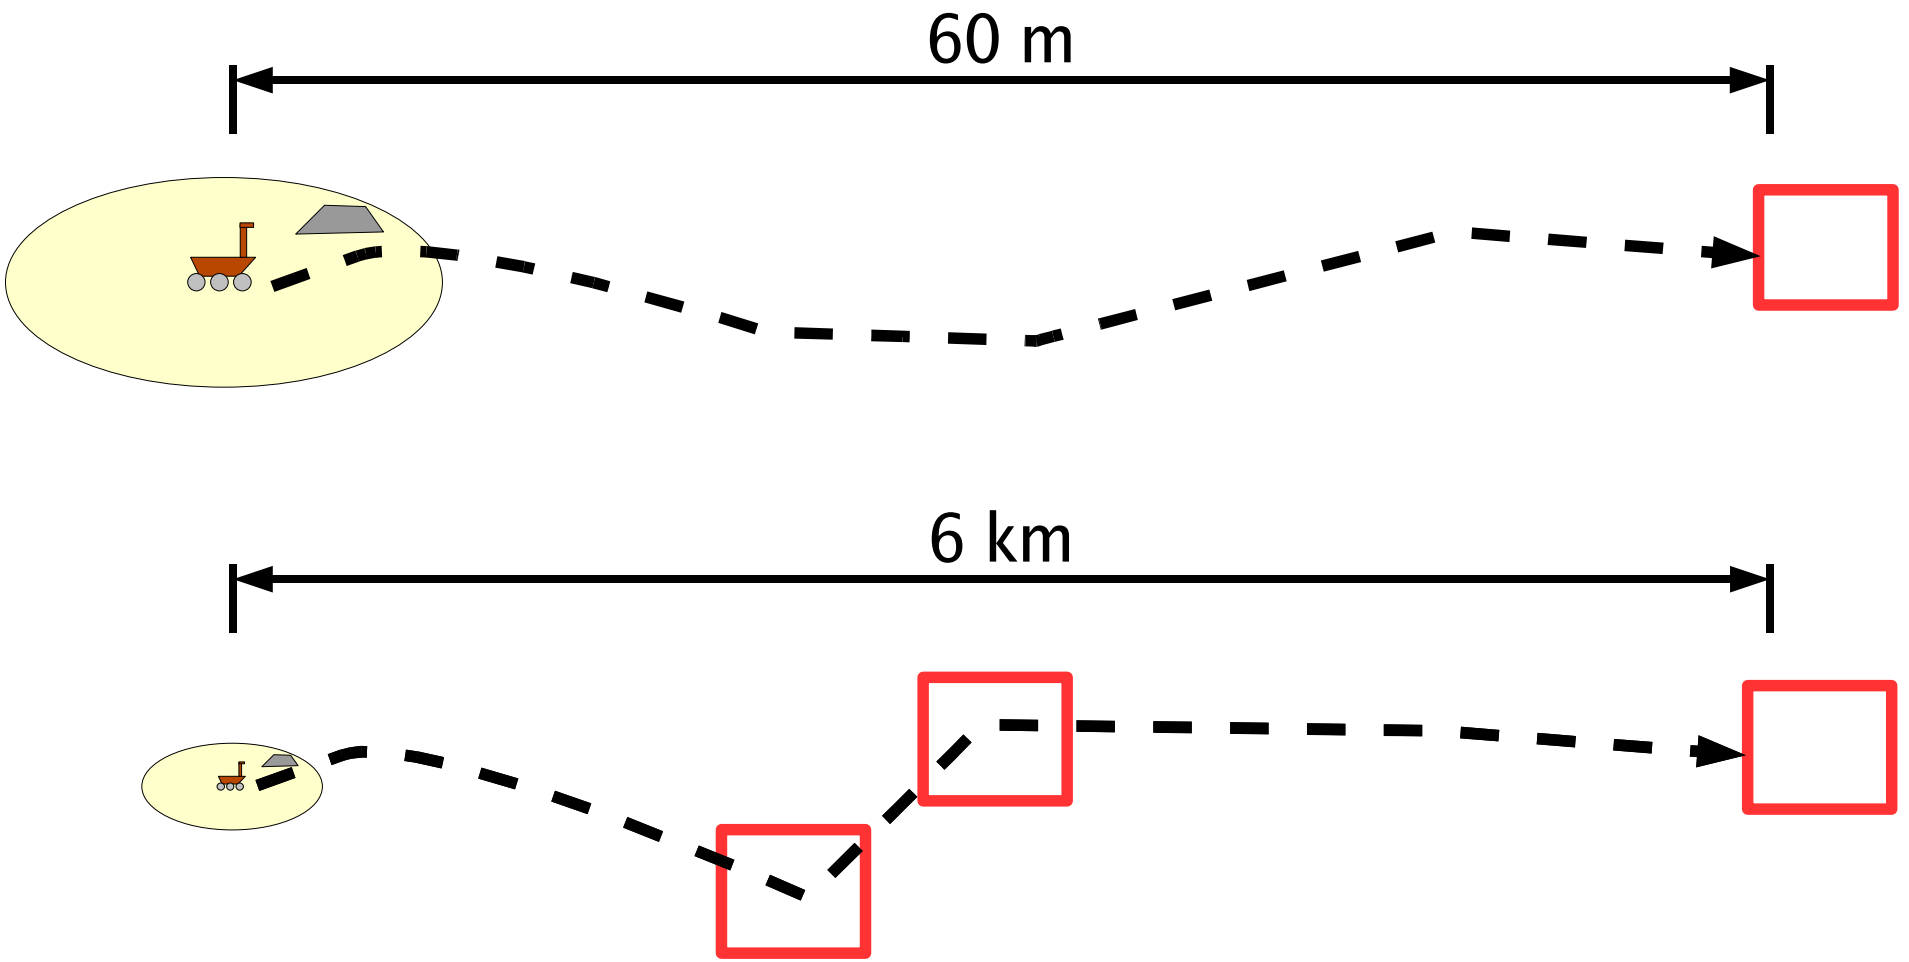

Science Autonomy

- ***The opportunity:*** Long-distance mobility gives the rover access to new sites
- ***The challenge:*** Do useful science in over-the-horizon situations
- ***The method:*** Enable the rover to reason about **science goals** and the **science data** it collects

New Operational Modes

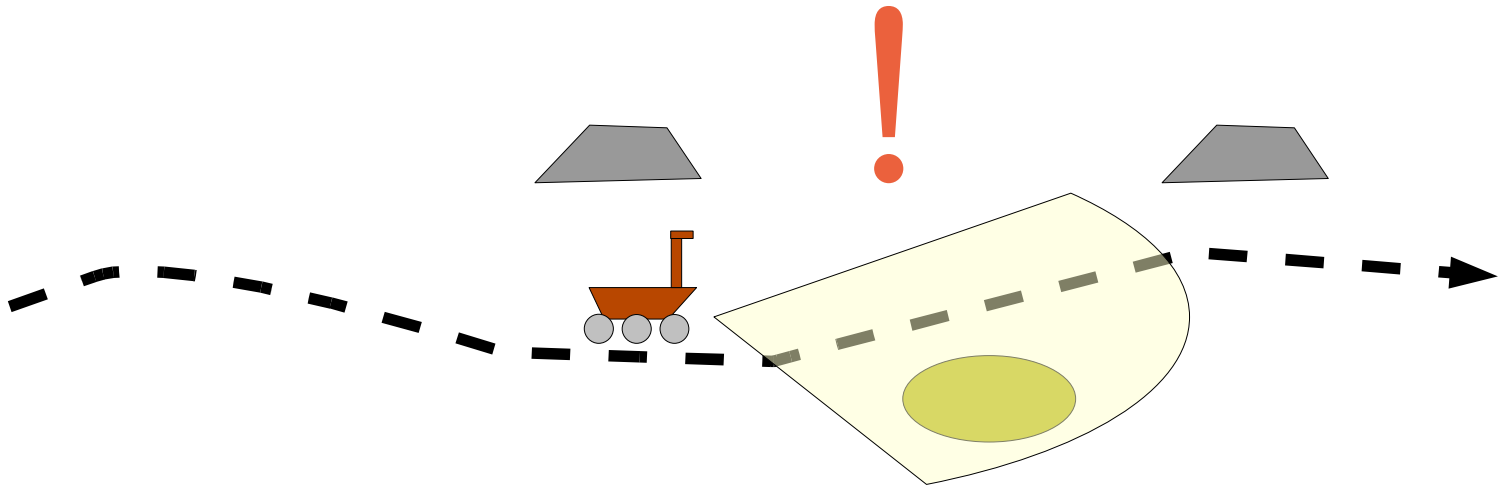

Sensing Uncertainty

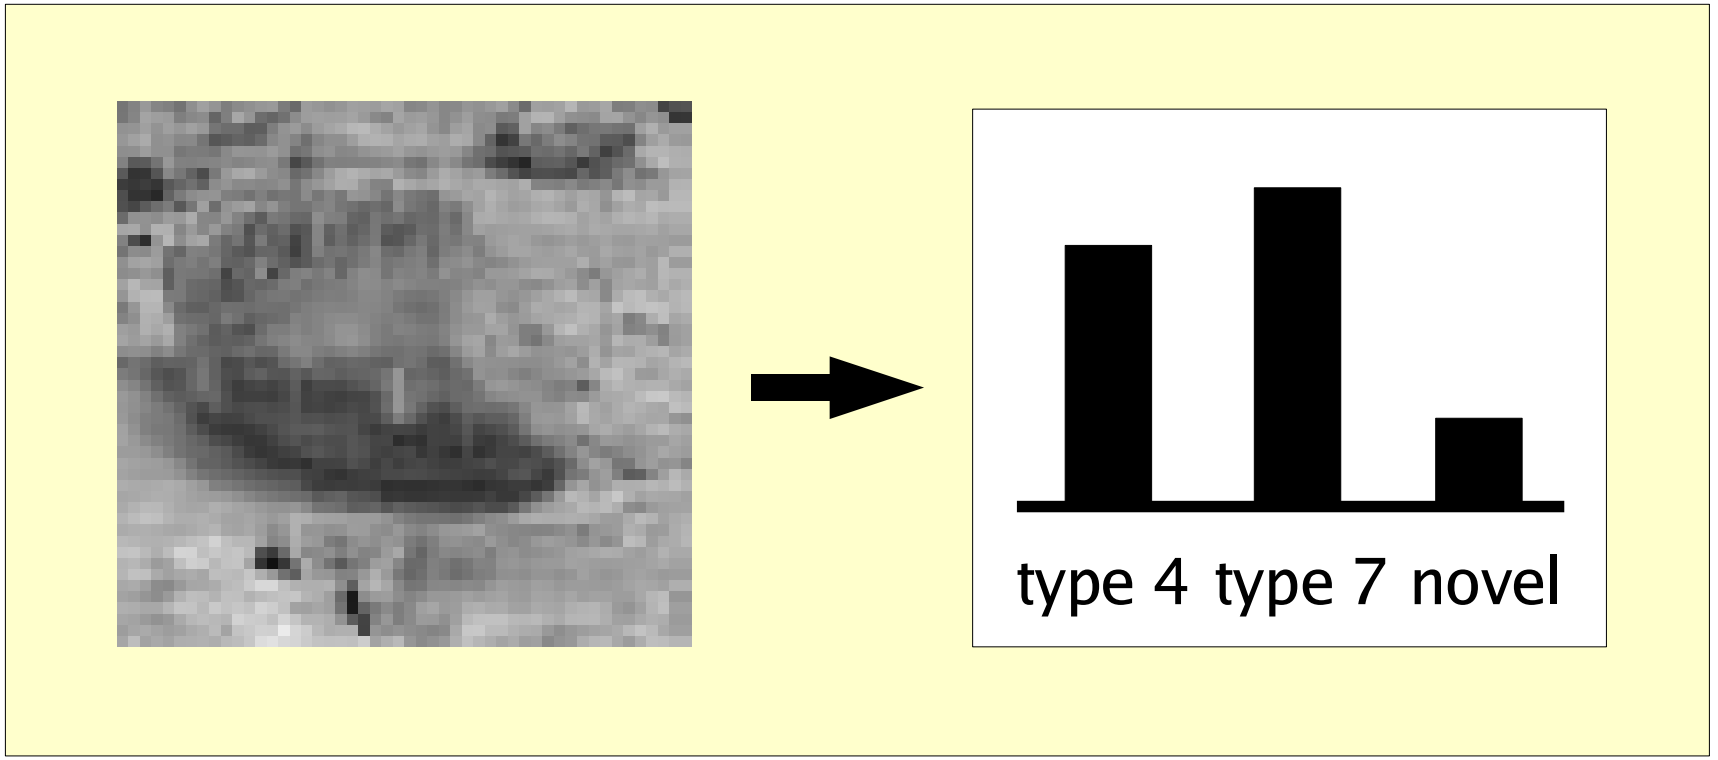

Probabilistic Sensor Planning

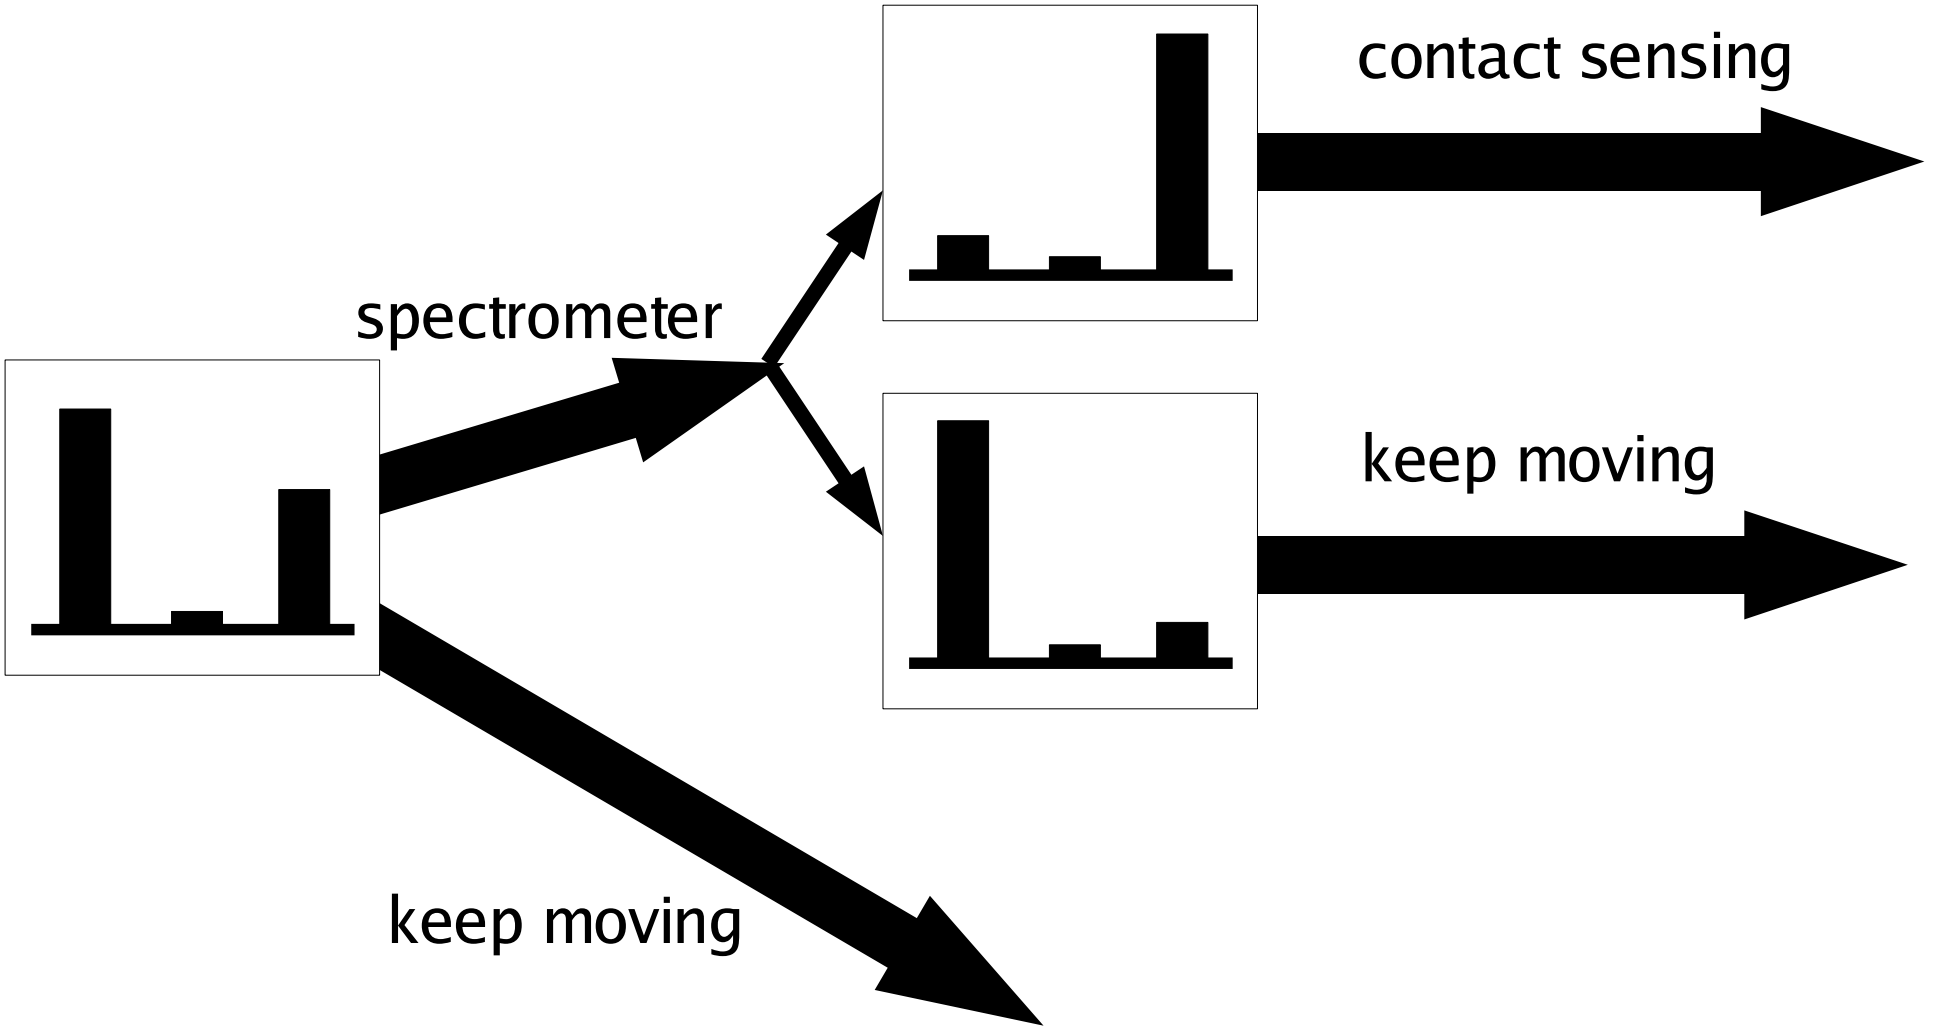

Probabilistic Sensor Planning

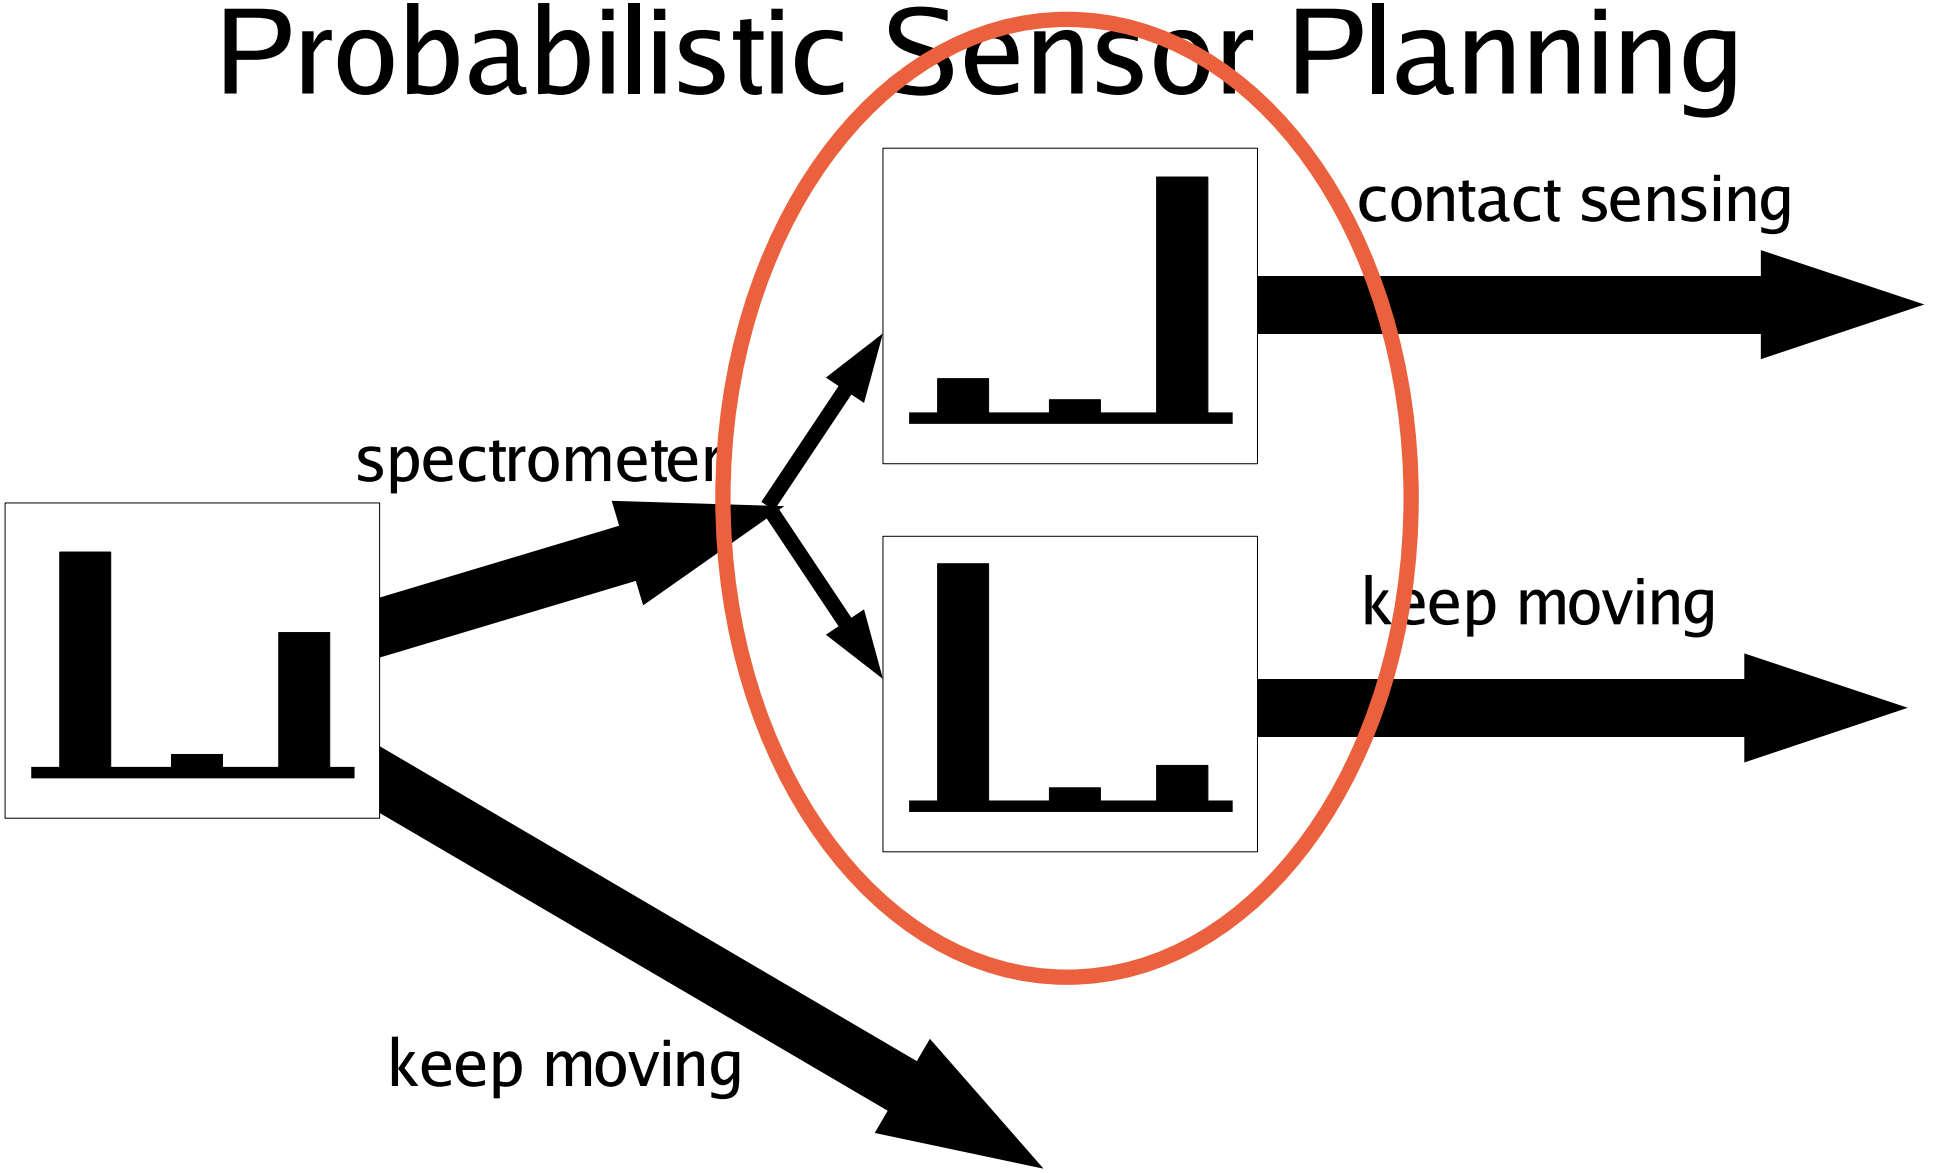

POMDP Model

- Principled method for reasoning about information gain and future beliefs

POMDP Model

- Principled method for reasoning about information gain and future beliefs
- **Horribly, horribly intractable**

POMDP Model

- Principled method for reasoning about information gain and future beliefs
- **Horribly, horribly intractable**
- Develop approximation techniques that take advantage of structure in the SA domain to speed up planning

Thesis Statement

- Enabling a robot to reason about science goals and the science data it collects will significantly improve its exploration efficiency,

Thesis Statement

- Enabling a robot to reason about science goals and the science data it collects will significantly improve its exploration efficiency, **and POMDP methods can significantly improve the quality of science autonomy plans.**

Outline

- Motivation
- ***Background and related work***
- Technical approach
- Preliminary work
- Proposed research

Science Autonomy Related Work

- Image and spectral analysis
  [Gulick et al., 2001] [Gazis and Roush, 2001]
- Balancing exploration priorities
  2001]
- Sensor planning to detect meteorites
  et al., 2001]
- Orbiter reacting to science events
  et al., 2003]
- OASIS project: closed-loop SA system
  et al., 2003]

[Wagner

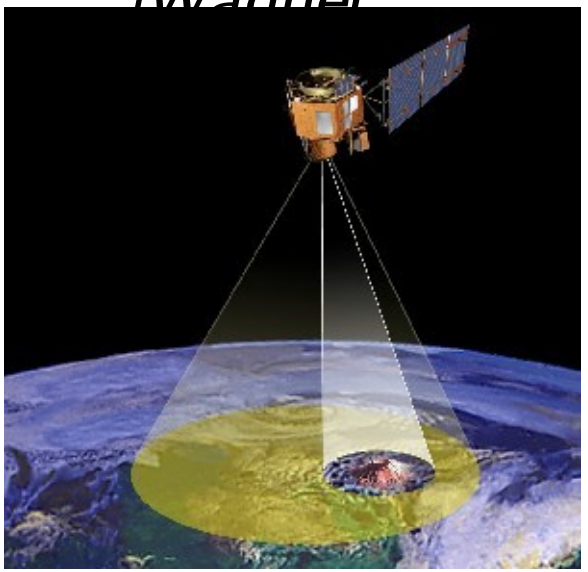

POMDP Introduction

- Recall:
  - Great for sensor planning
  - Tend to be massively intractable

Example Problem: *RockSample*

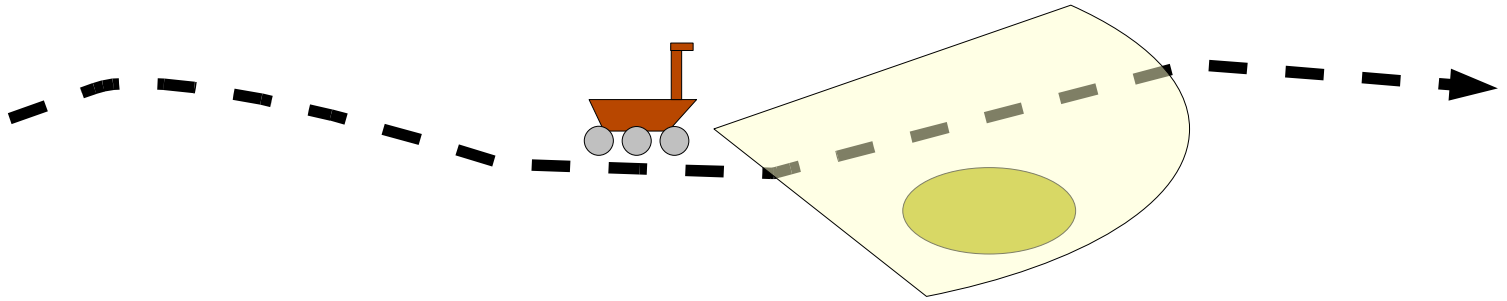

Example Problem: *RockSample*

sample

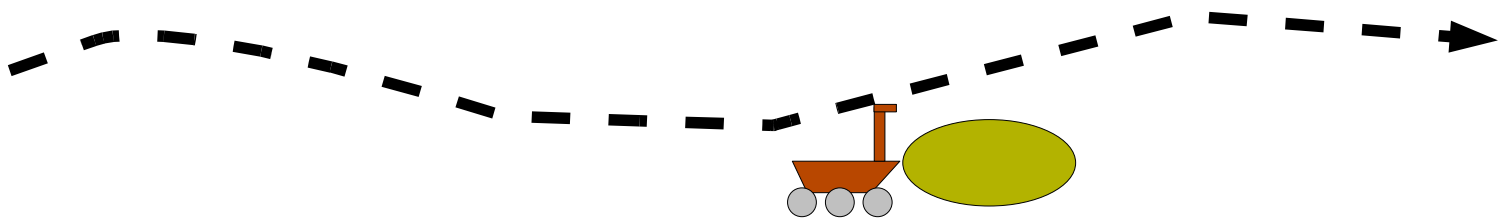

Example Problem: *RockSample*

sample
exit

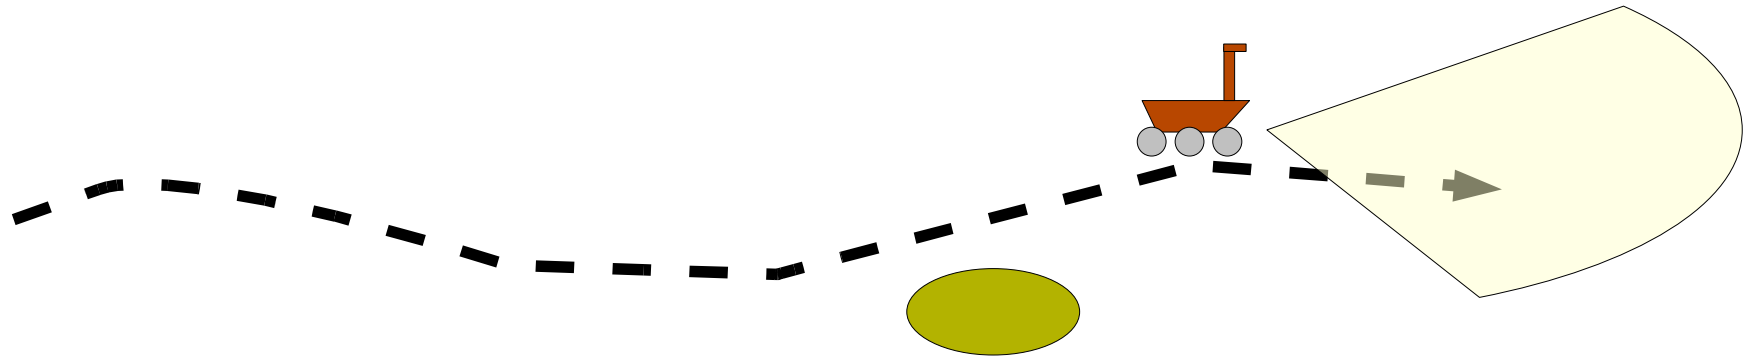

Example Problem: *RockSample*

sample
exit
check

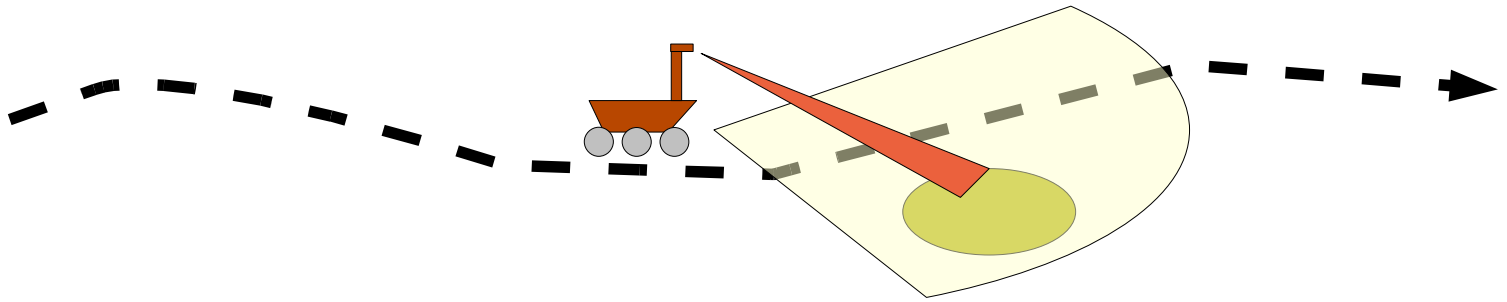

Example Problem: *RockSample*

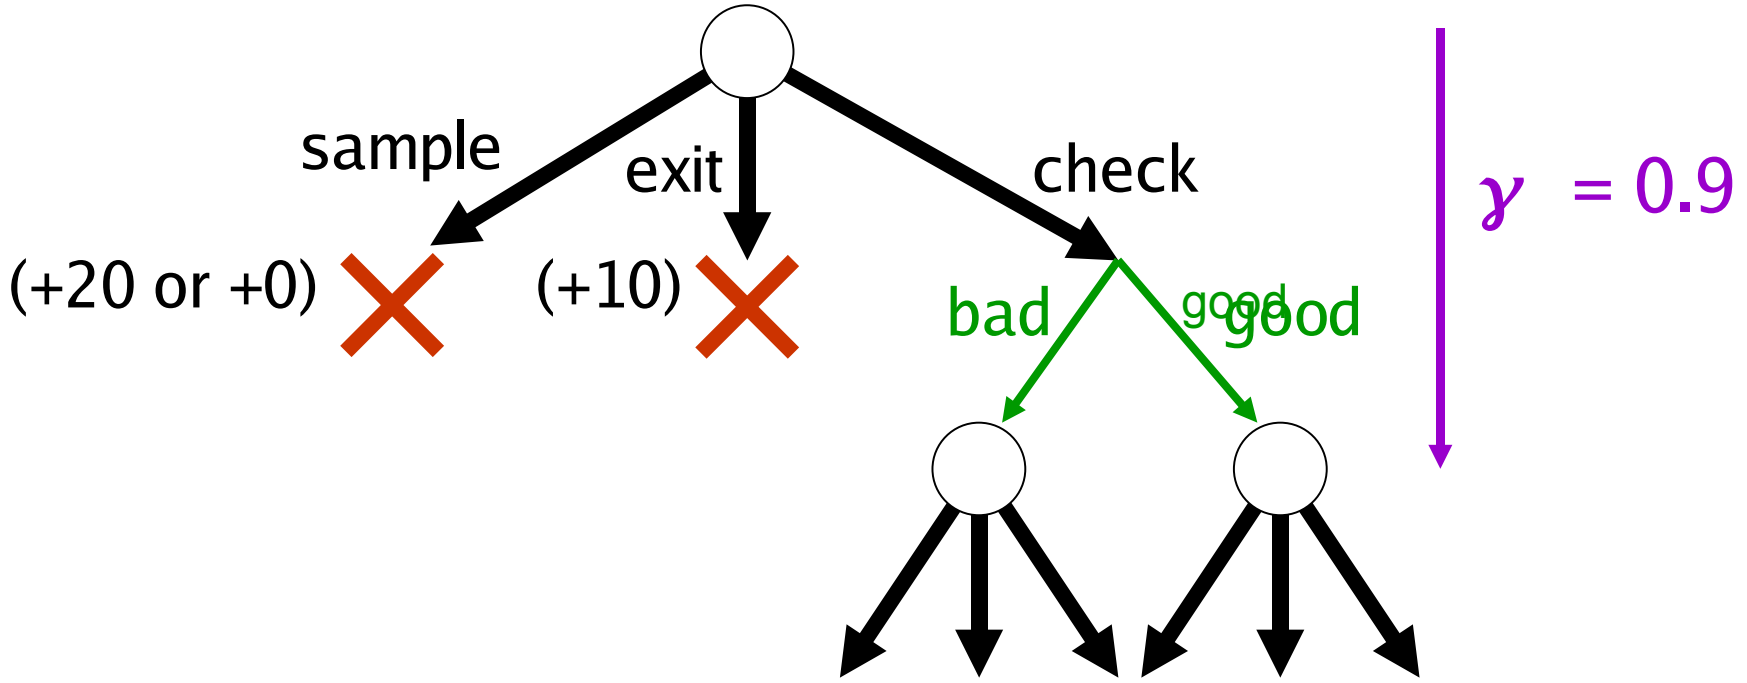

History and Belief

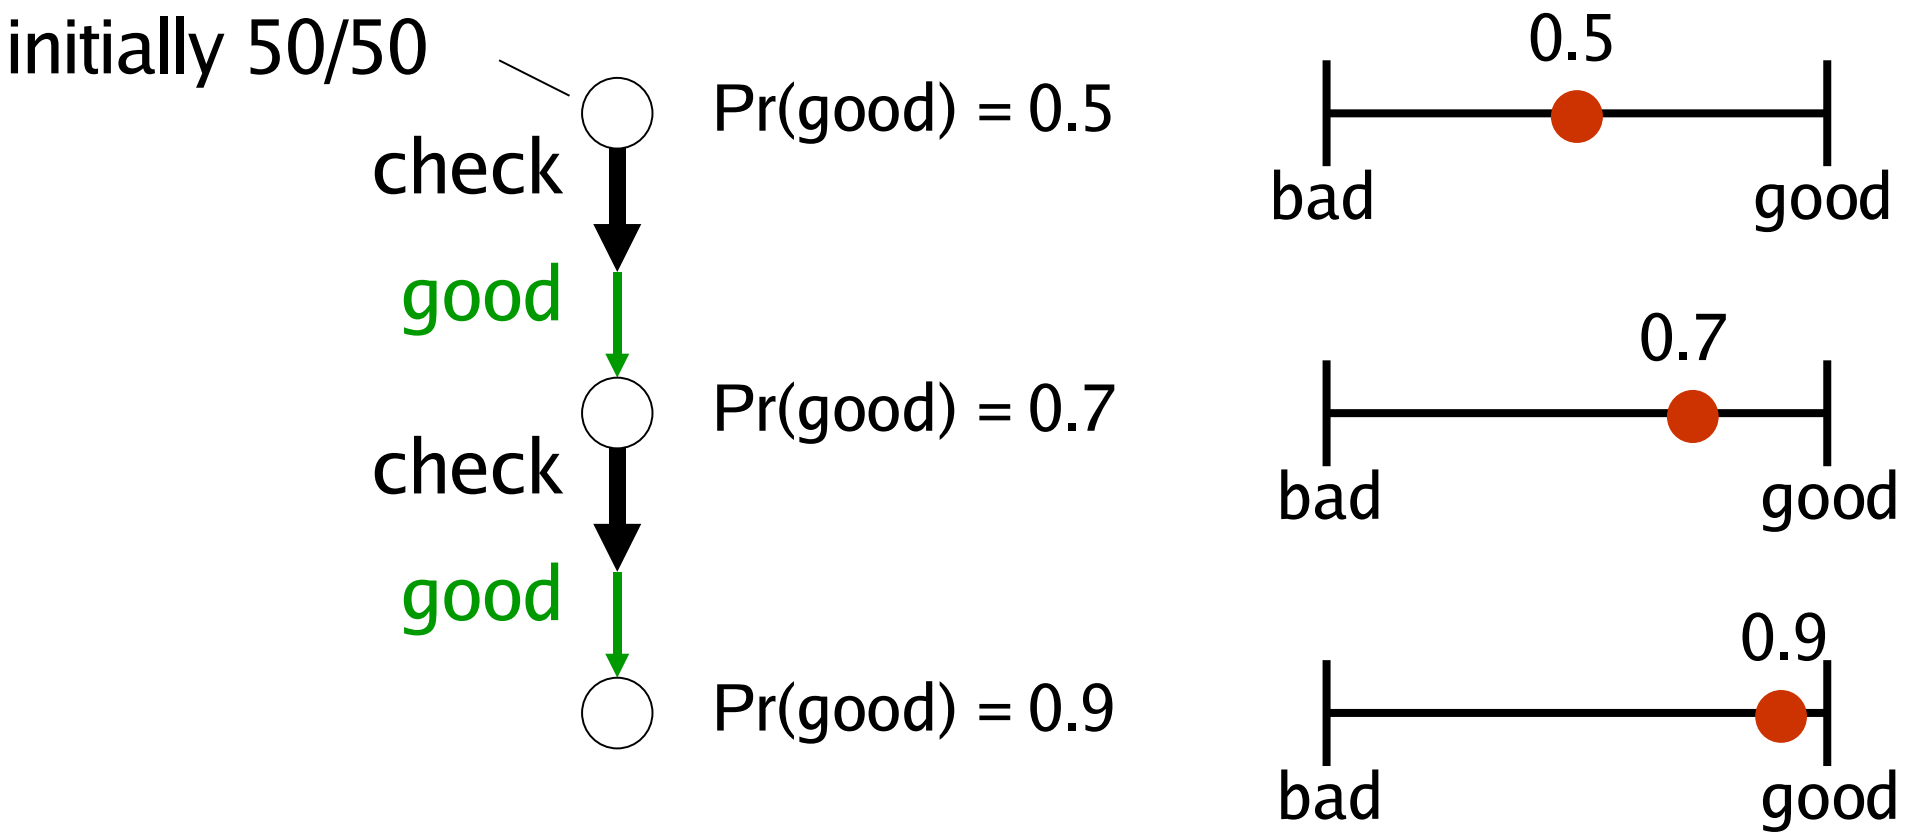

Policy

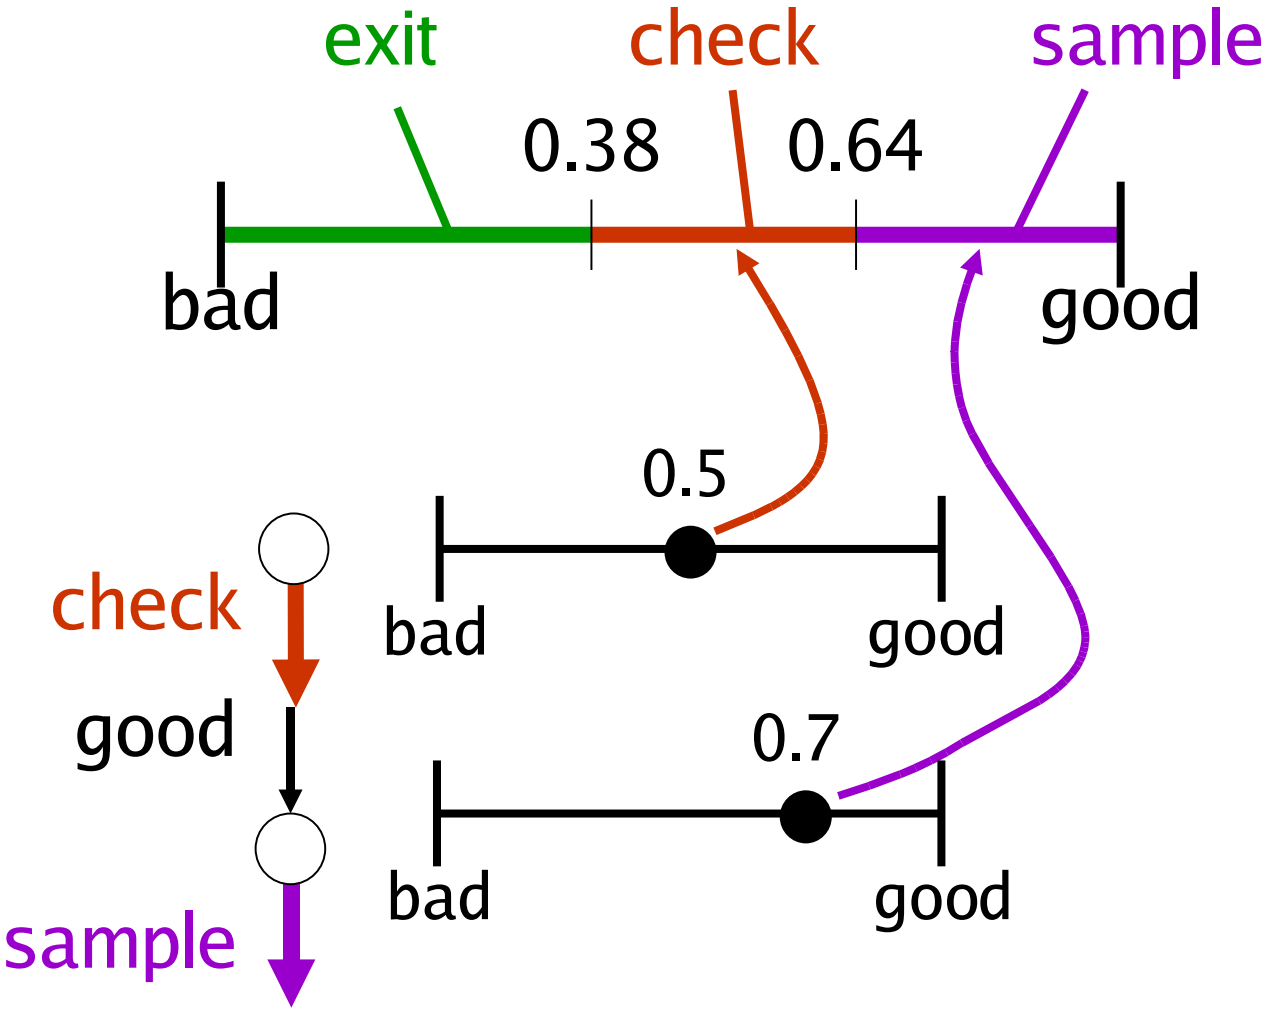
exit
check
sample
0.38
0.64
bad
good
0.5
check
bad
good
good
0.7
bad
good
sample

Value Function: Horizon 1

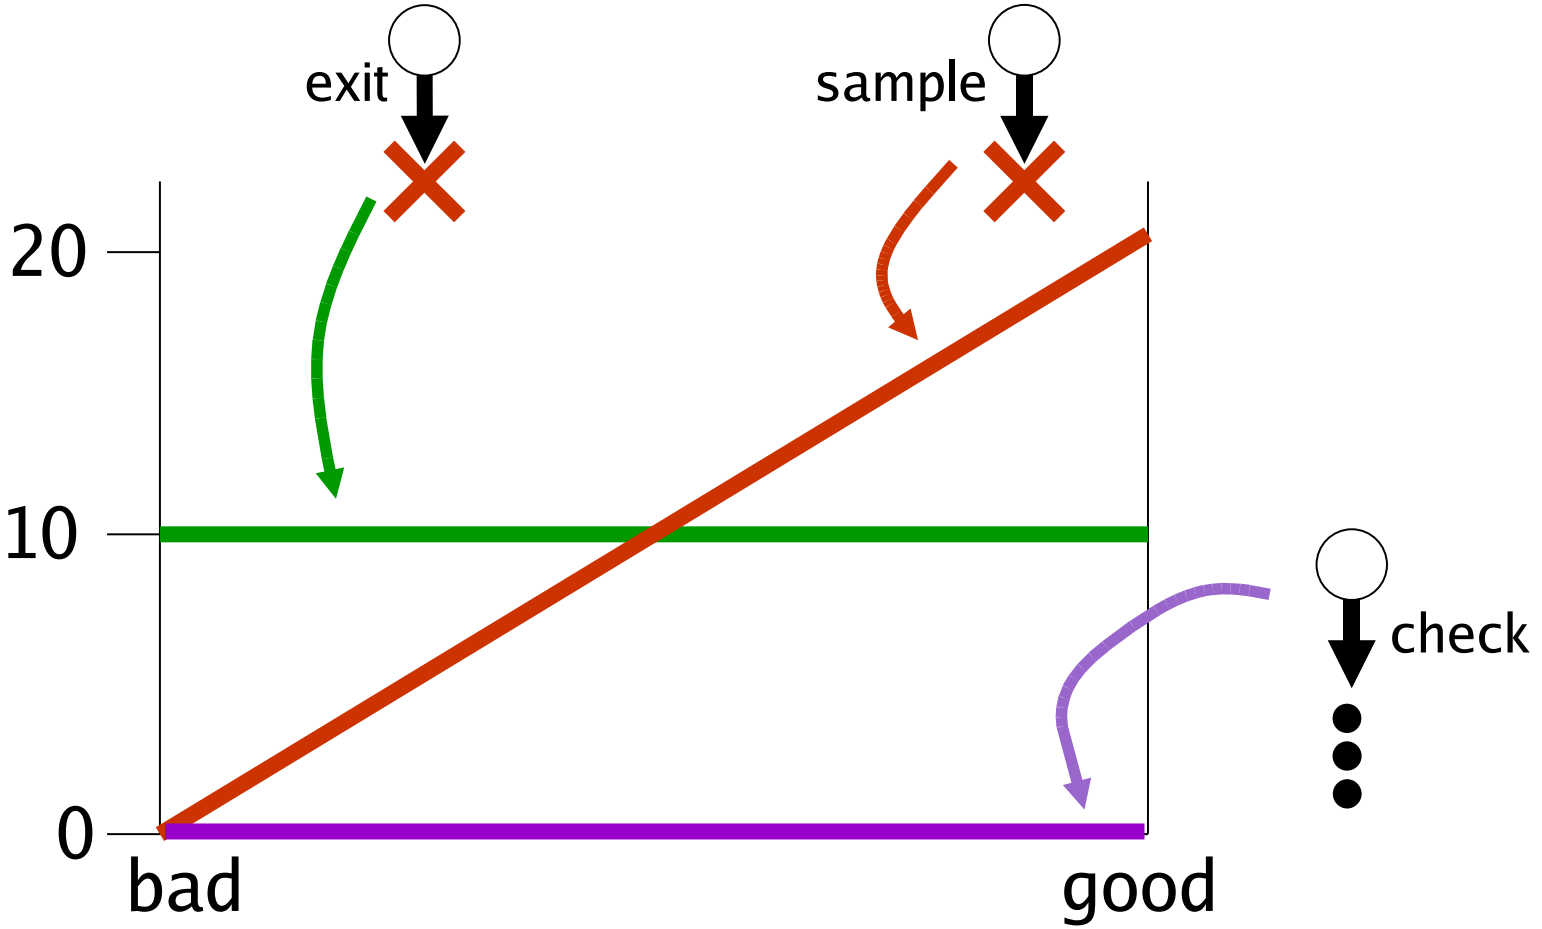

Value Function: Horizon 1

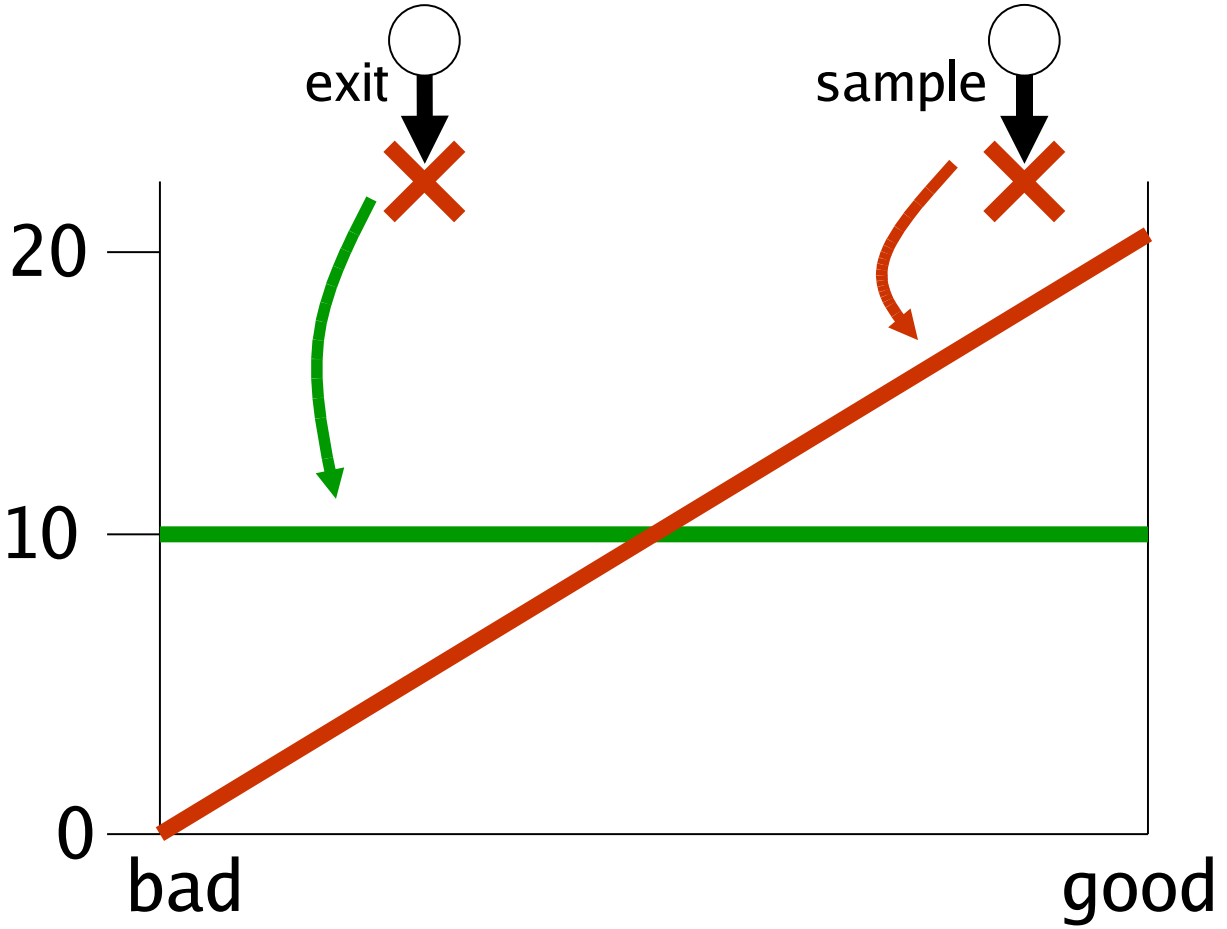

Value Function: Horizon 2

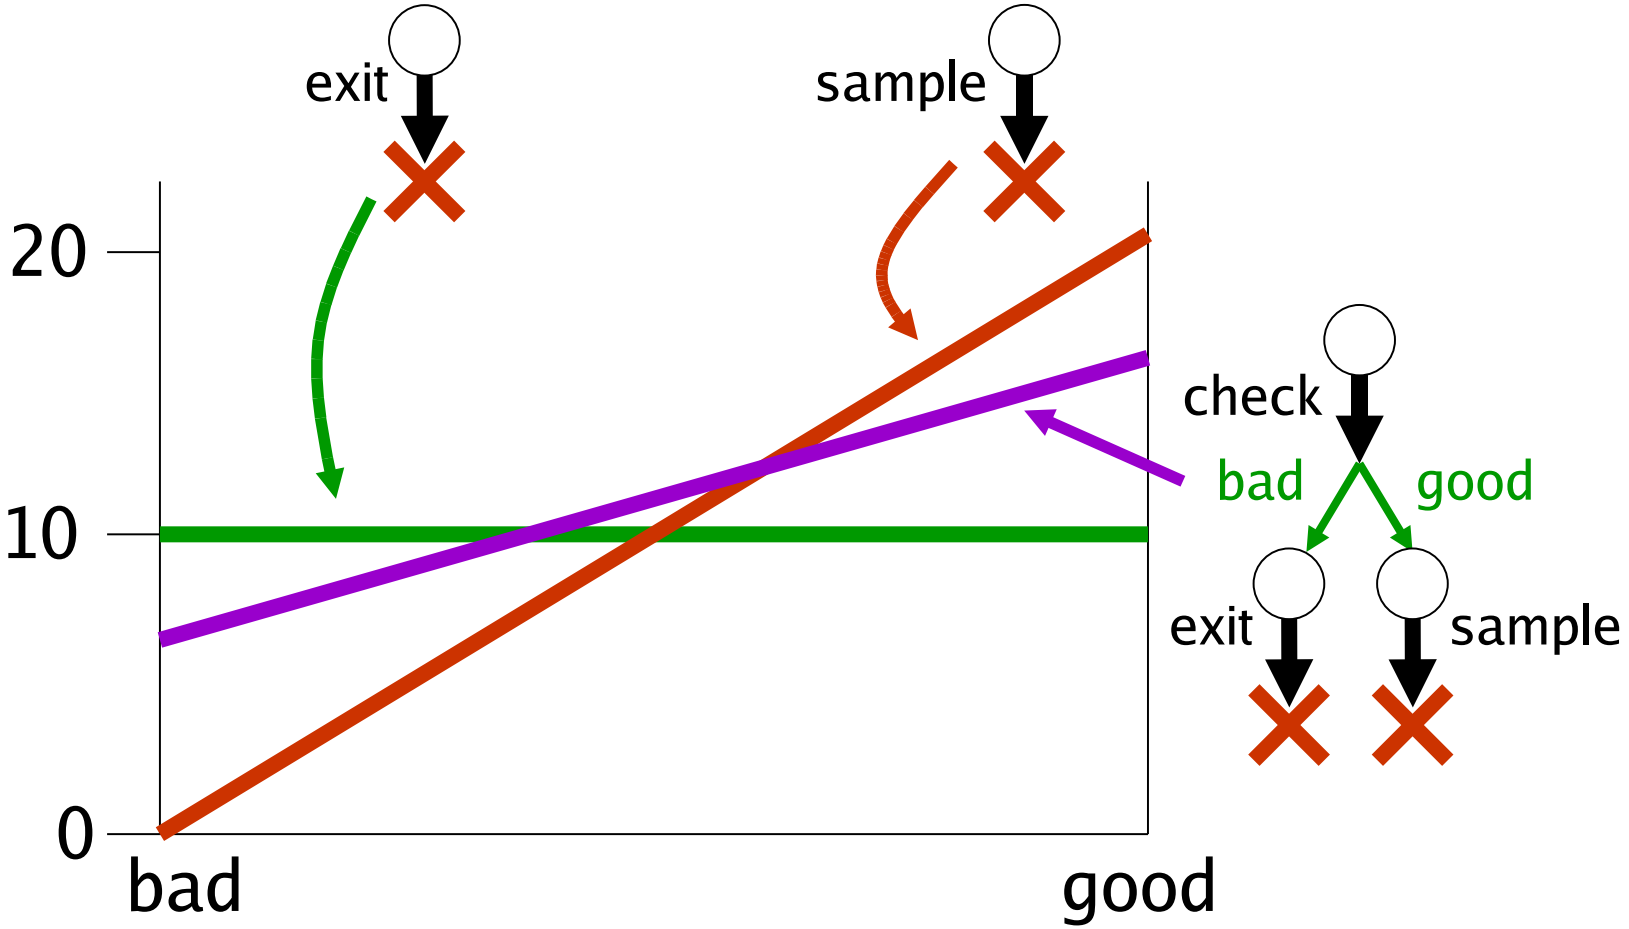

Computing the Policy

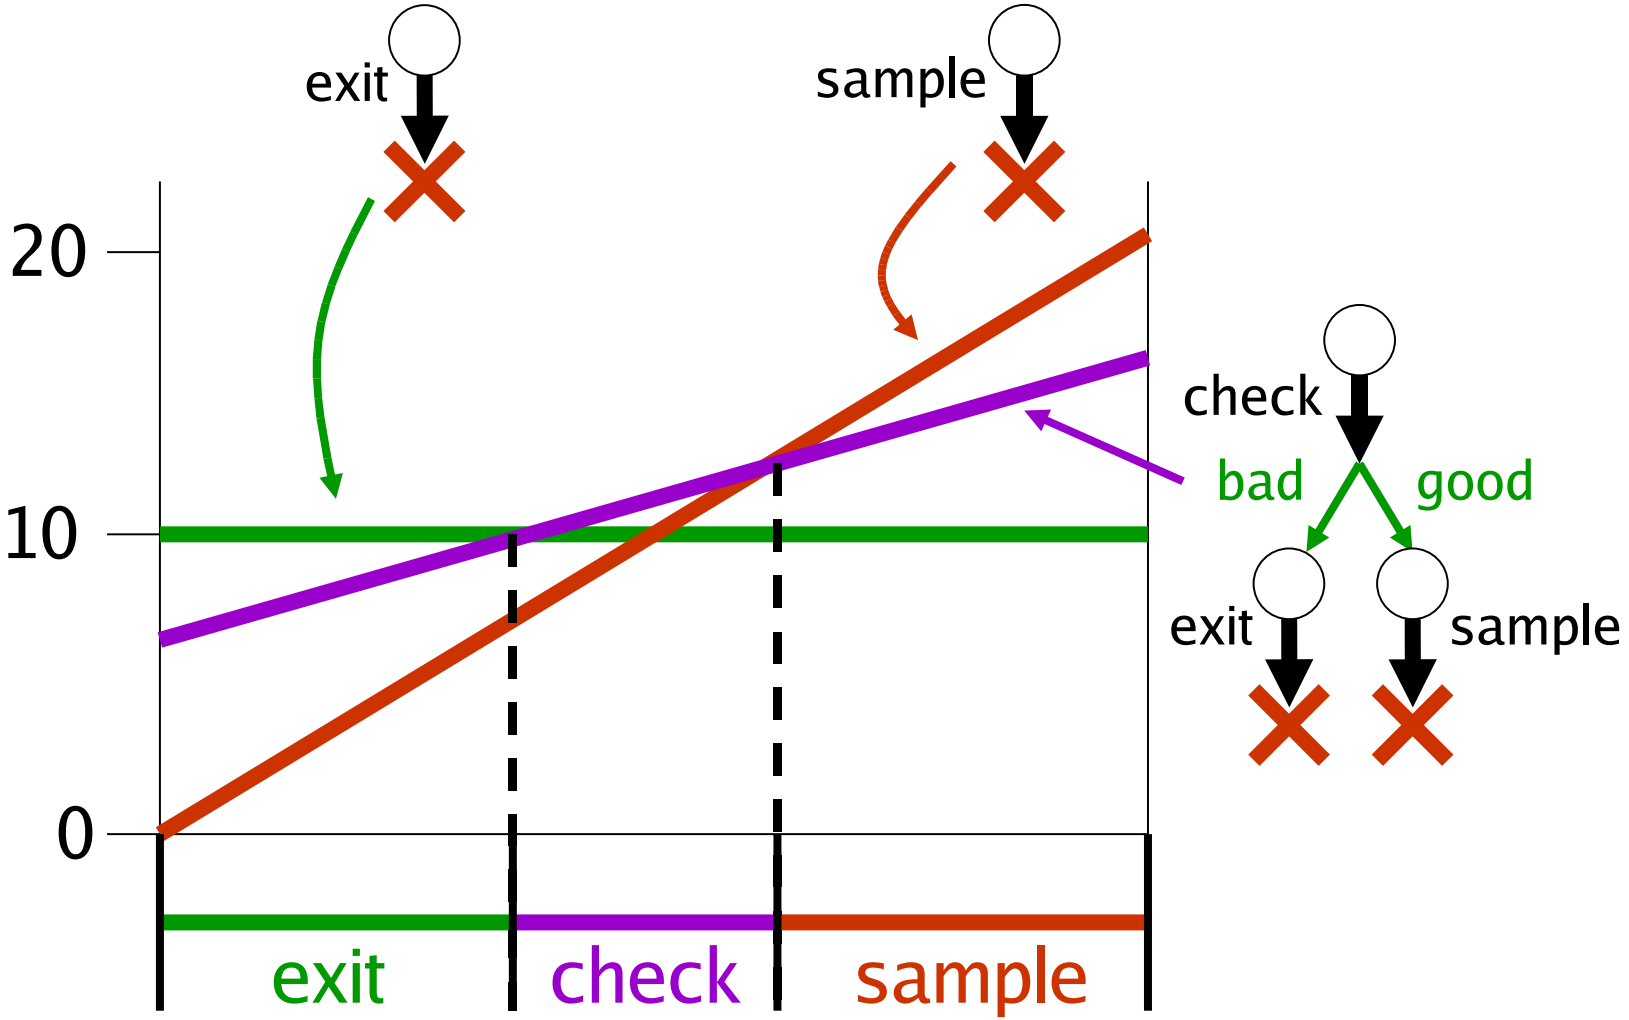

Value Iteration

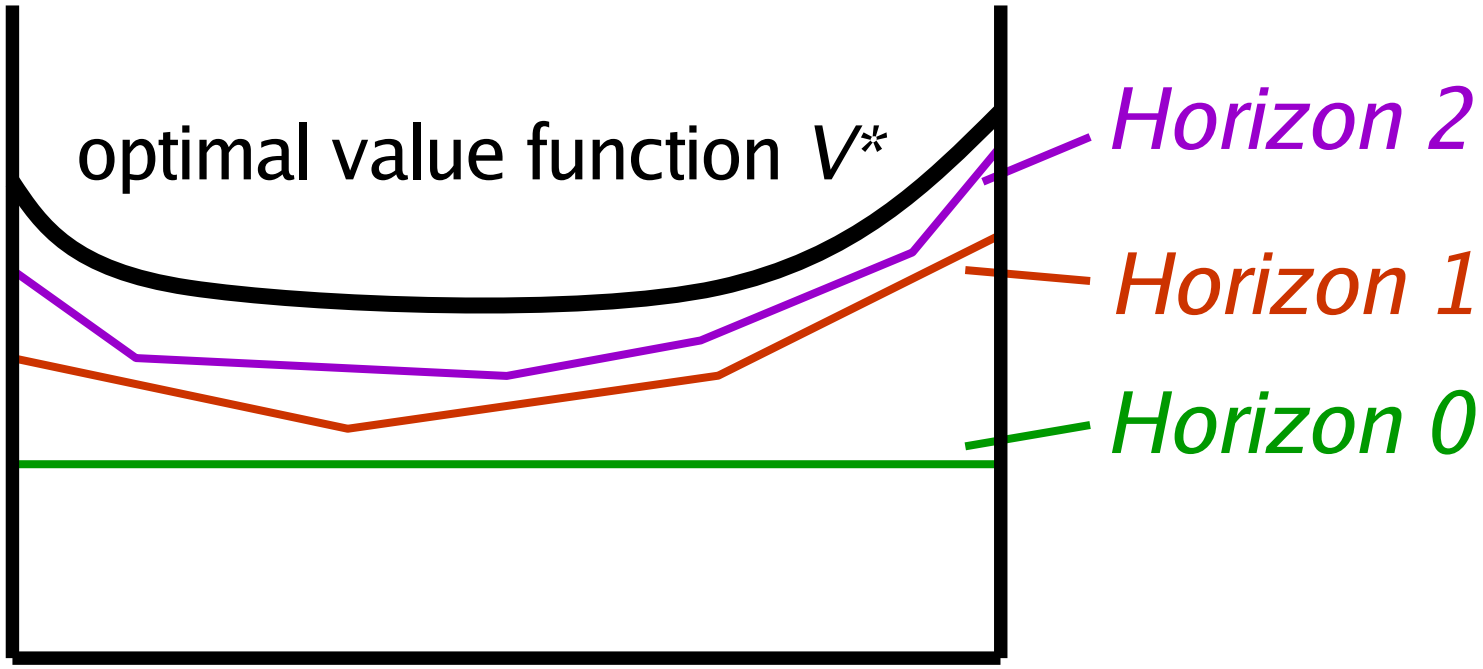

POMDP Related Work

- **Policy iteration**
  [Hansen, 1998] [Poupart and Boutilier, 2003]
- **Gradient ascent policy search**
  [Baxter and Bartlett, 2000]
- **Value iteration** *[Sondik, 1971]*

Value Iteration Variants

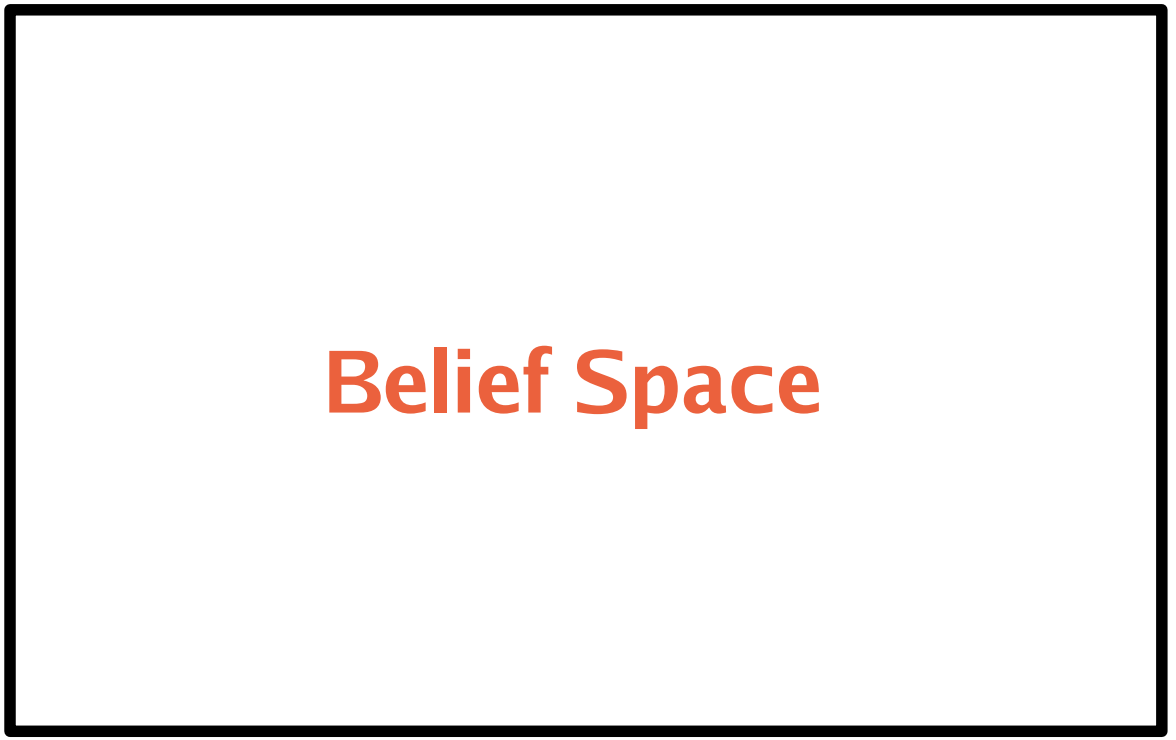

Value Iteration Variants

Exact value iteration

[Littman, 1994] [Cassandra et al., 1997]

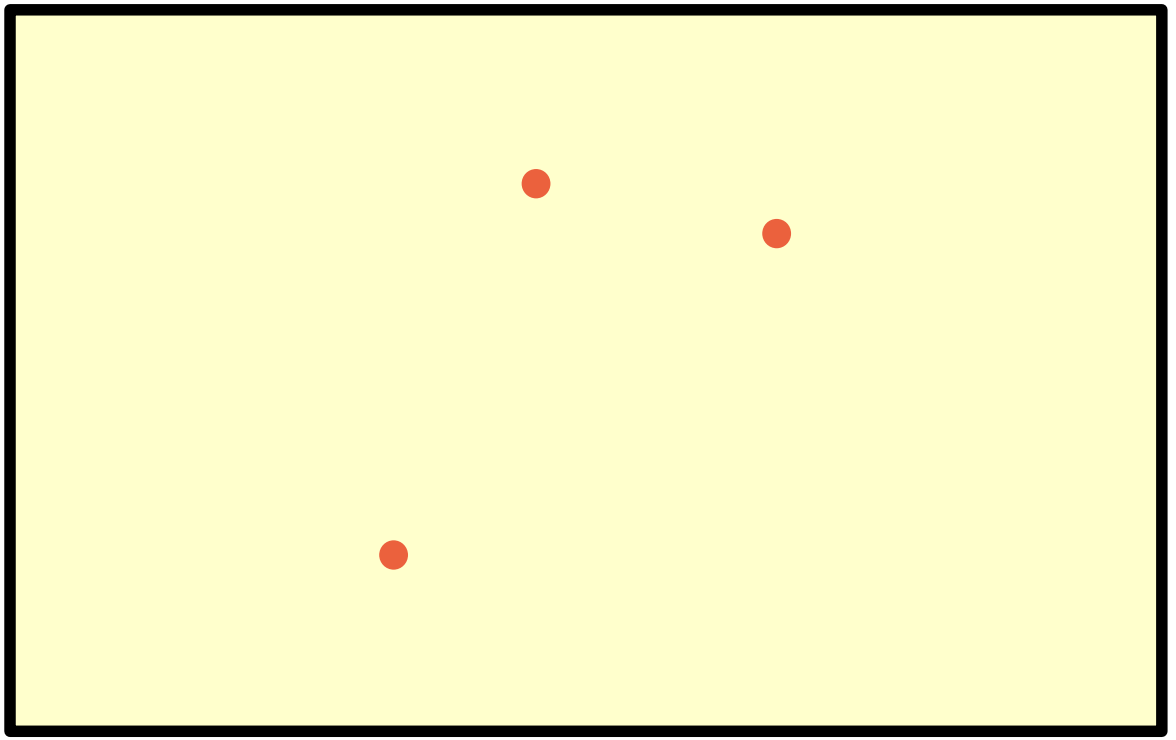

Value Iteration Variants

Exact value iteration

[Littman, 1994] [Cassandra et al., 1997]

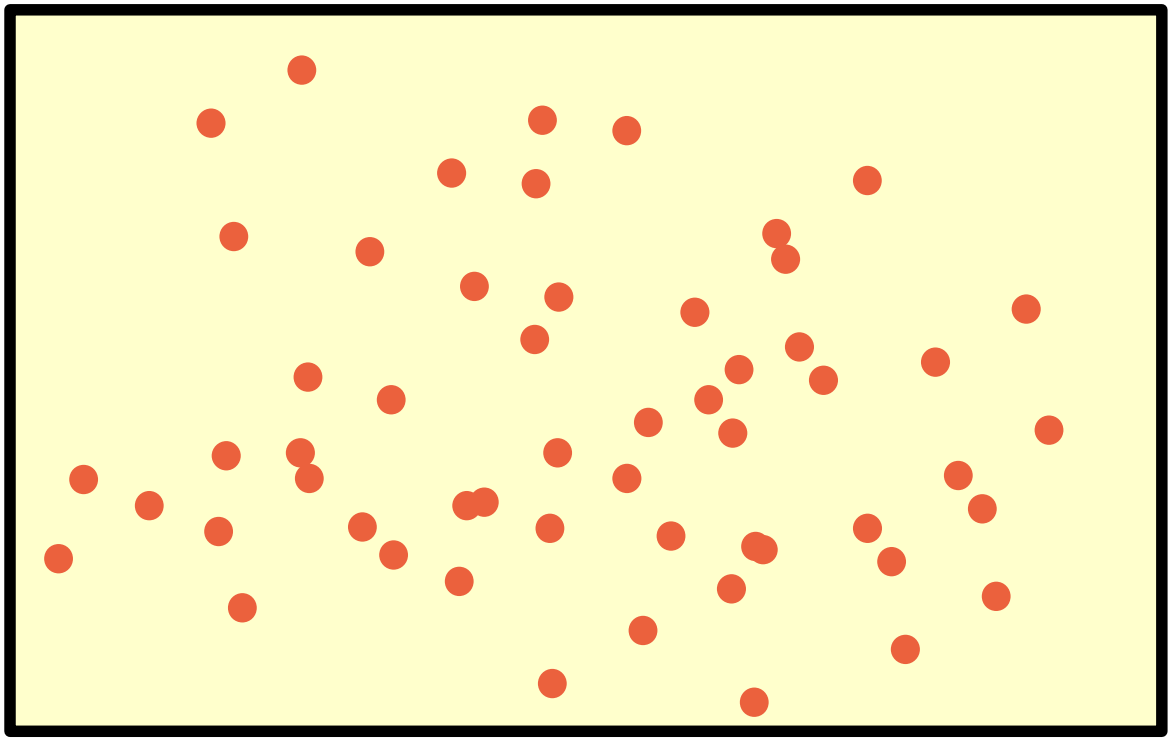

Value Iteration Variants

Value Iteration Variants

Grid-based

[Brafman, 1997] [Hauskrecht, 2000]

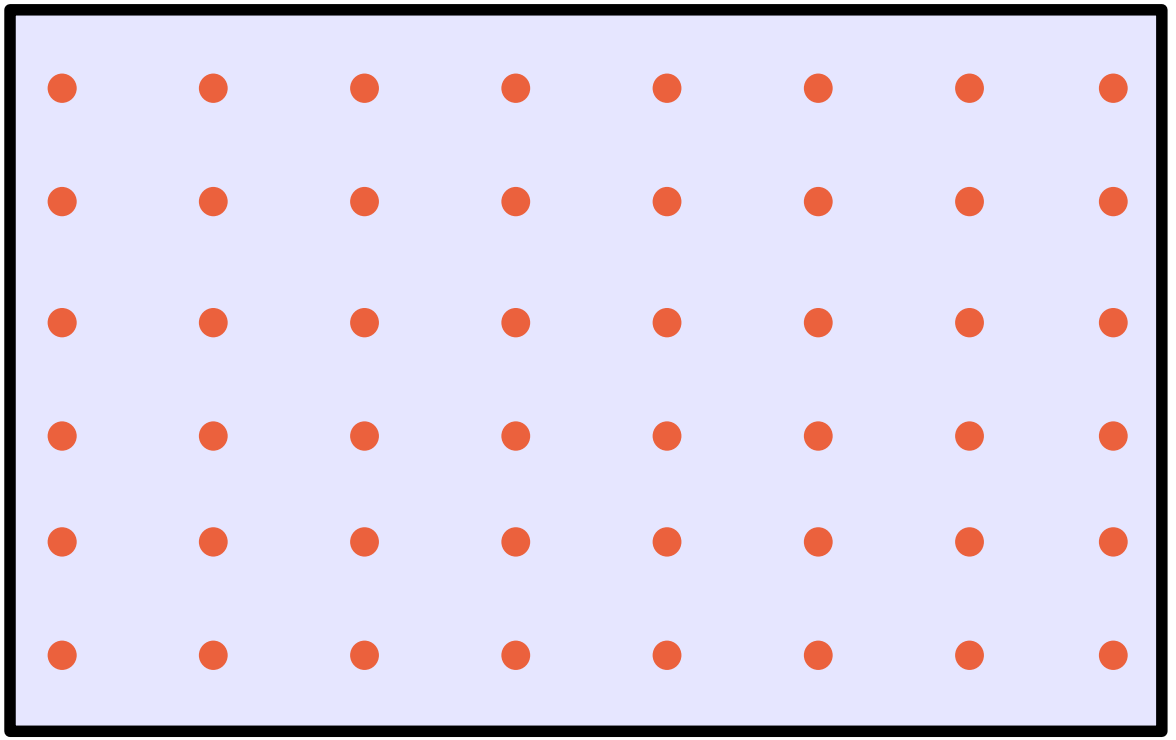

Value Iteration Variants

What next?

Value Iteration Variants

Focused value iteration

Value Iteration Variants

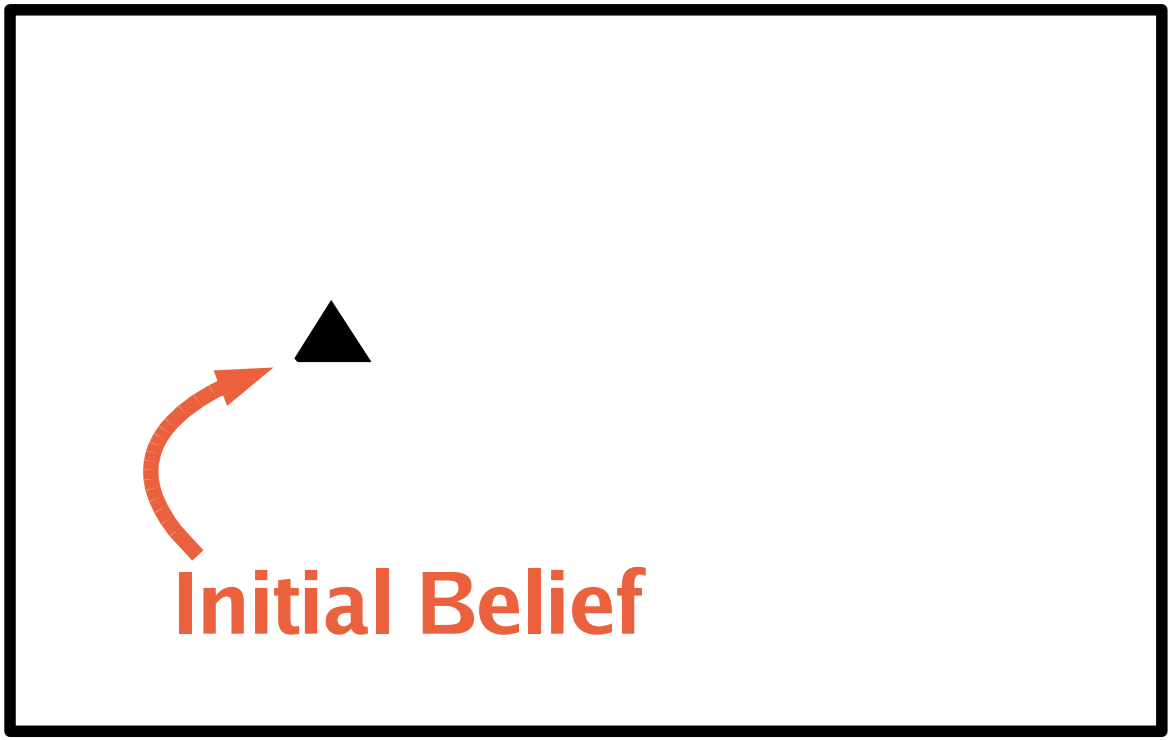

Value Iteration Variants

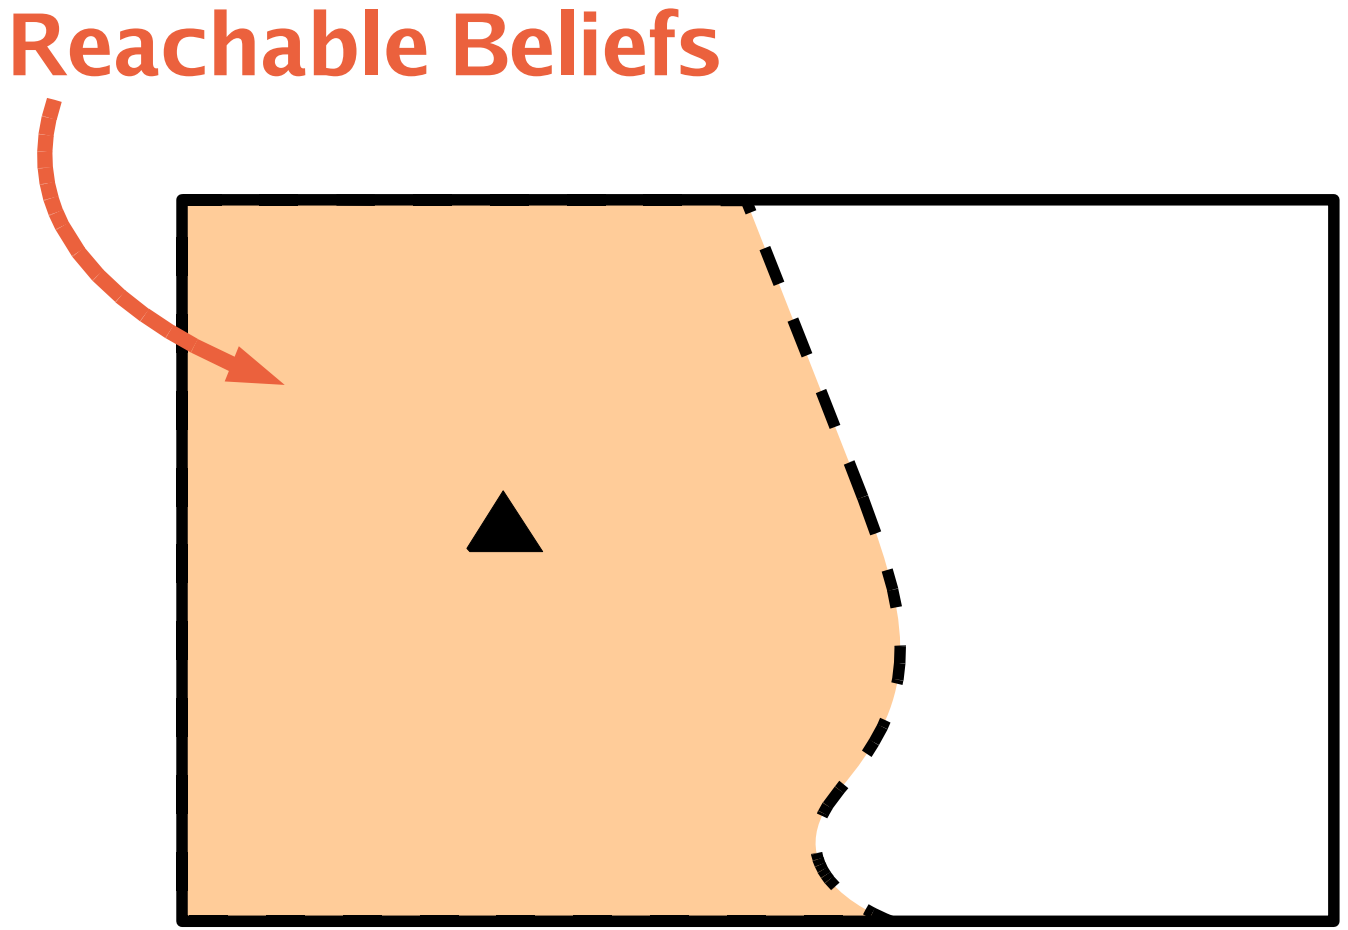

Value Iteration Variants

Relevant Beliefs

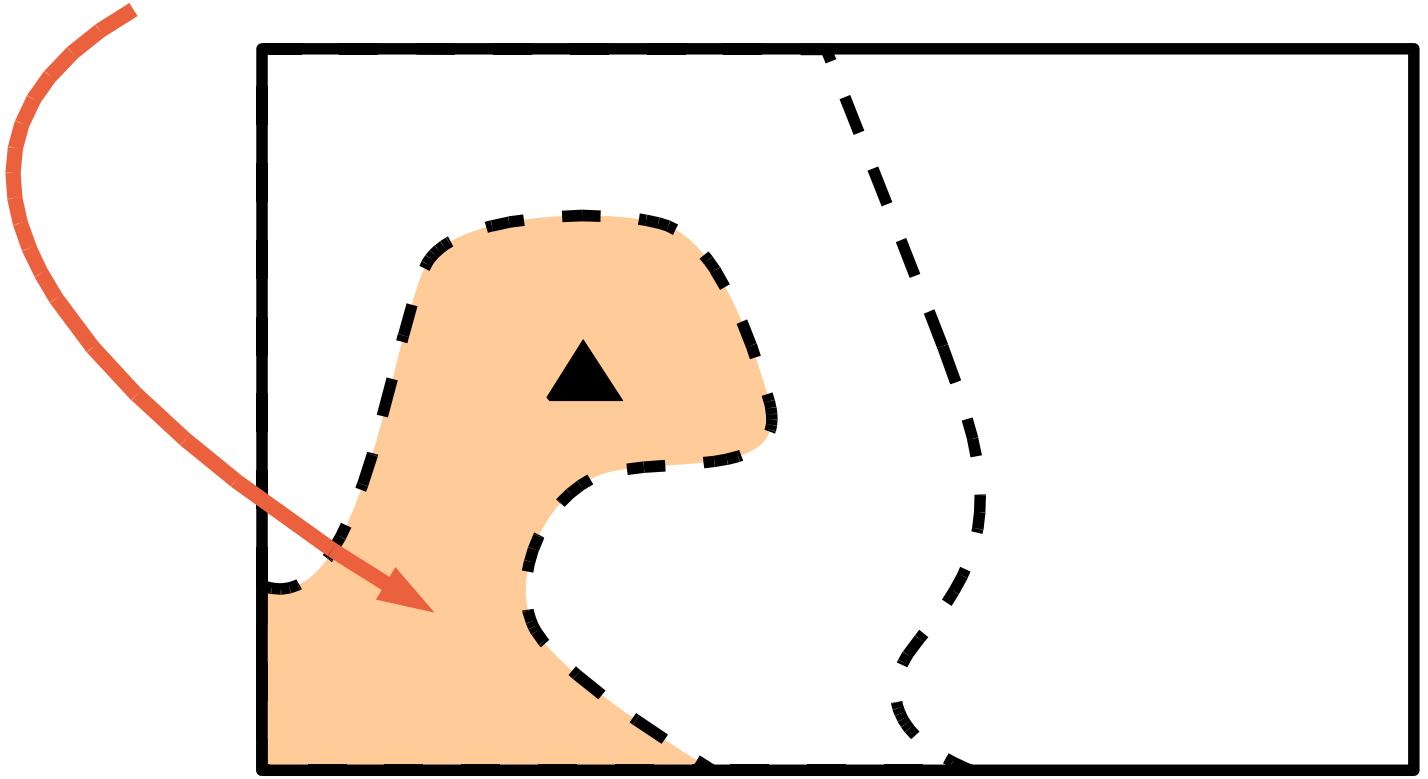

Value Iteration Variants

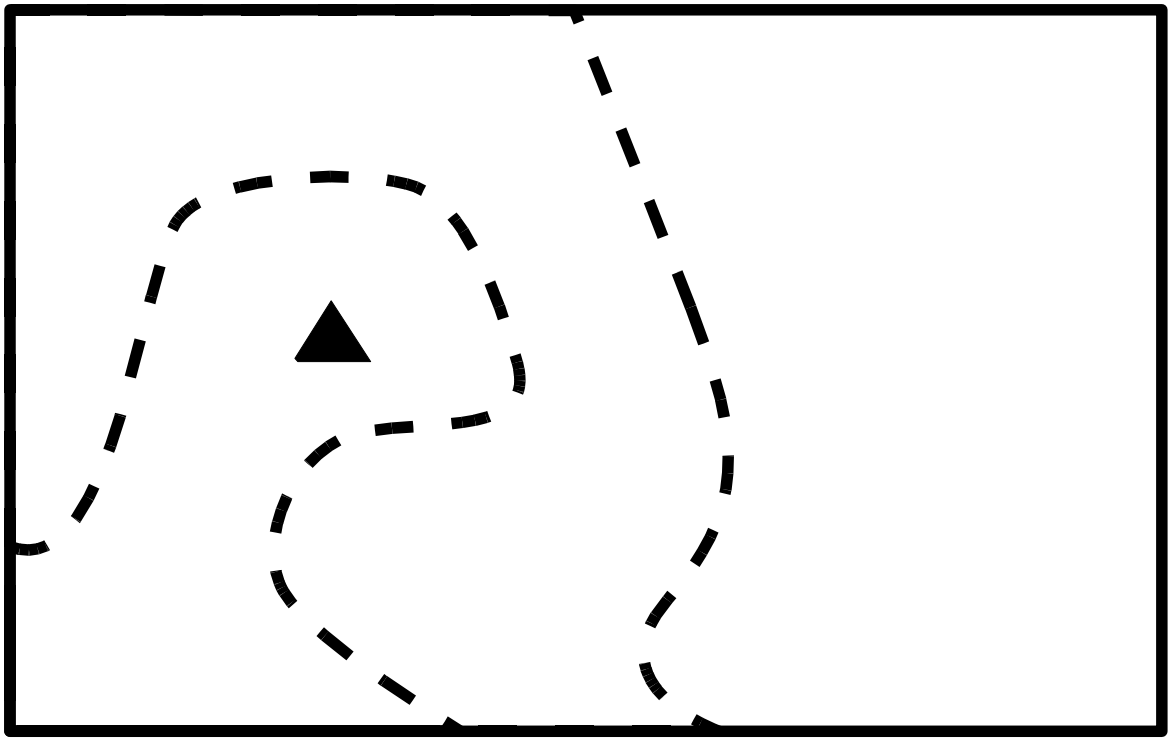

Value Iteration Variants

Point-based (PBVI)

[Pineau et al., 2003]

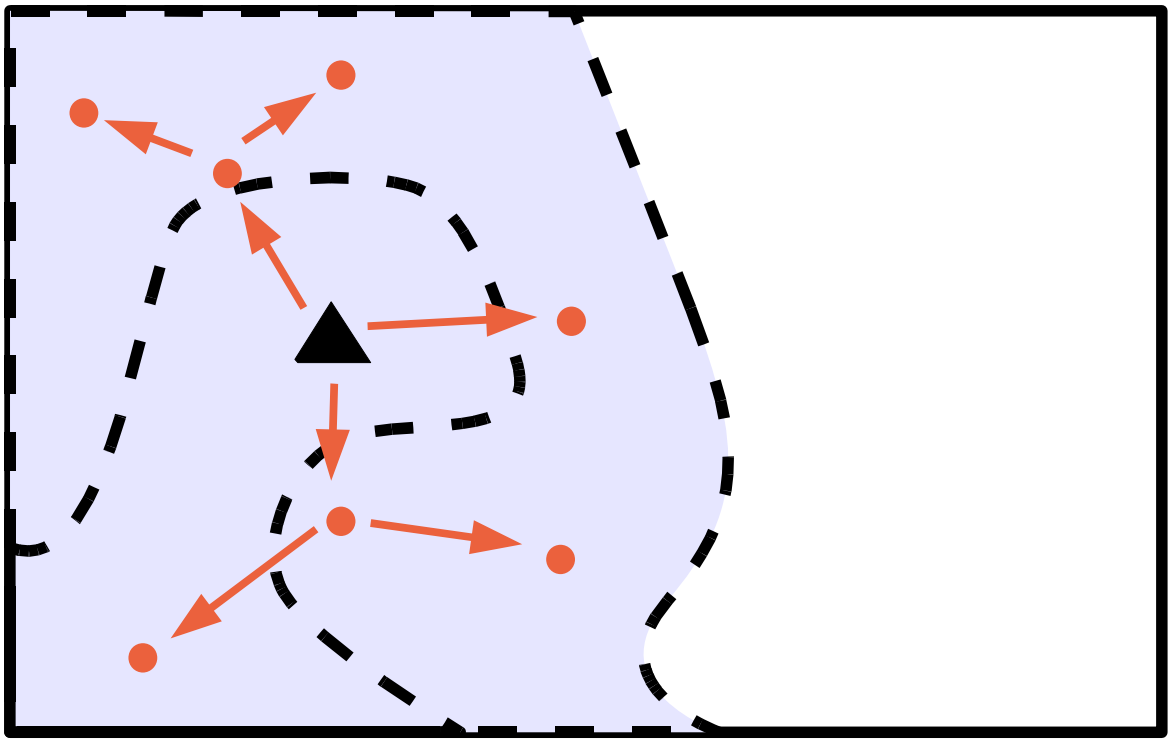

Value Iteration Variants

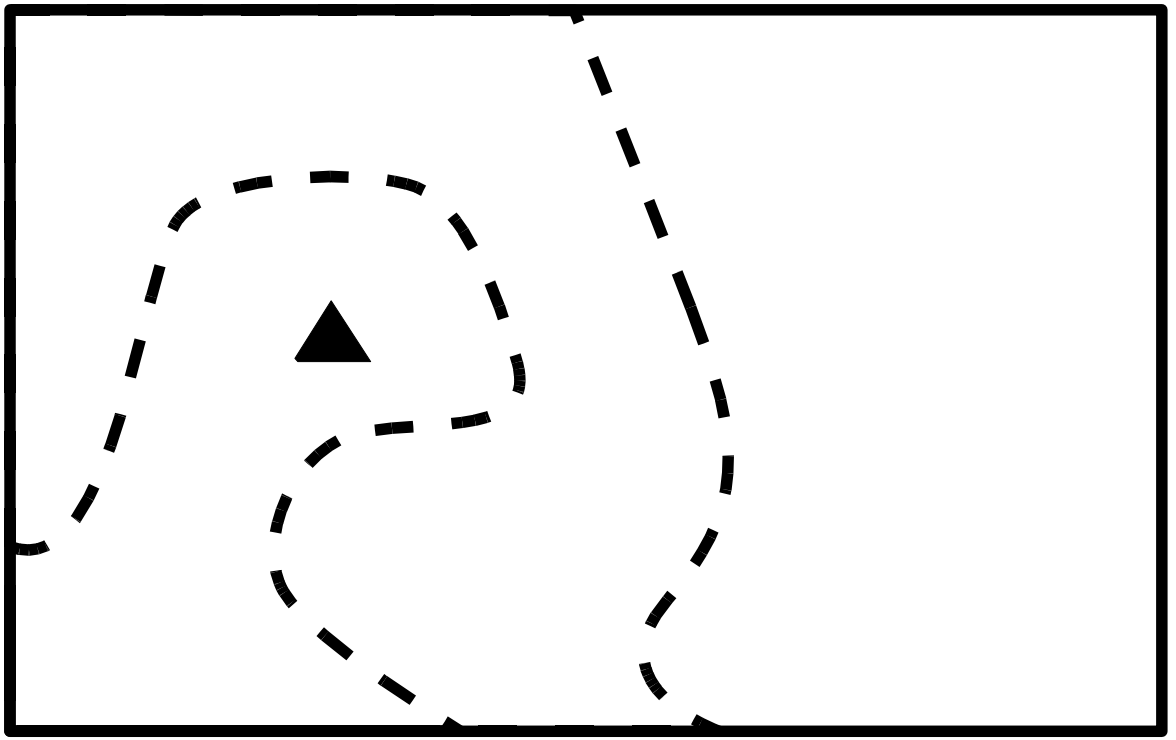

Value Iteration Variants

Heuristic Search (HSVI)

[Smith and Simmons, 2004]

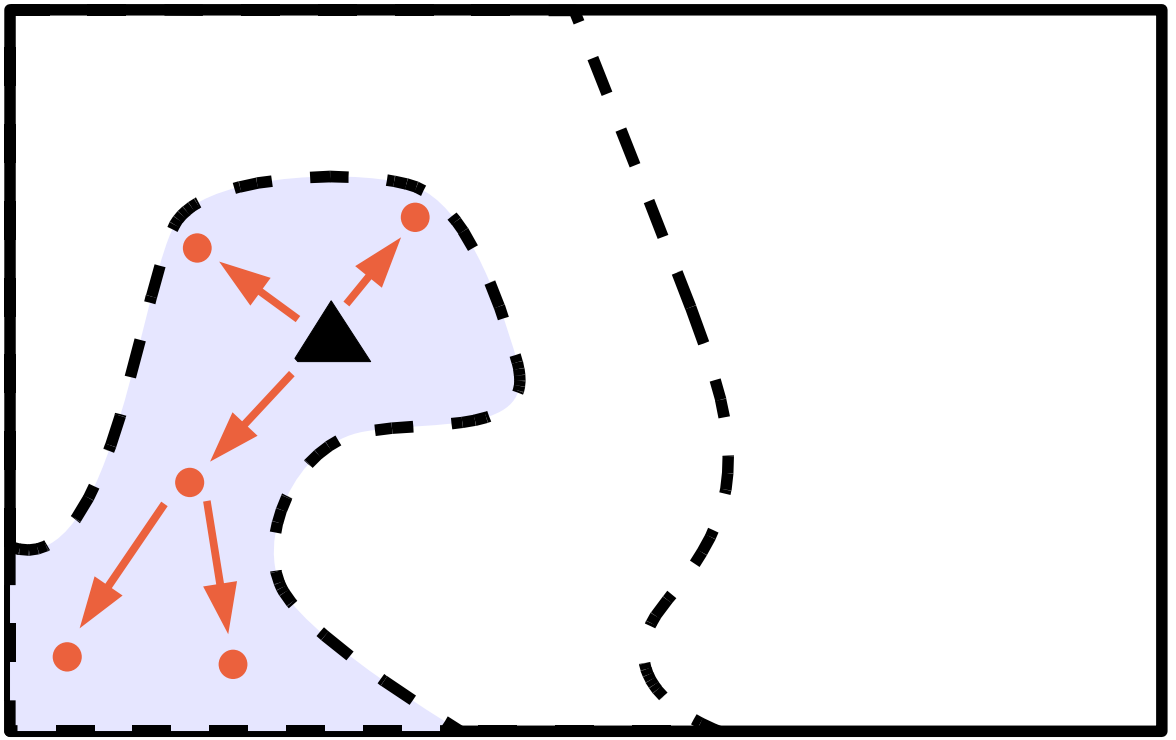

Outline

- Motivation
- Background and related work
- ***Technical approach***
- Preliminary work
- Proposed research

Intelligent Site Survey

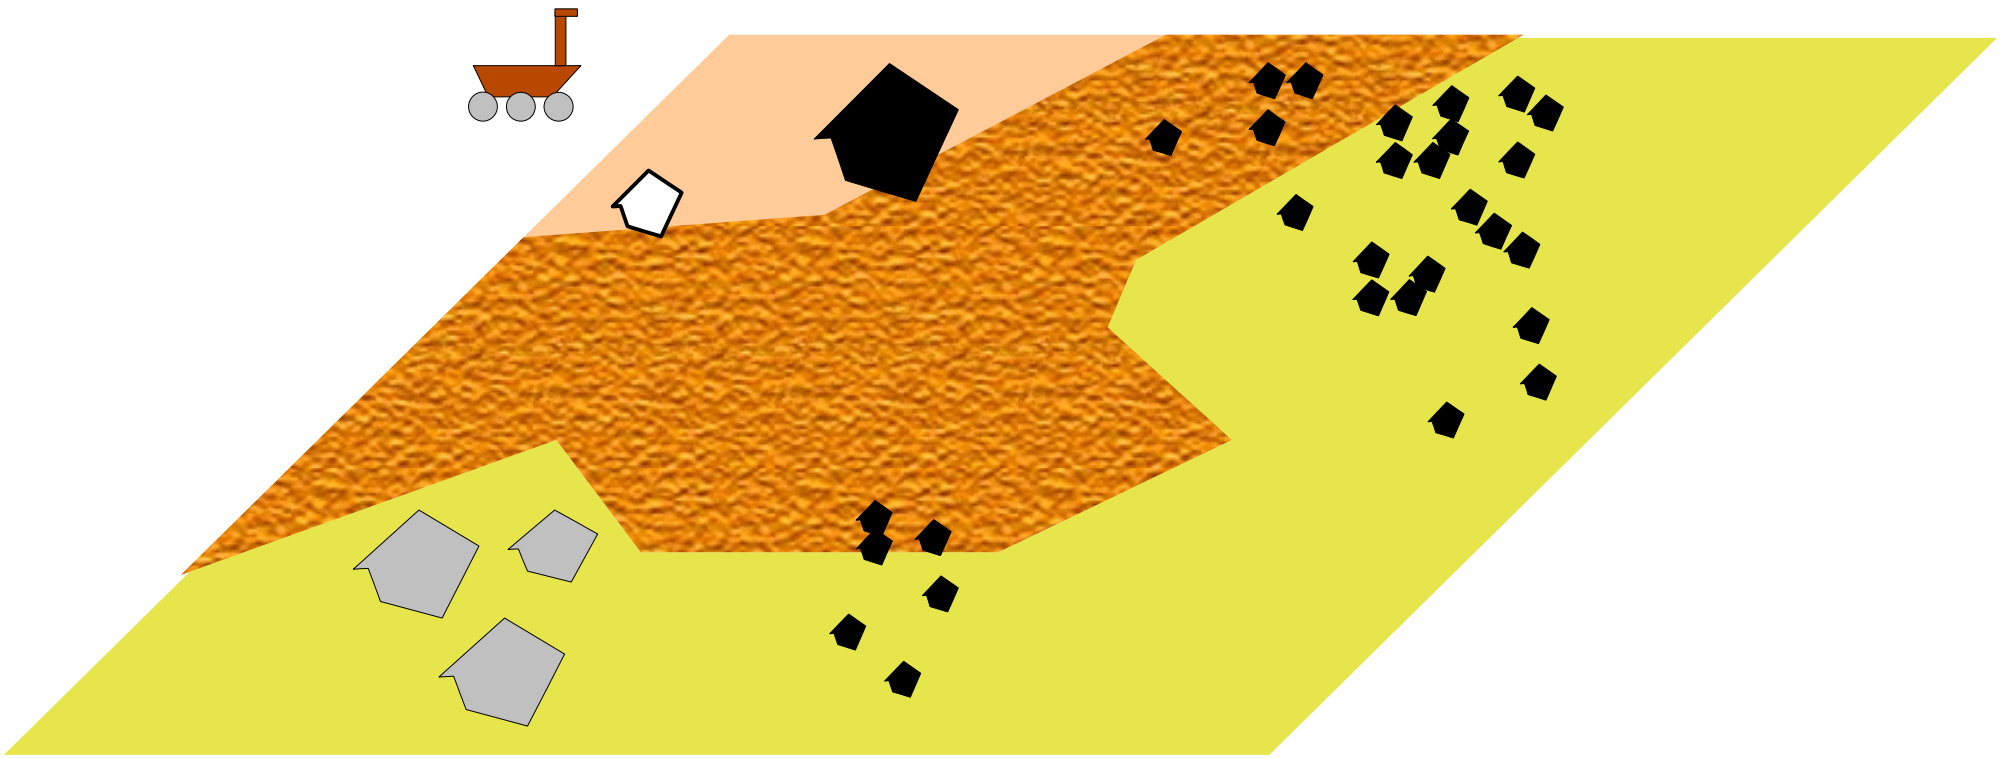

Intelligent Site Survey

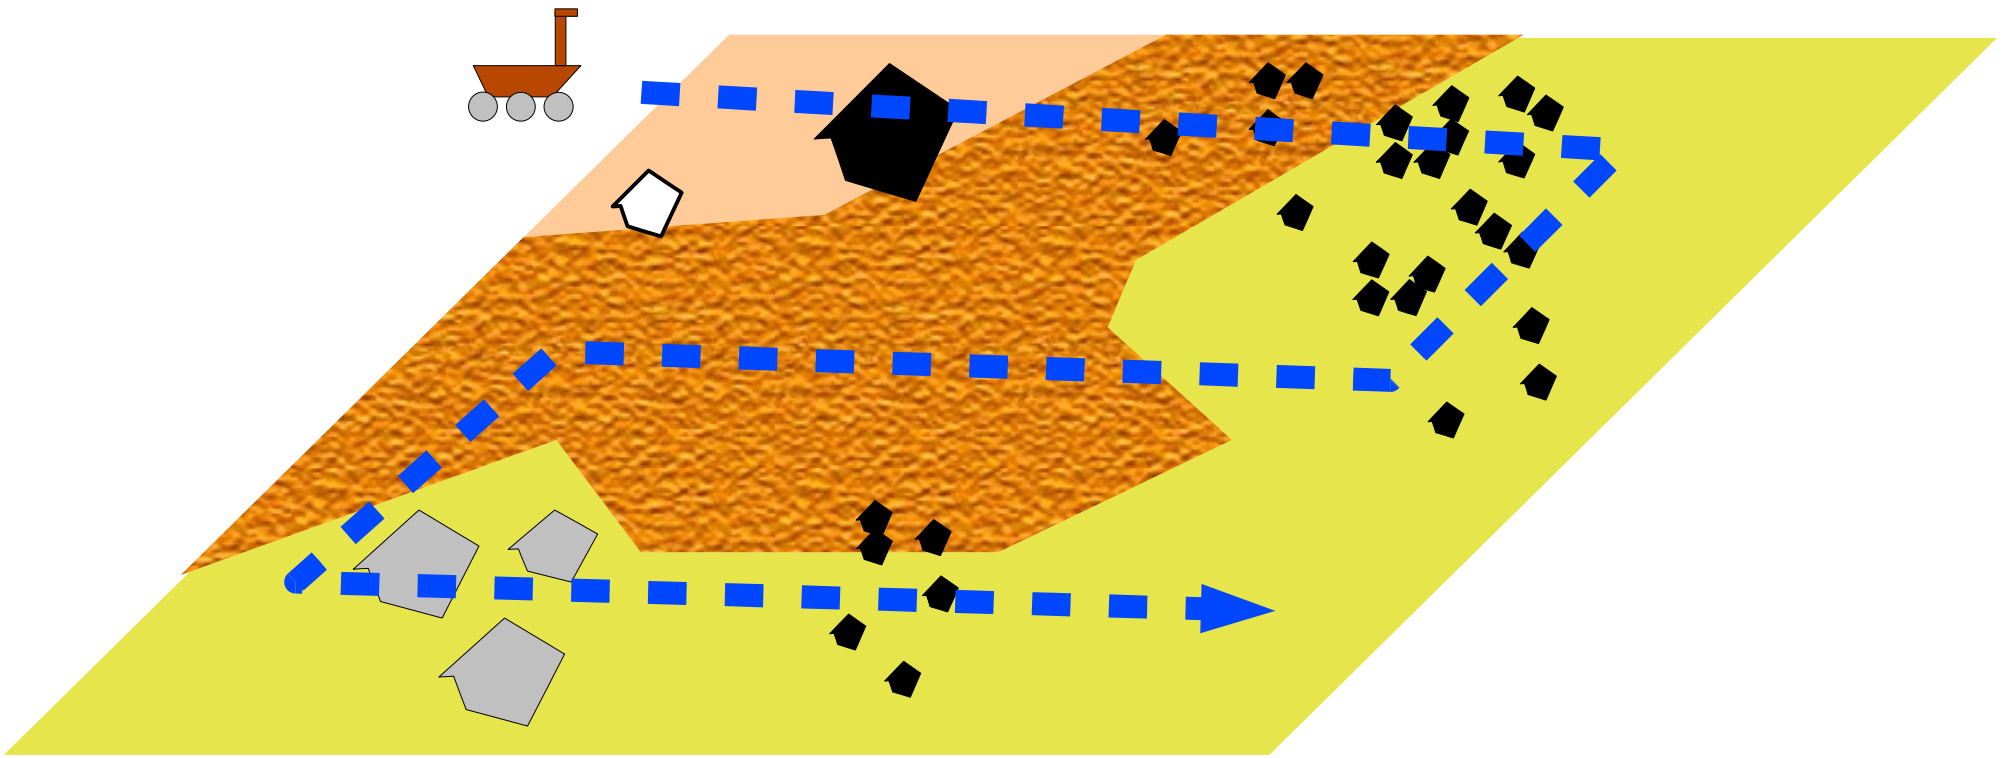

Intelligent Site Survey

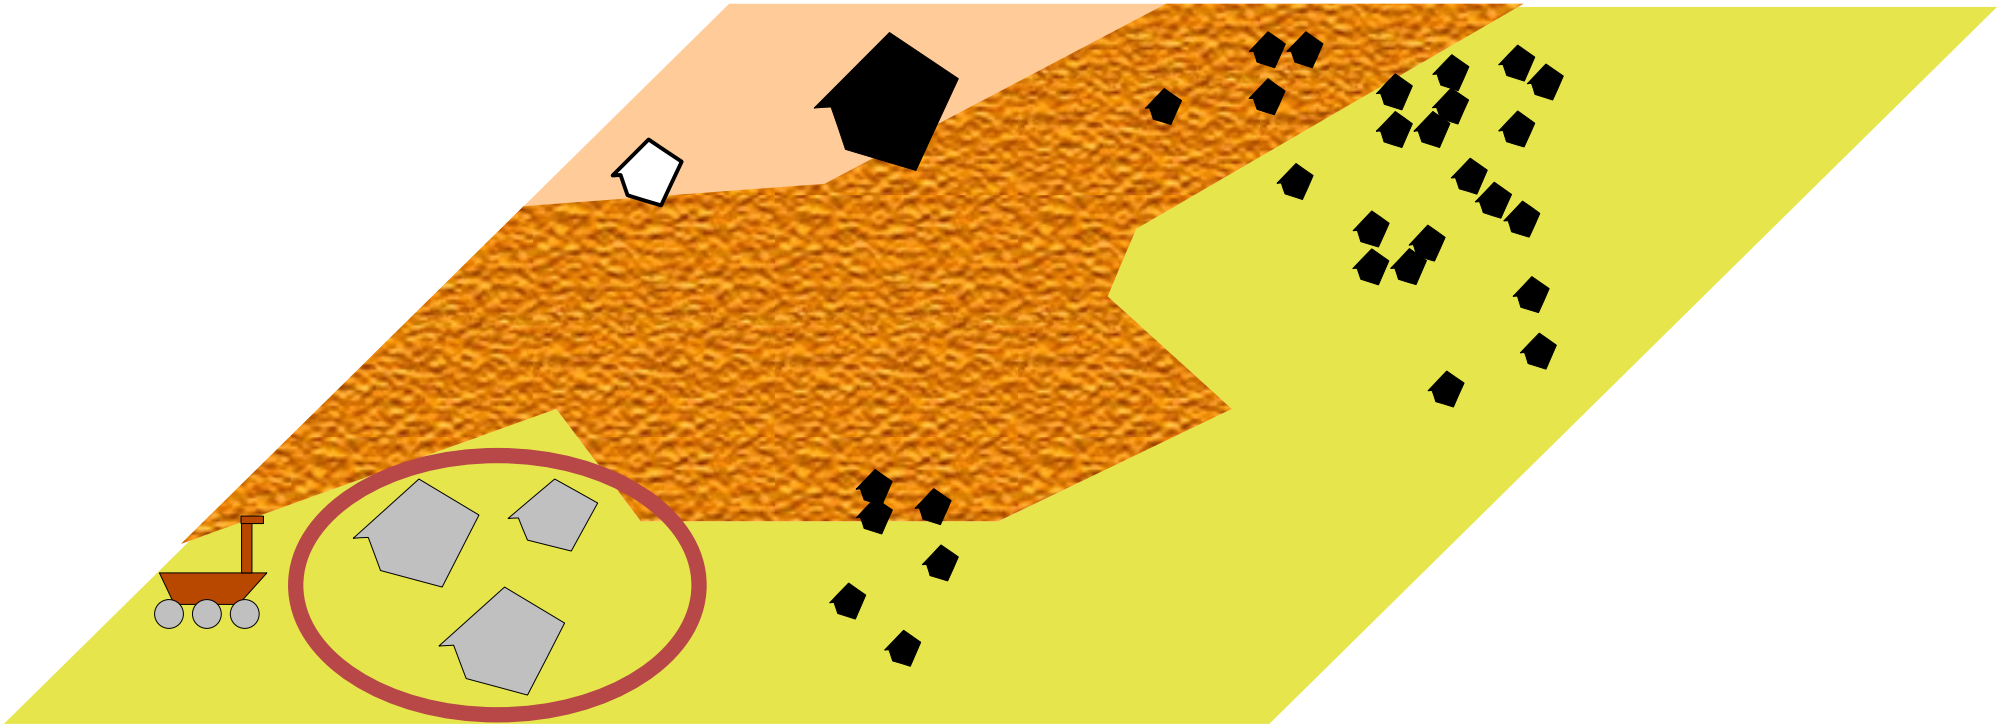

Intelligent Site Survey

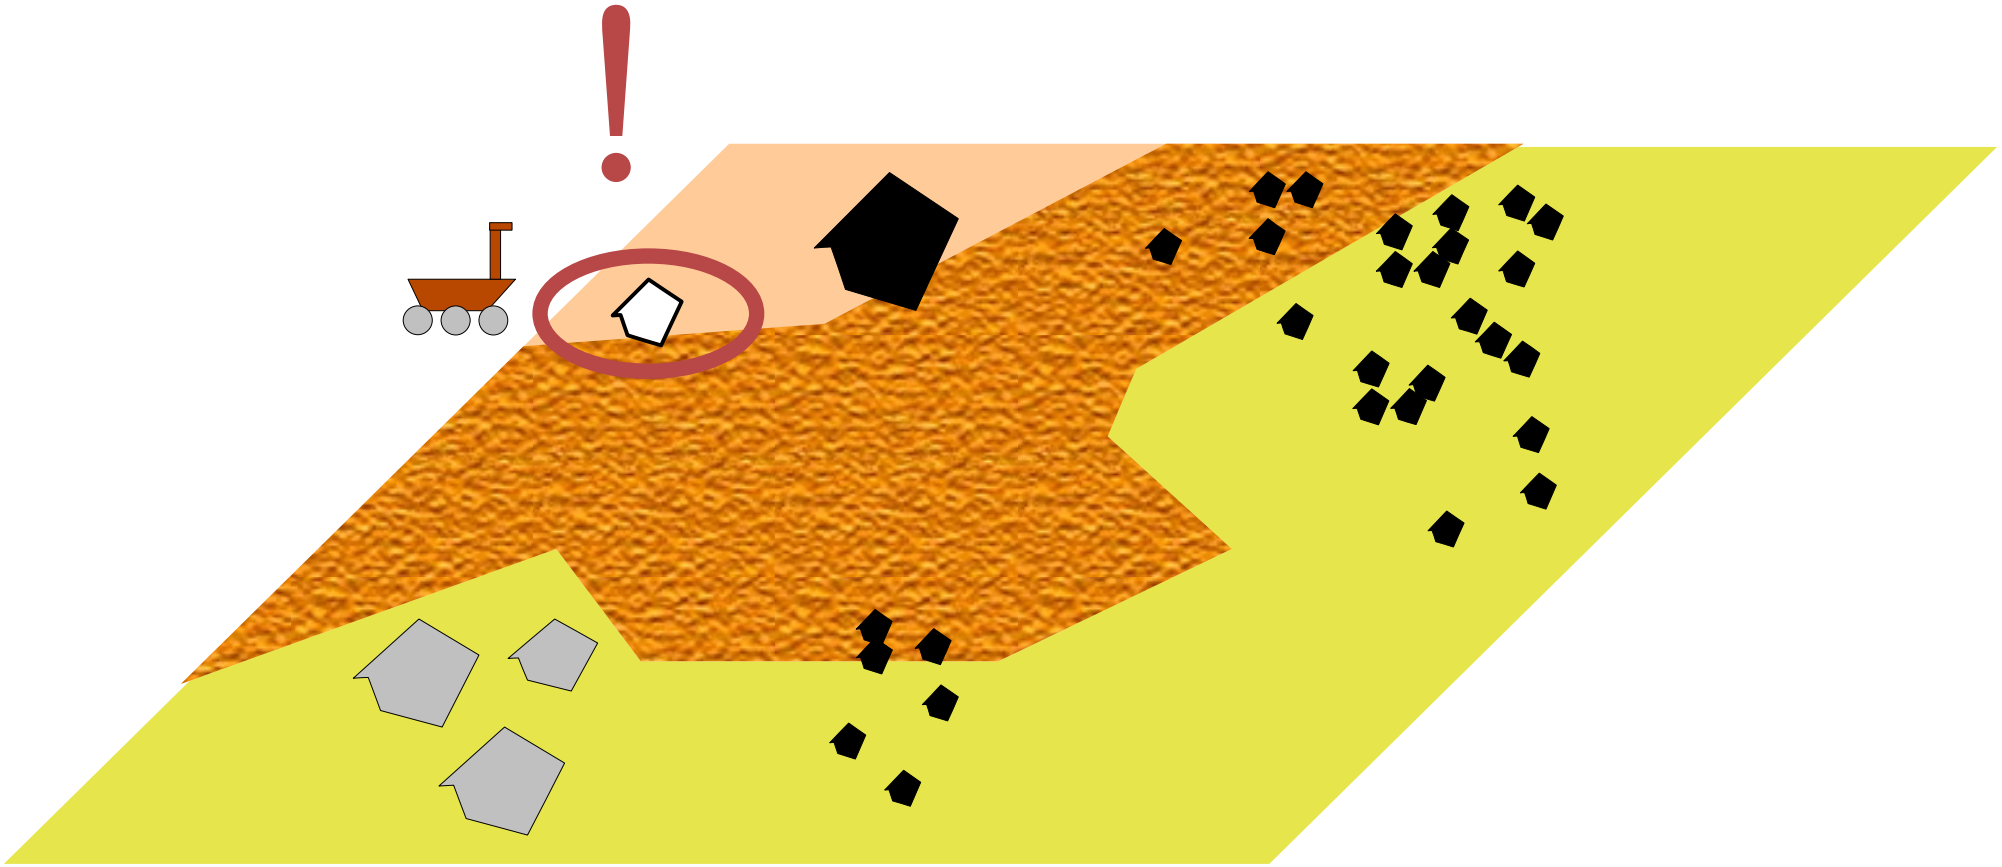

Intelligent Site Survey

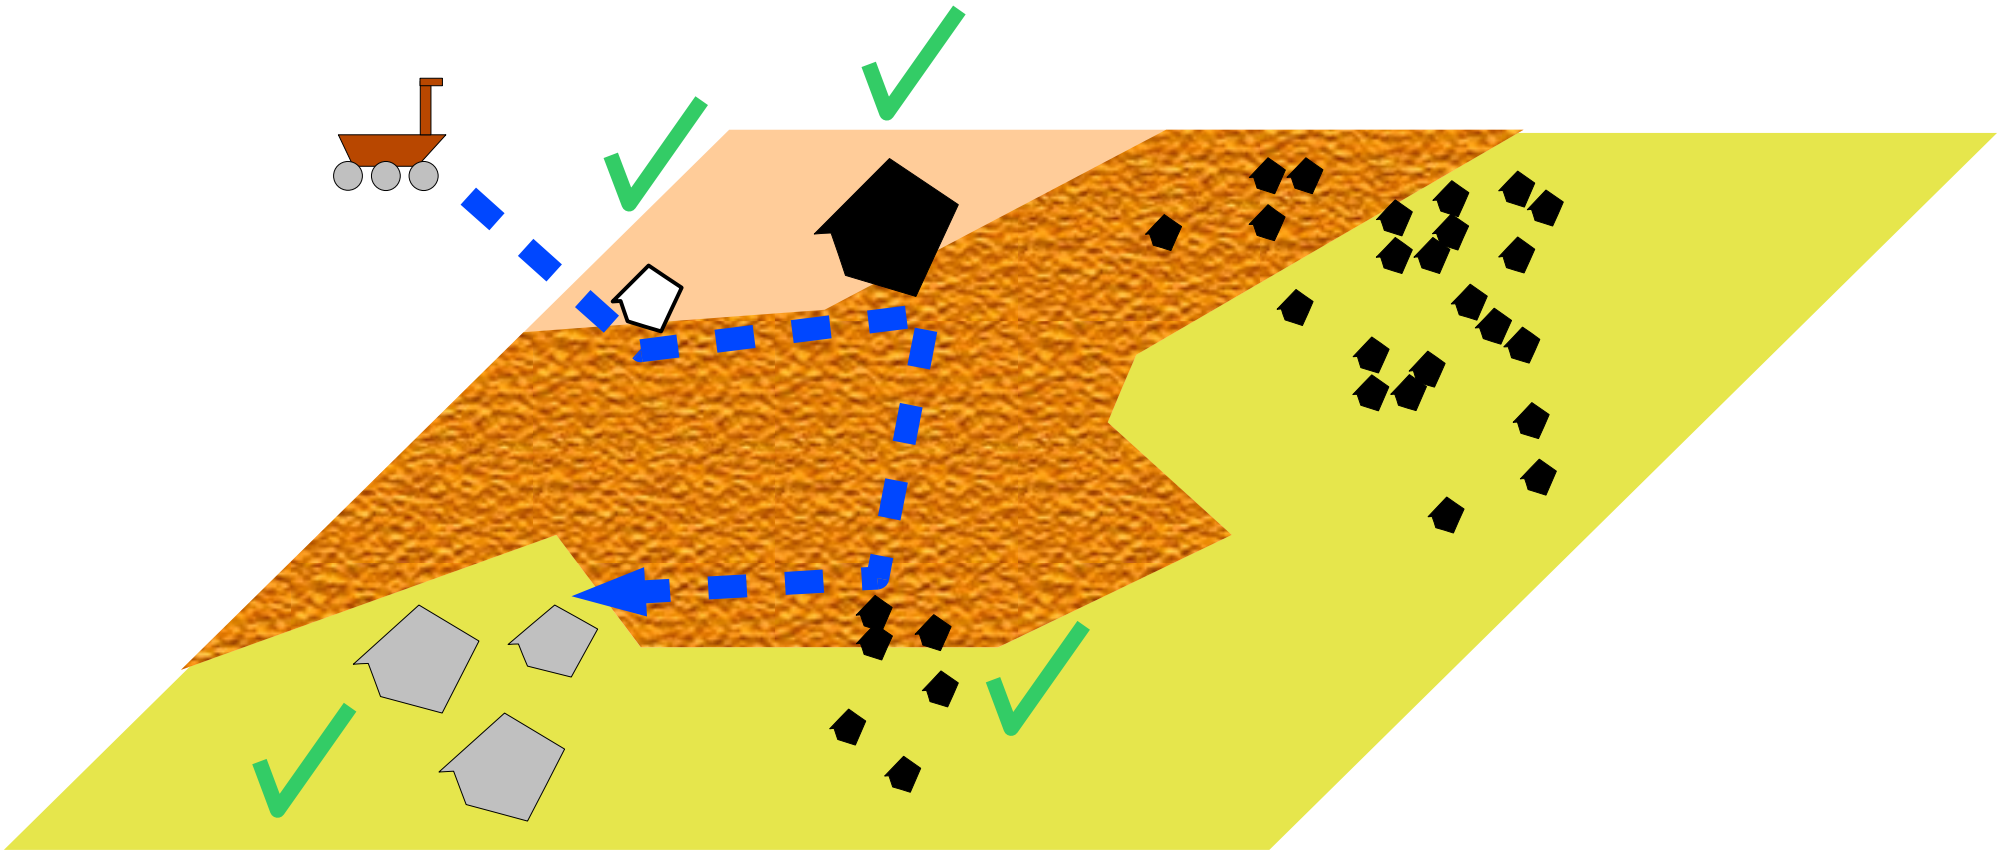

System Architecture

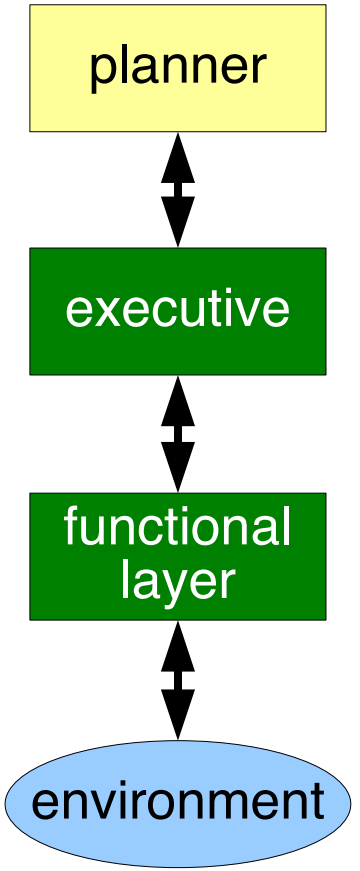

System Architecture

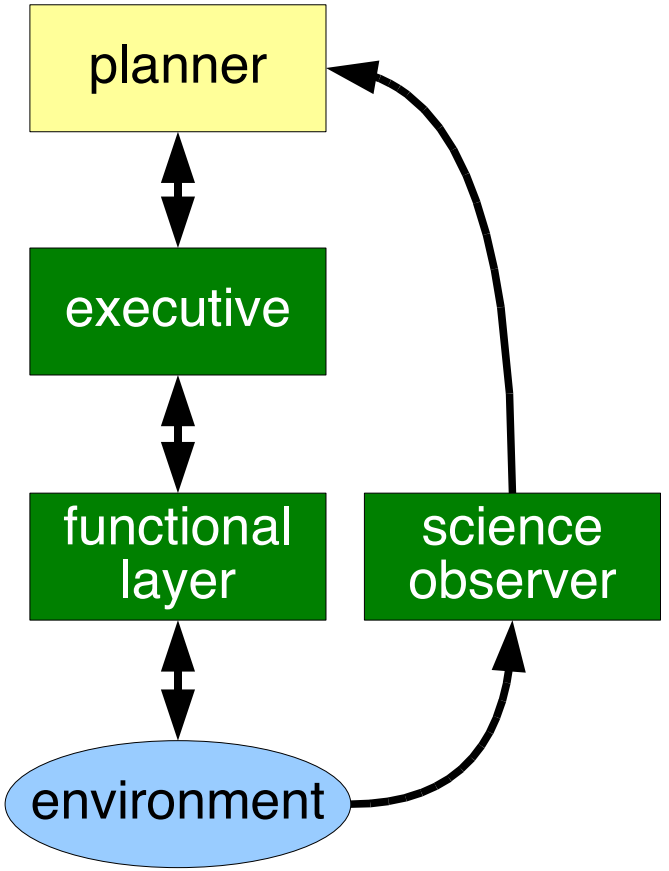

System Architecture

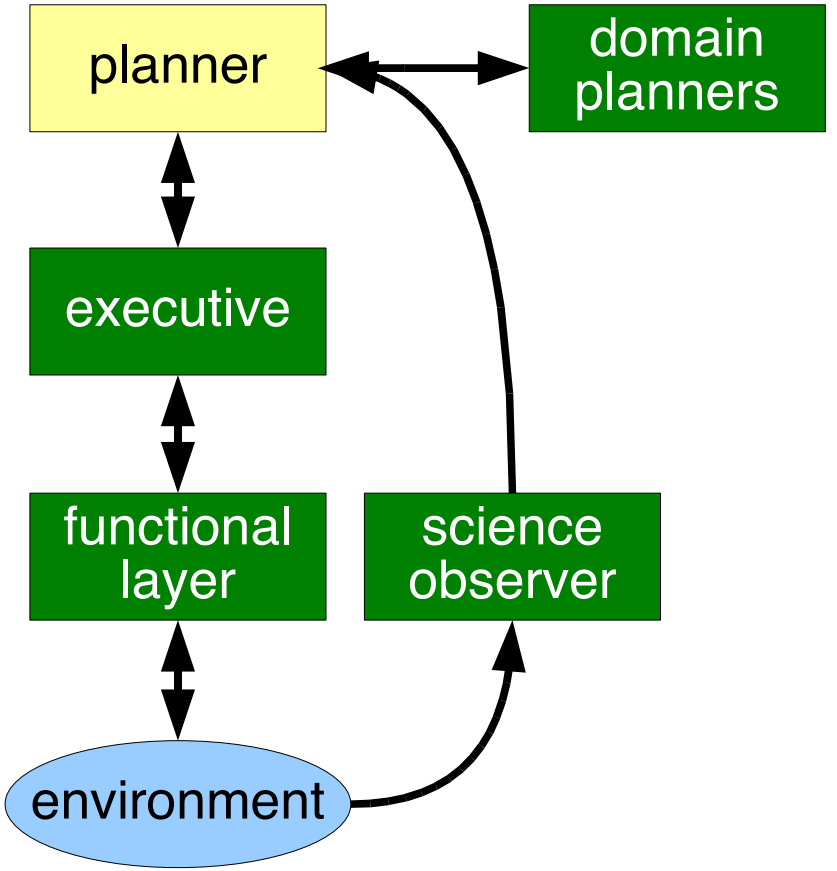

System Architecture

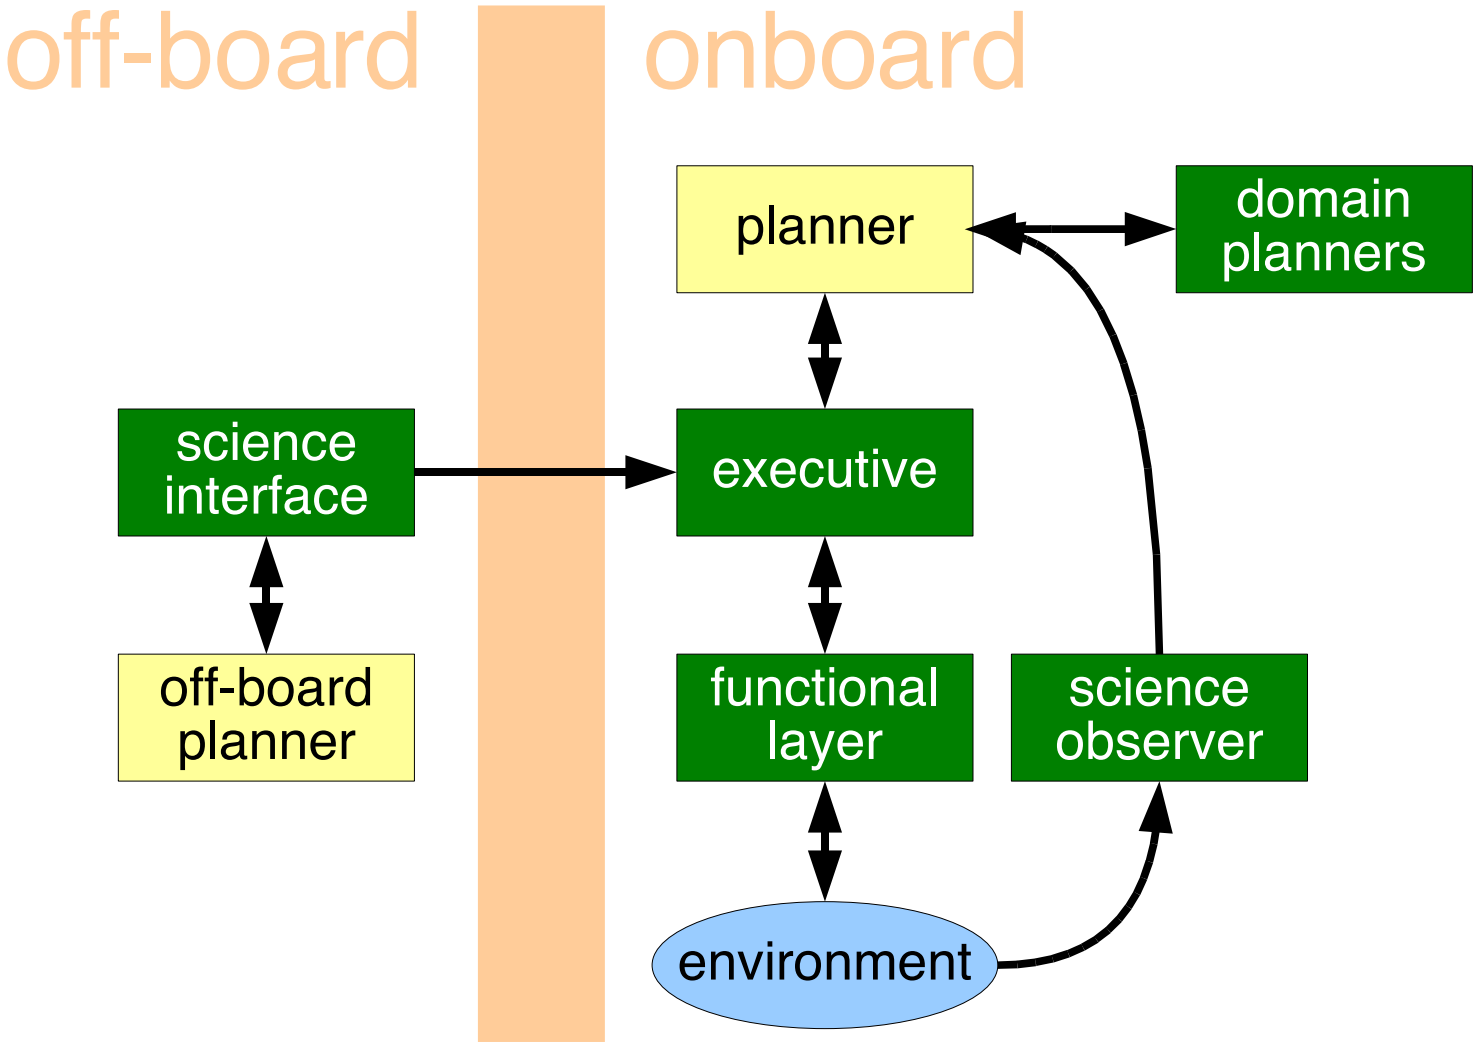

Pushing POMDP Technology

- **Heuristic search**
  - Combine heuristic search with efficient value function representations
- **Factored state**
  - Do better when the state is **mostly observable**
  - Data structures that leverage **independence**
- **Continuous planning**
  - Fast replanning

Outline

- Motivation
- Background and related work
- Technical approach
- ***Preliminary work***
- Proposed research

Heuristic Search Value Iteration

- Solves POMDPs
- Fast anytime algorithm with provable bounds on plan quality
- On some benchmark problems, >100x speedup

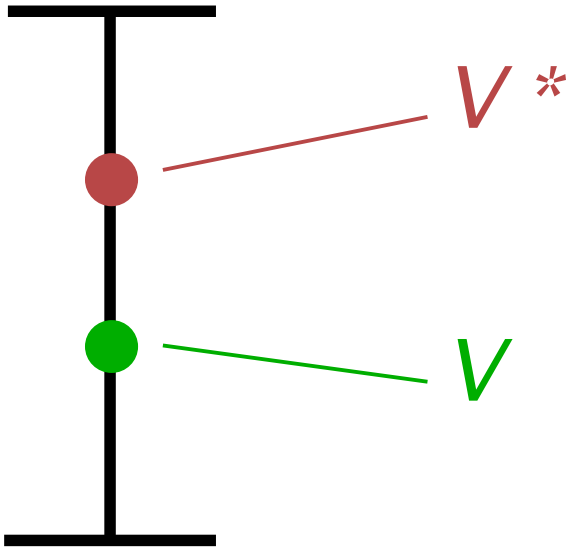

Heuristic Search Value Iteration

- Focus on reachable beliefs
- Avoid considering foolish actions and unlikely observations

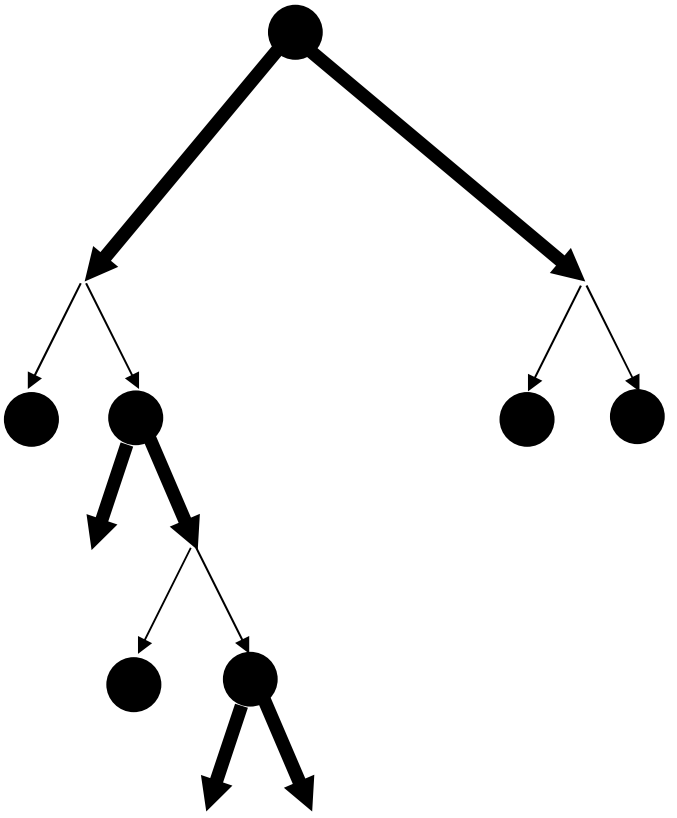

Heuristic Search Value Iteration

- Focus on reachable beliefs
- Avoid considering foolish actions and unlikely observations

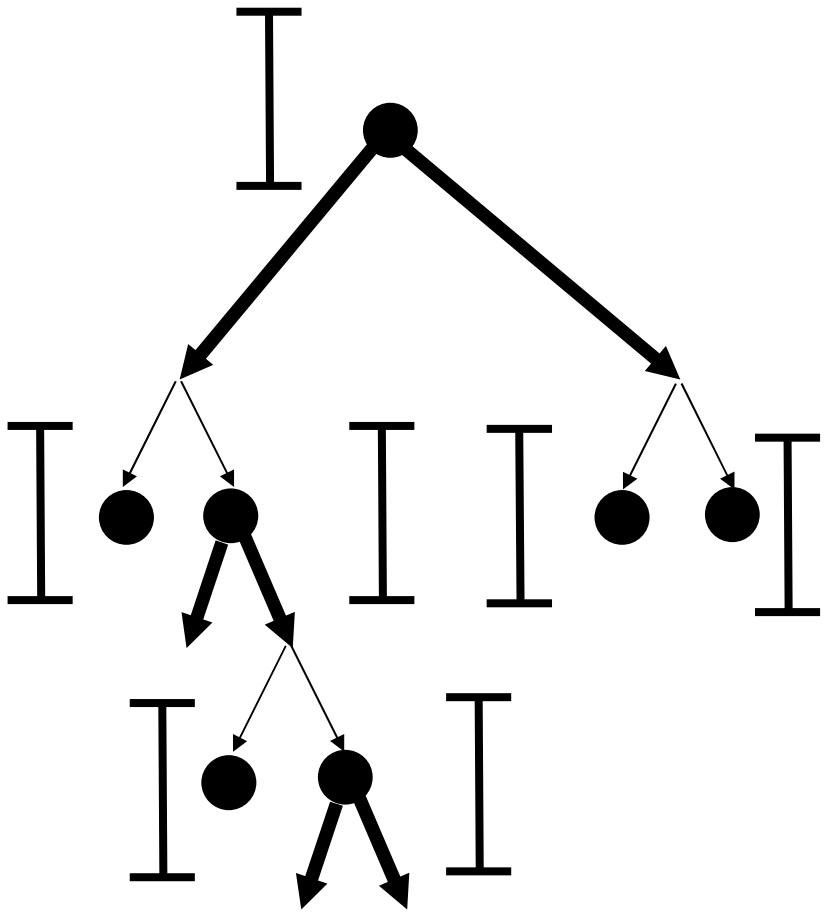

Heuristic Search Value Iteration

- Focus on reachable beliefs
- Avoid considering foolish actions and unlikely observations

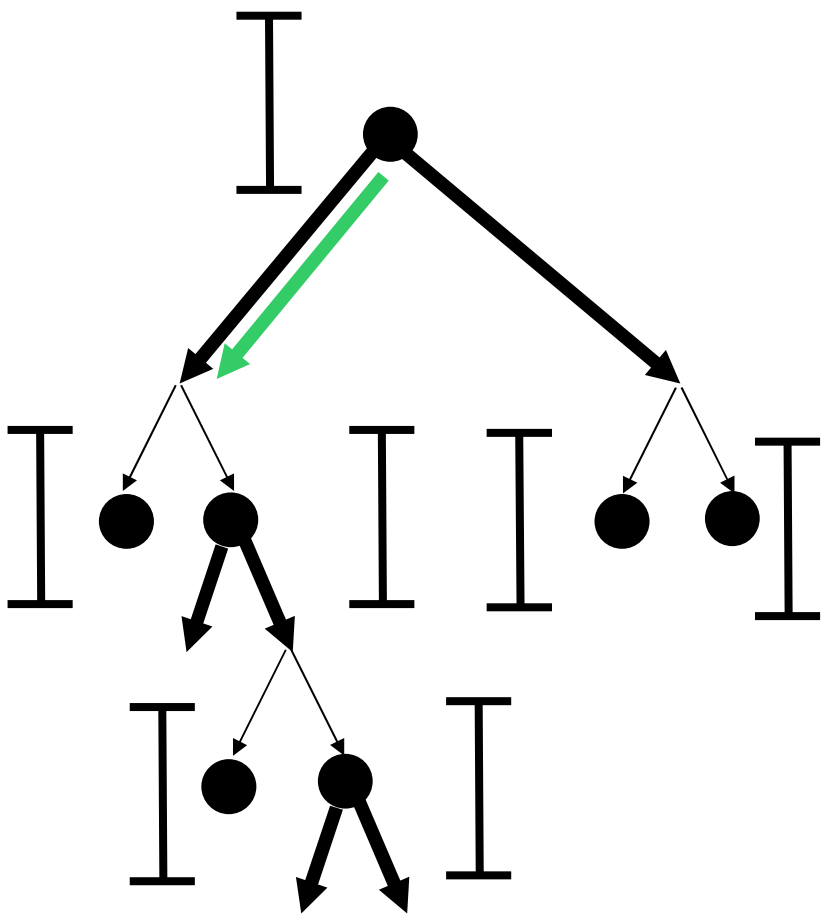

Heuristic Search Value Iteration

- Focus on reachable beliefs
- Avoid considering foolish actions and unlikely observations

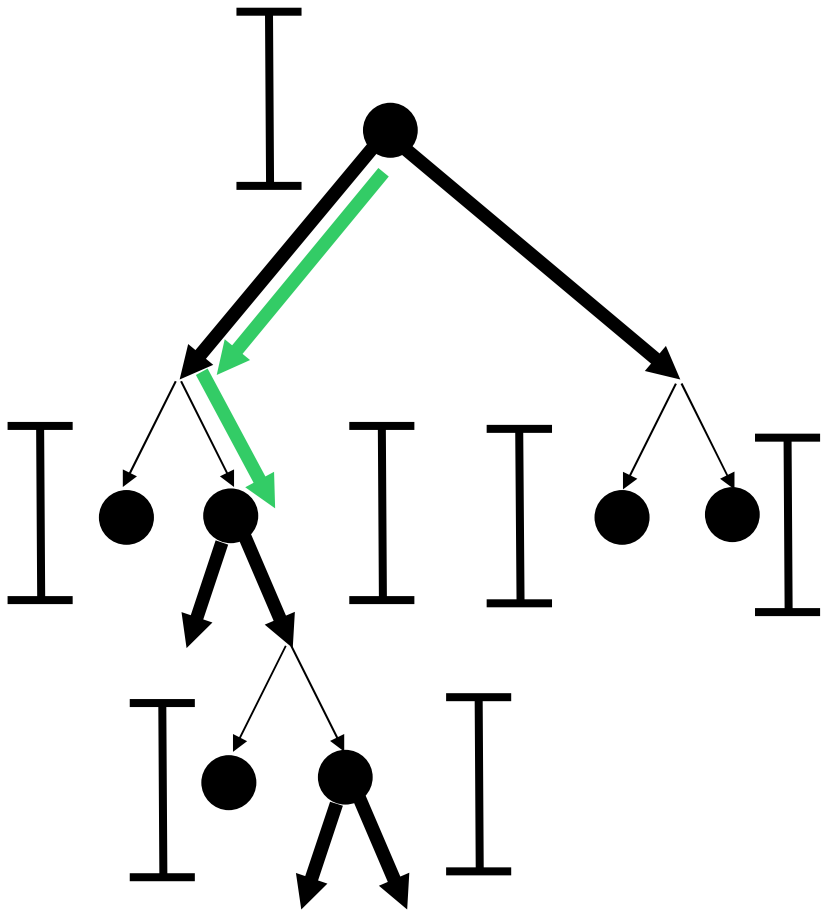

Heuristic Search Value Iteration

- Focus on reachable beliefs
- Avoid considering foolish actions and unlikely observations

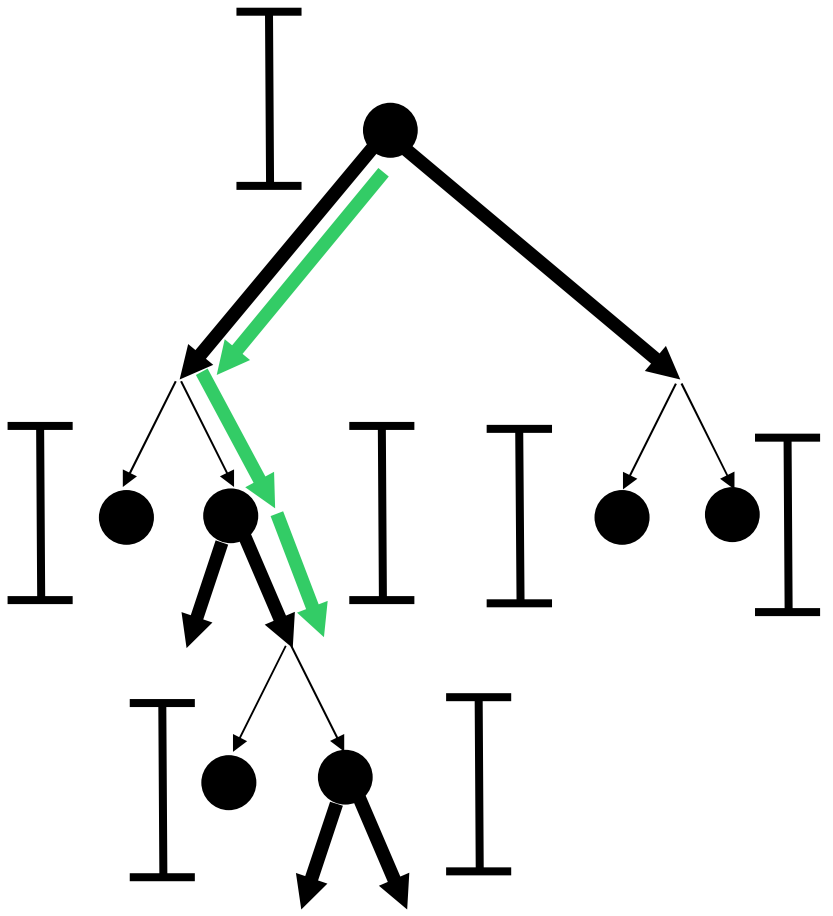

Heuristic Search Value Iteration

- Focus on reachable beliefs
- Avoid considering foolish actions and unlikely observations

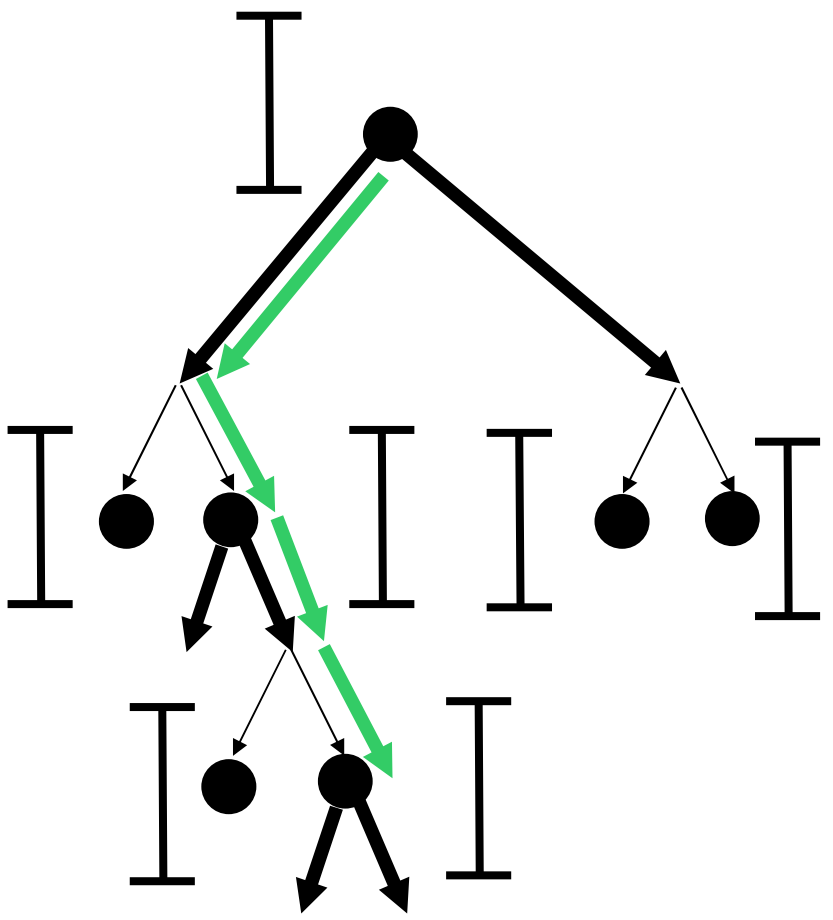

Heuristic Search Value Iteration

- Focus on reachable beliefs
- Avoid considering foolish actions and unlikely observations

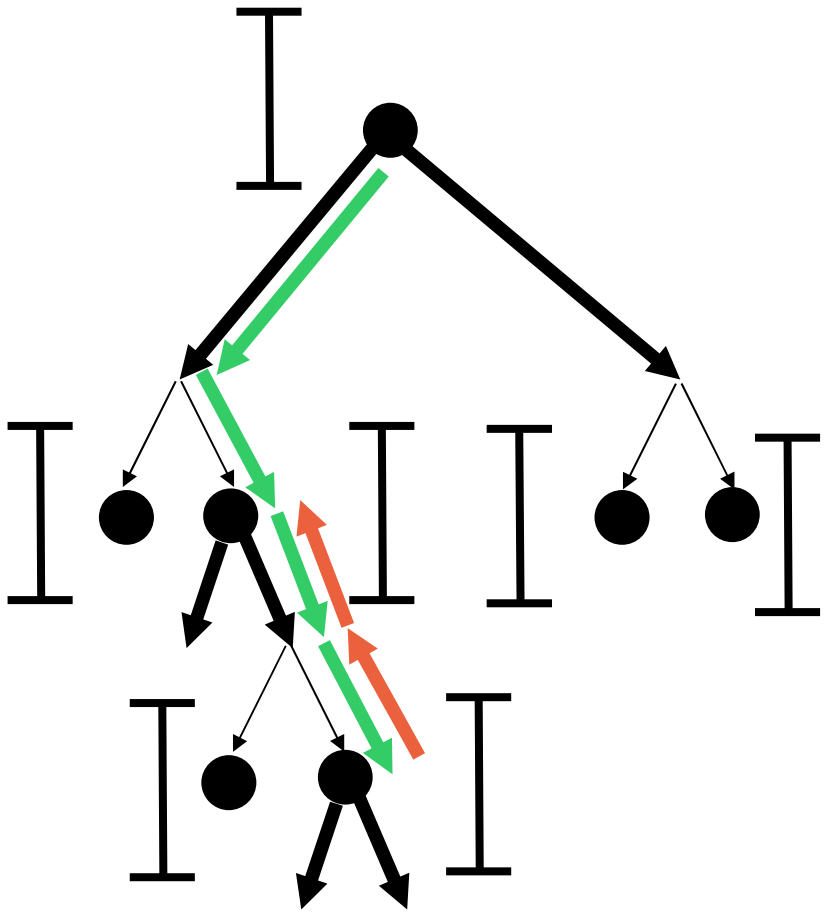

Heuristic Search Value Iteration

- Focus on reachable beliefs
- Avoid considering foolish actions and unlikely observations

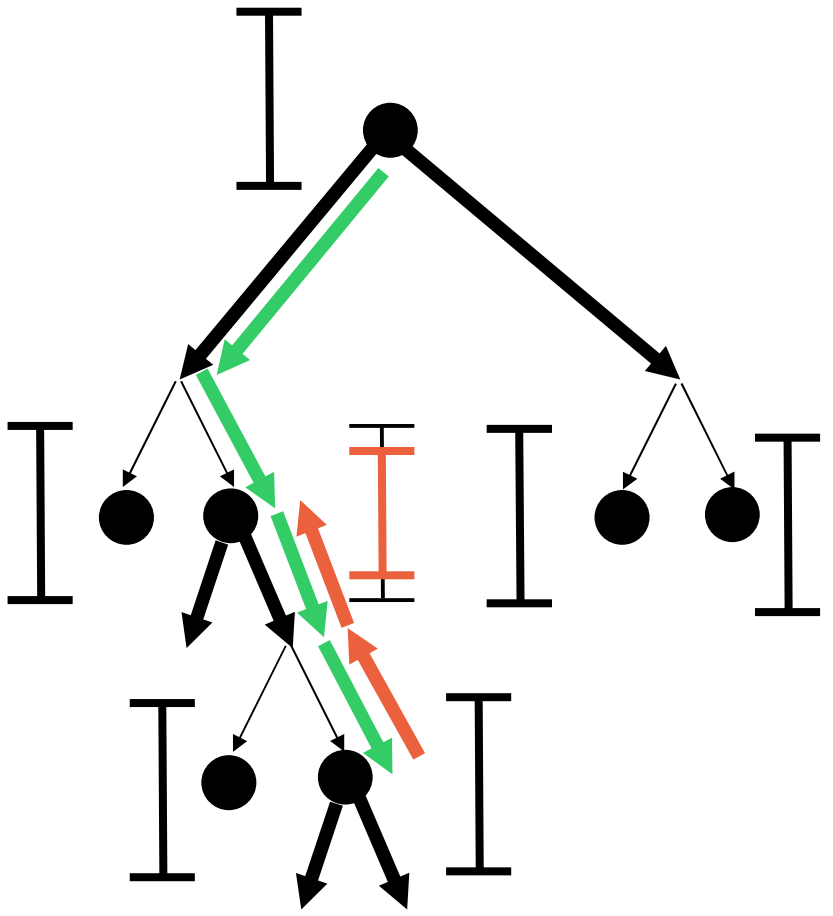

Heuristic Search Value Iteration

- Focus on reachable beliefs
- Avoid considering foolish actions and unlikely observations

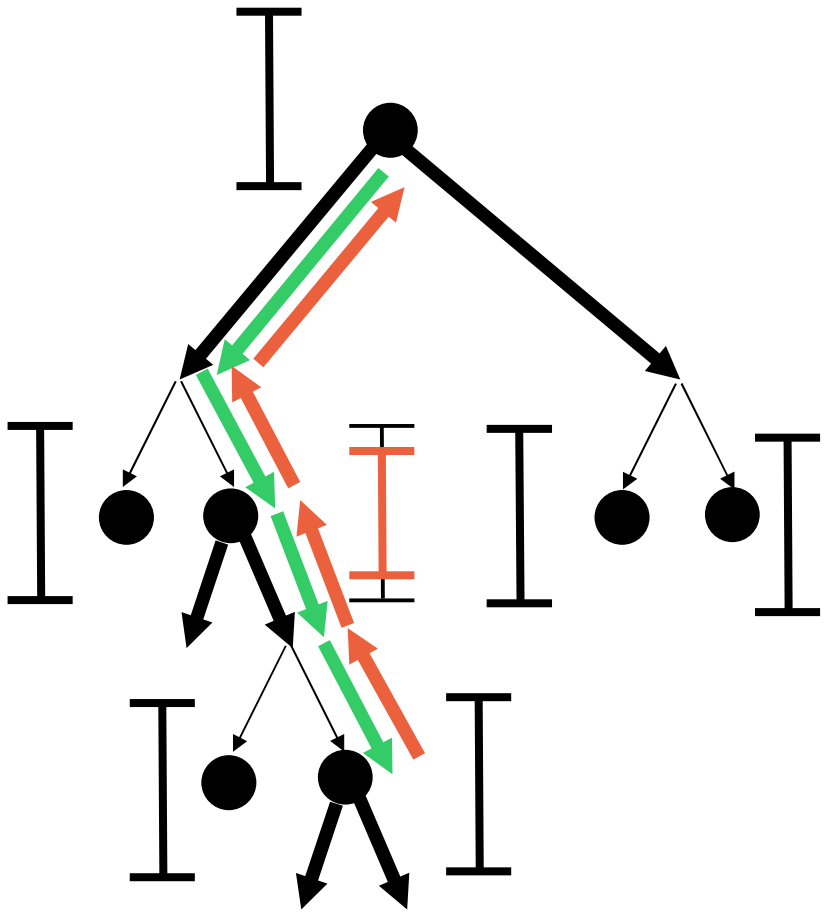

Heuristic Search Value Iteration

- Focus on reachable beliefs
- Avoid considering foolish actions and unlikely observations

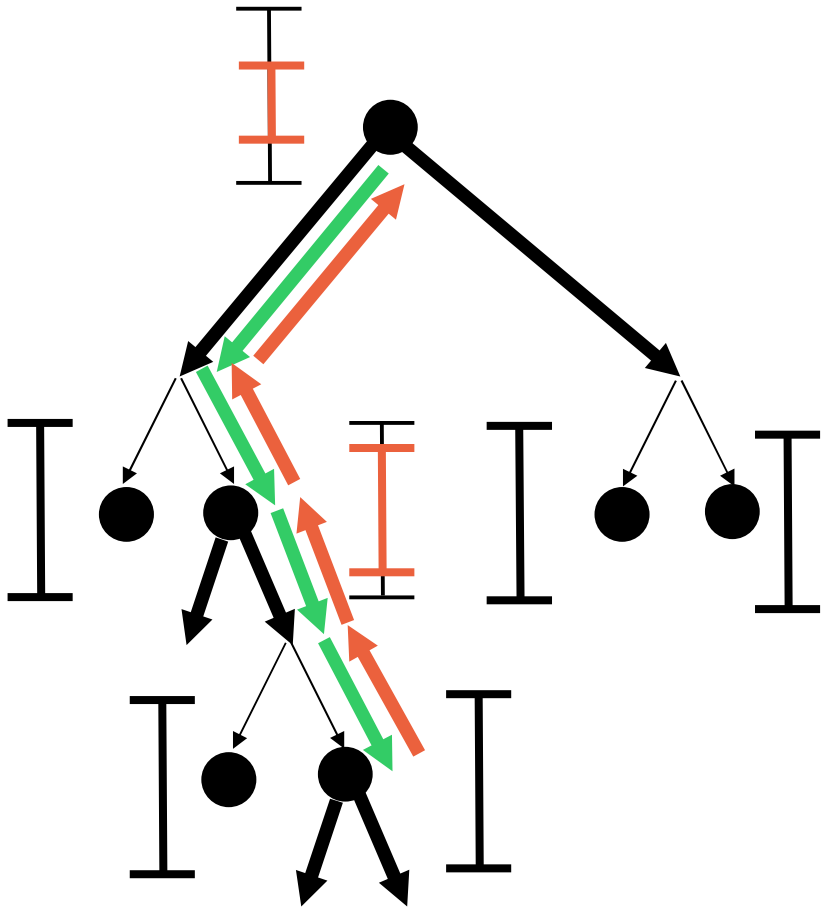

Heuristic Search Value Iteration

- Focus on reachable beliefs
- Avoid considering foolish actions and unlikely observations

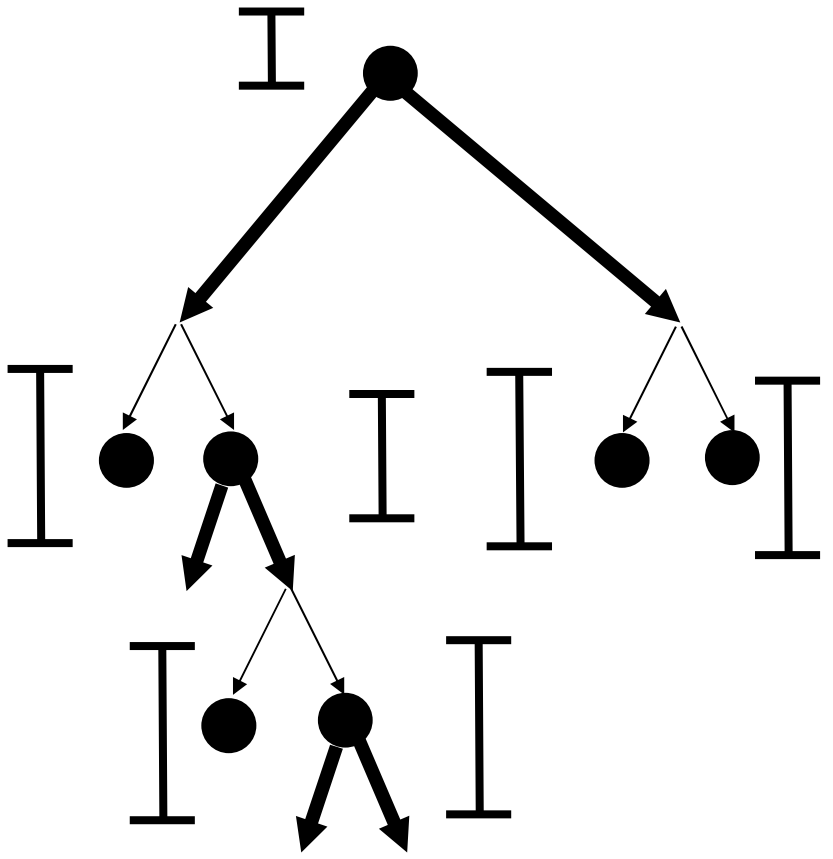

Heuristic Search Value Iteration

- Focus on reachable beliefs
- Avoid considering foolish actions and unlikely observations

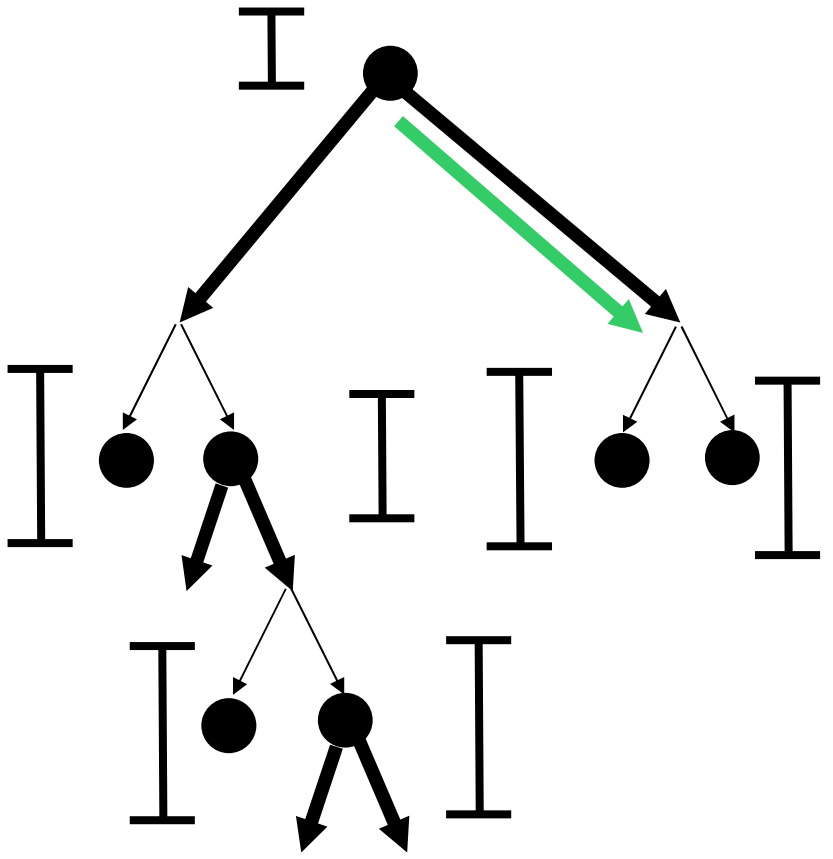

Heuristic Search Value Iteration

- Focus on reachable beliefs
- Avoid considering foolish actions and unlikely observations

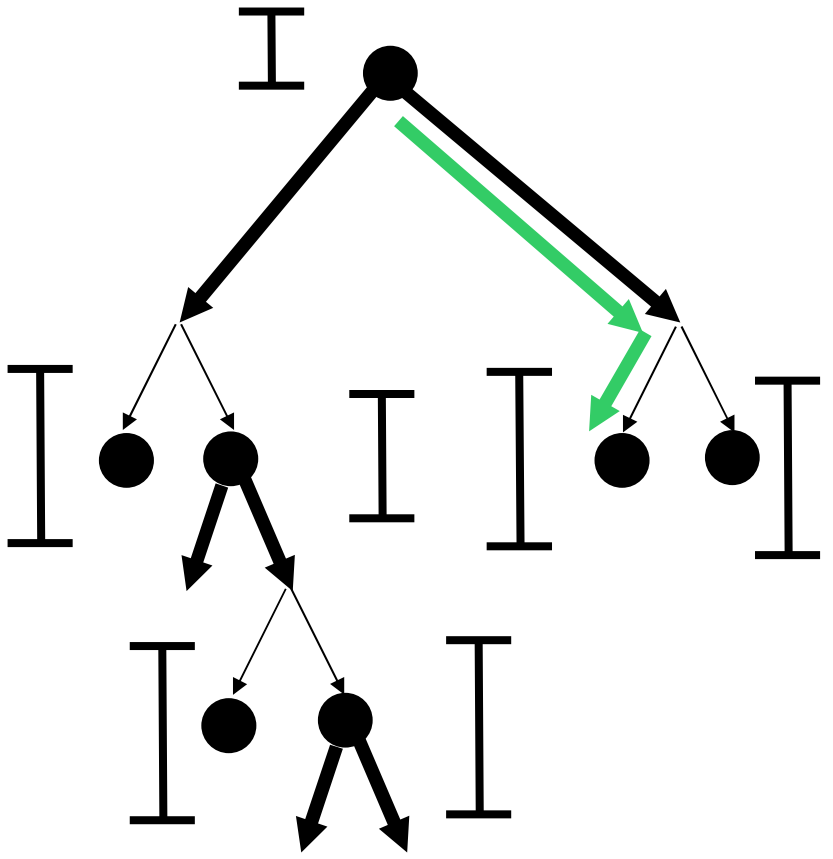

Heuristic Search Value Iteration

- Focus on reachable beliefs
- Avoid considering foolish actions and unlikely observations

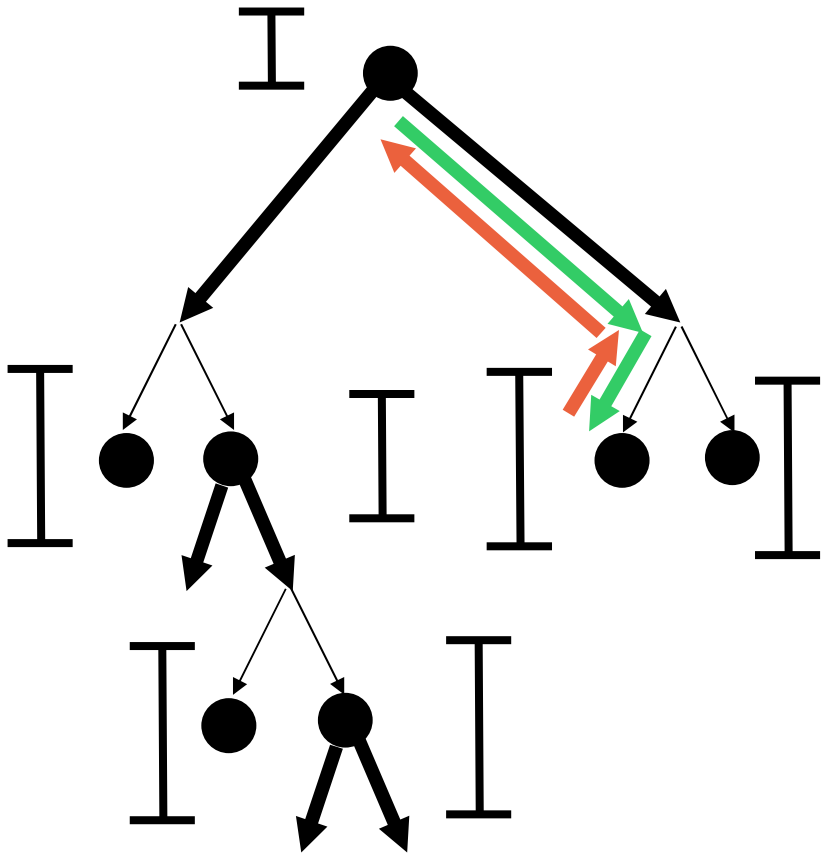

Heuristic Search Value Iteration

- Focus on reachable beliefs
- Avoid considering foolish actions and unlikely observations

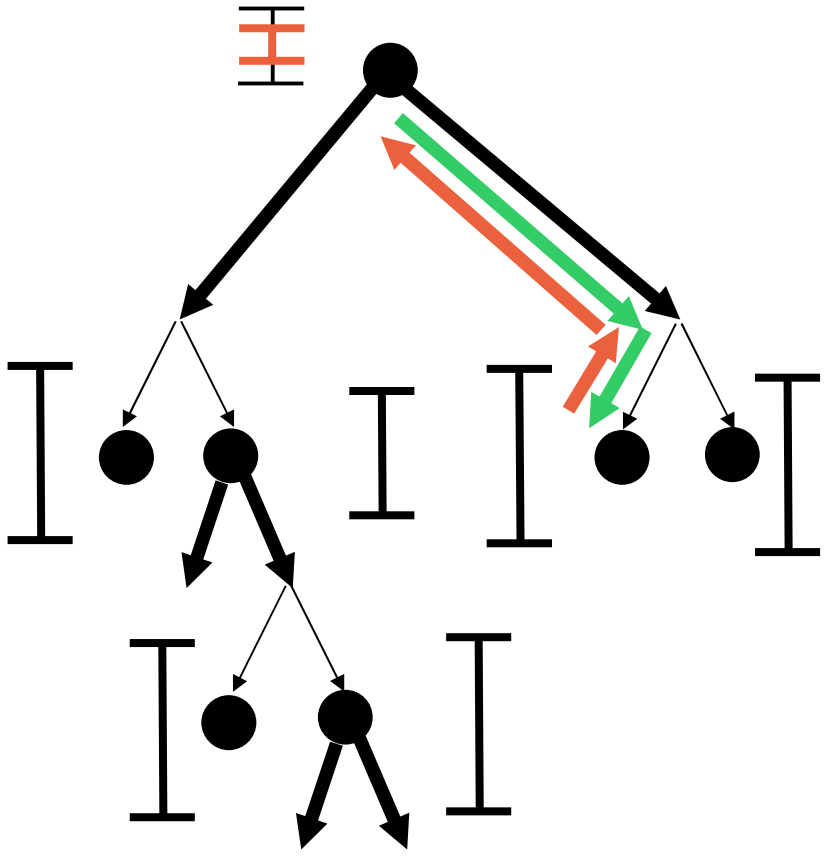

Heuristic Search Value Iteration

- Focus on reachable beliefs
- Avoid considering foolish actions and unlikely observations

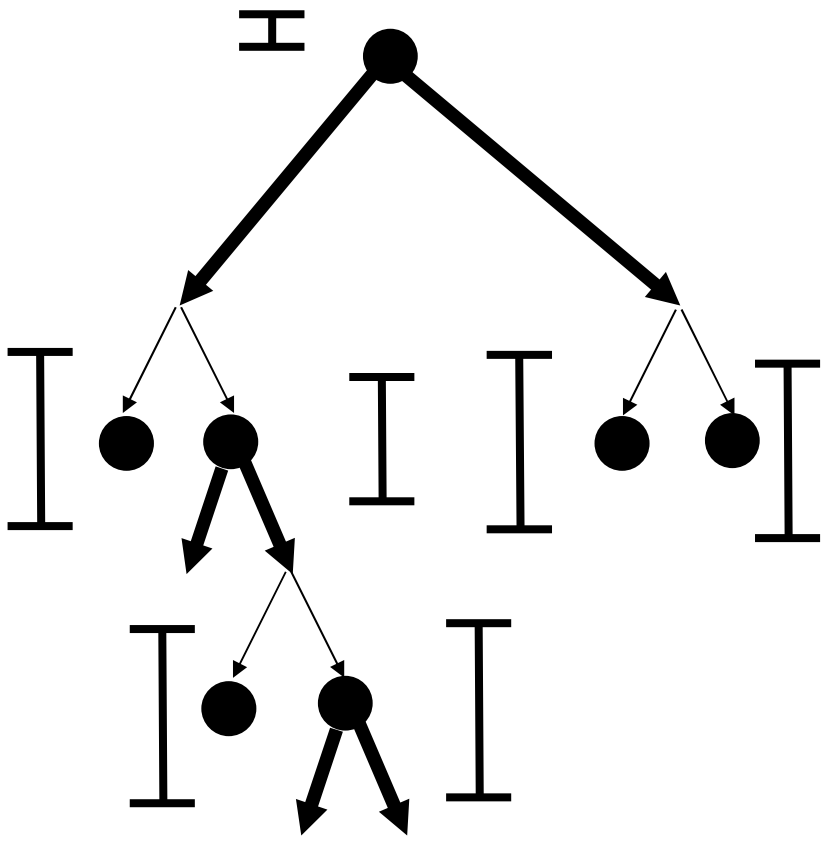

Bounds Update at *b*

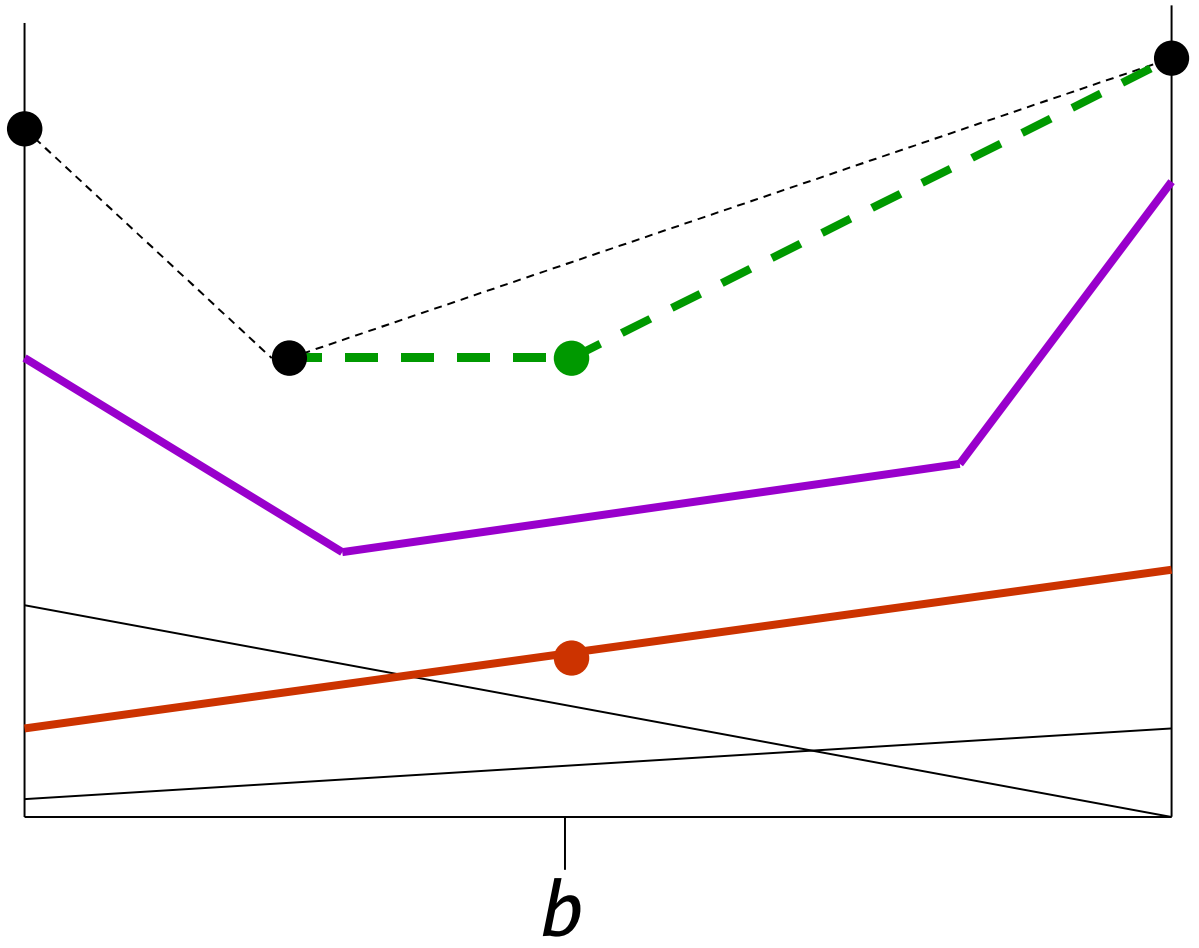

HSVI Performance

Tag (870s 5a 30o)

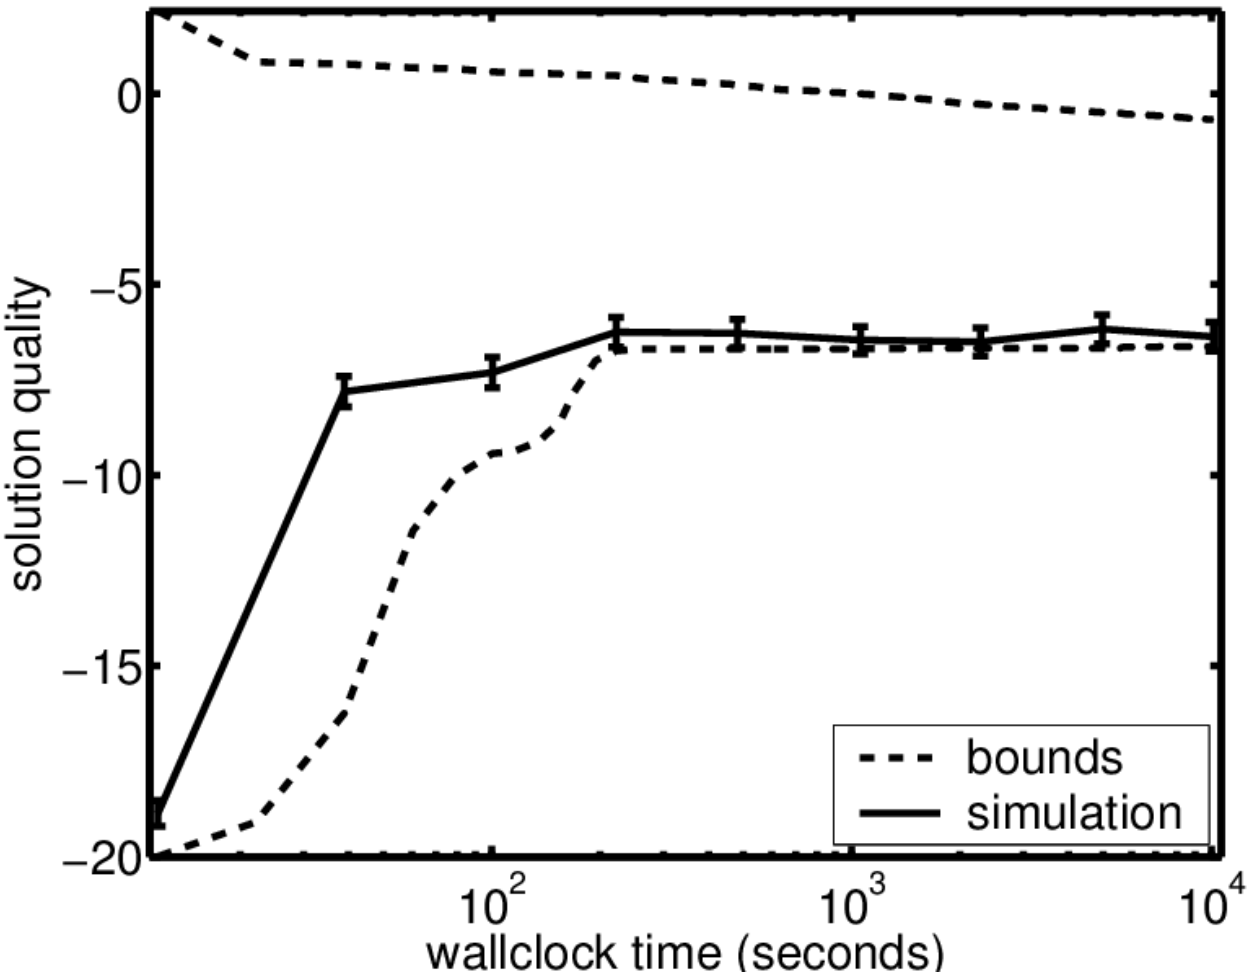

Multi-Algorithm Comparison

Problem (num. states/actions/observations)	Goal%	Reward	Time (s)	$\|\Gamma\|$
Tiger-Grid (36s 5a 17o)				
QMDP [Pineau et al., 2003]	n.a.	0.198	0.19	n.a.
Grid [Brafman, 1997]	n.a.	0.94	n.v.	174
PBUA [Poon, 2001]	n.a.	2.30	12116	660
PBVI [Pineau et al., 2003]	n.a.	2.25	3448	470
HSVI	n.a.	2.35	10341	4860
Hallway (61s 5a 21o)				
QMDP [Littman et al., 1995]	47.4	n.v.	n.v.	n.a.
PBUA [Poon, 2001]	100	0.53	450	300
PBVI [Pineau et al., 2003]	96	0.53	288	86
HSVI	100	0.52	10836	1341
Hallway2 (93s 5a 17o)				
QMDP [Littman et al., 1995]	25.9	n.v.	n.v.	n.a.
Grid [Brafman, 1997]	98	n.v.	n.v.	337
PBUA [Poon, 2001]	100	0.35	27898	1840
PBVI [Pineau et al., 2003]	98	0.34	360	95
HSVI	100	0.35	10010	1571
Tag (870s 5a 30o)				
QMDP [Pineau et al., 2003]	17	-16.769	13.55	n.a.
PBVI [Pineau et al., 2003]	59	-9.180	180880	1334
HSVI	100	-6.37	10113	1657
RockSample[4,4] (257s 9a 2o)				
PBVI [Pineau, personal communication]	n.a.	17.1	~2000	n.v.
HSVI	n.a.	18.0	577	458

Comparison with PBVI

[Pineau et al., 2003]

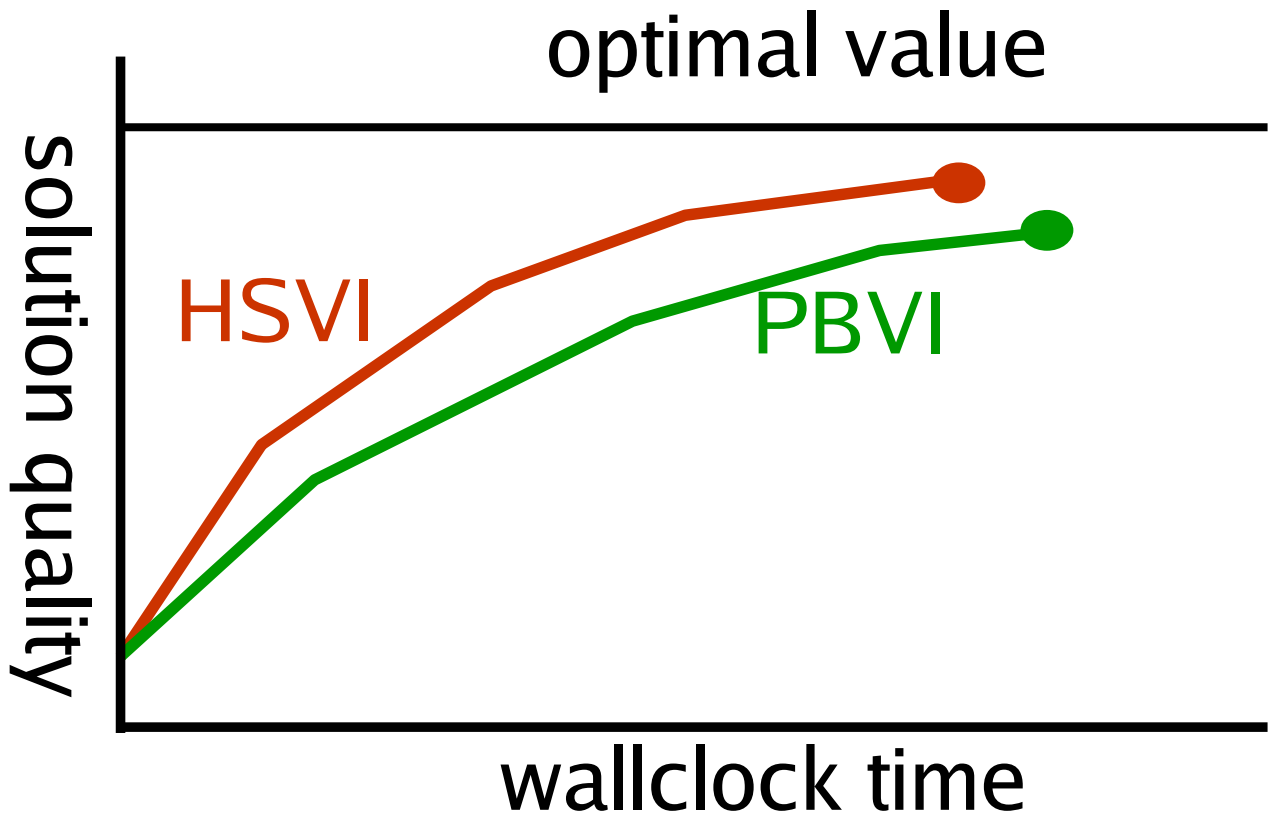

Comparison with PBVI

[Pineau et al., 2003]

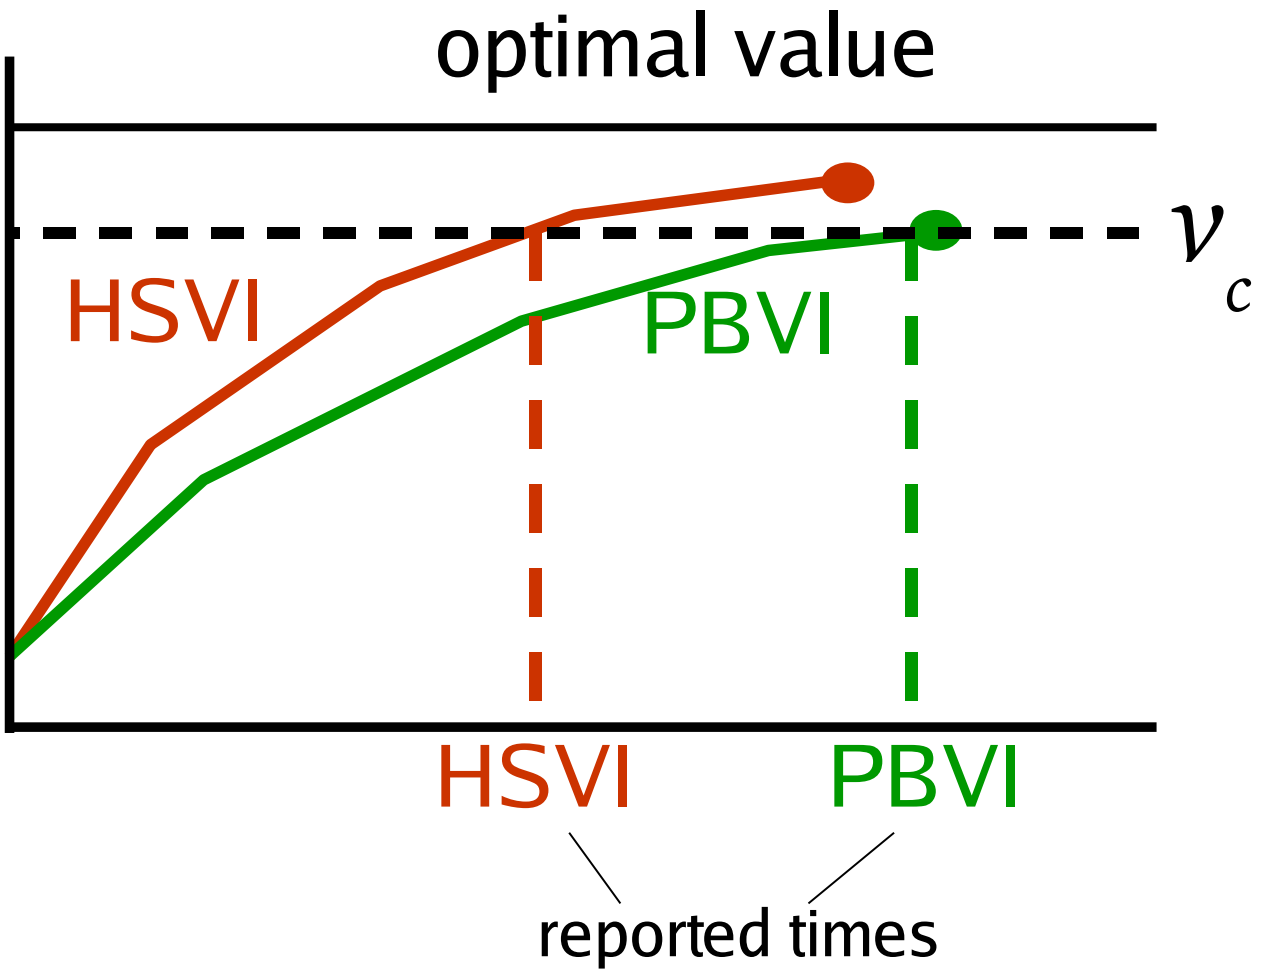

Comparison with PBVI

[Pineau et al., 2003]

Problem (num. states/actions/observations)	v_c	Time PBVI	Time HSVI	Speedup
Tiger-Grid (36 s 5a 17o)	2.25	3448	1053	3.3
Hallway (61s 5a 21o)	0.52	100-200	163	~1
Hallway2 (93s 5a 17o)	0.34	360	181	2.0
Tag (870s 5a 30o)	-9.18	180880	39	4600
RockSample[4,4] (256s 9a 2o)	17.1	~2000	23	~87

HSVI Related Work

	Applied to POMDPs	Asynchronous updates	Examines only reachable states/beliefs	Uses action heuristic	Uses observation/outcome heuristic	Leverages value function convexity	Keeps upper and lower bounds
HSVI	Y	Y	Y	Y	Y	Y	Y
ICUB/ICUL [Hauskrecht, 1997]	Y	Y	-	Y	-	Y	Y
BI-POMDP [Washington, 1997]	Y	Y	Y	Y	Y	-	Y
RTDP-BEL [Geffner and Bonet, 1998]	Y	Y	Y	Y	-	-	-
[Brafman, 1997]	Y	Y	-	Y	Y	Y	-
[Dearden and Boutilier, 1994]	-	Y	Y	Y	-	-	Y
LAO* [Hansen and Zilberstein, 2001]	-	Y	Y	Y	-	-	Y
PBVI [Pineau et al., 2003]	Y	-	Y	-	-	Y	-
PBDP [Zhang and Zhang, 2001]	Y	-	-	-	-	Y	-
Incremental pruning [Cassandra et al., 1997]	Y	-	-	-	-	Y	-

Outline

- Motivation
- Background and related work
- Technical approach
- Preliminary work
- ***Proposed research***

Research Goals

- Develop and field test an SA system, includes developing
  - Problem definition and performance criteria
  - Overall architecture
  - Planning module
- Extend POMDP state-of-art to meet the needs of the SA problem

Performance Criteria: SA

- Overall performance: did the science team characterize the site accurately?
- Subgoals with quantitative metrics
  - Coverage, key signatures, anomalies, representative sampling
- Compare to **baseline operational modes**

Baseline Operational Modes

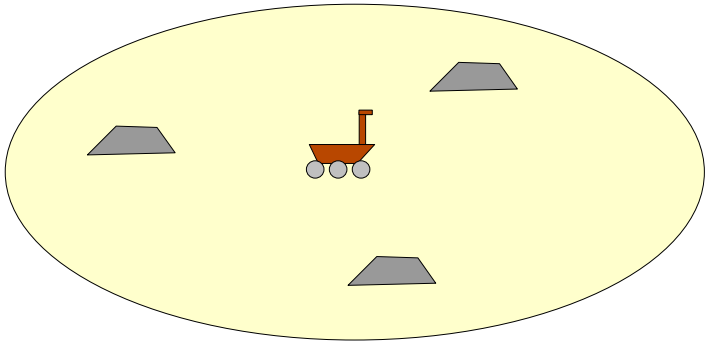

Baseline Operational Modes

directed sampling

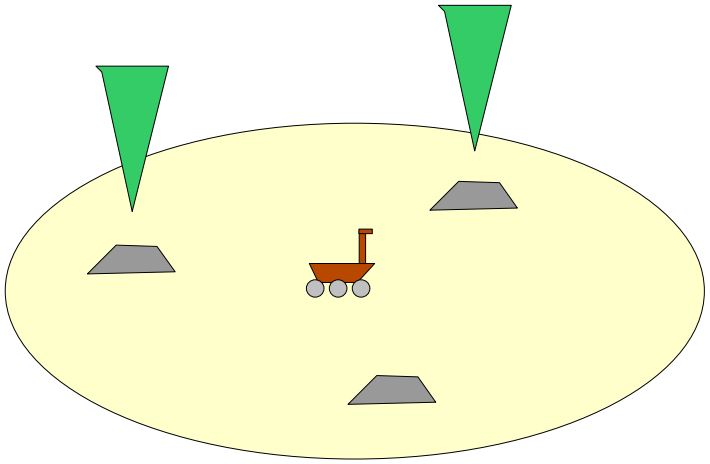

Baseline Operational Modes

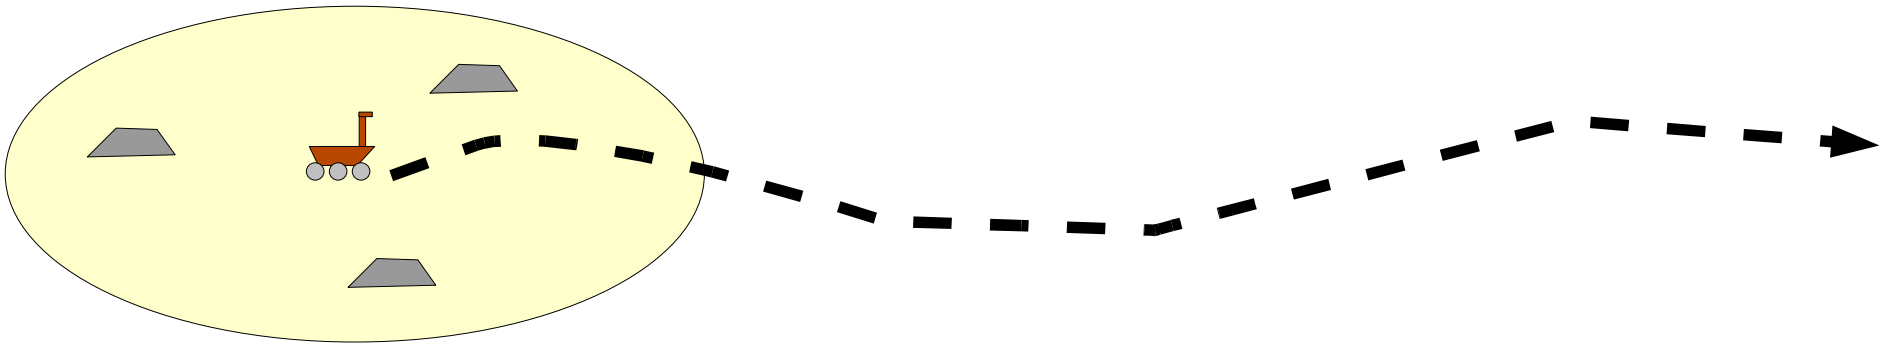

Baseline Operational Modes

periodic sampling

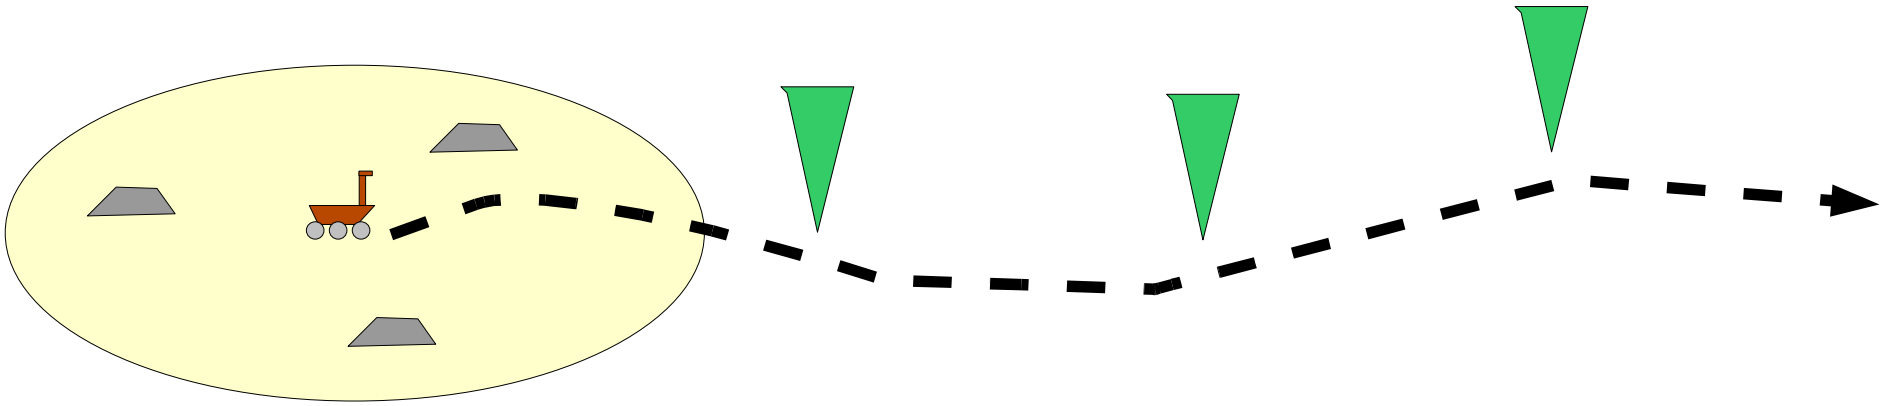

Performance Criteria: POMDP

- Well-established metrics for POMDP planning performance
- **Primarily test in simulation**
- Compare to **baseline planning techniques**
  - **Ignoring informational value (MDP)**
  - **Greedy or heuristic techniques (e.g. Q-MDP)**

Life in the Atacama

Robotic astrobiology: work under Mars-relevant operational constraints and do meaningful science on Earth.

Science Mission

The biodiversity and distribution of habitats in the Atacama is not yet measured. Where does life survive and where does it not?

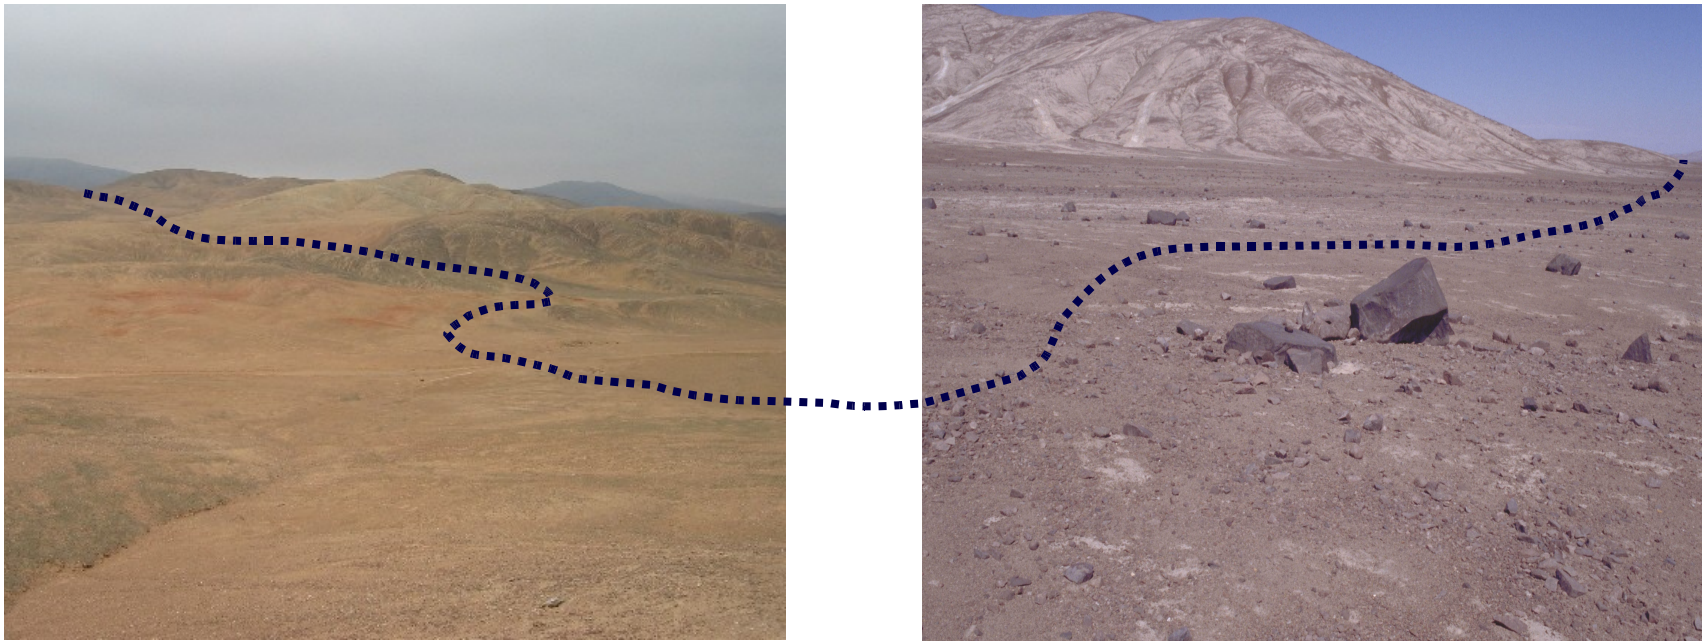

Coastal Range

Interior Desert

Planner Implementations

- **Atacama 2004**
  - Make it work: deterministic world model, heuristic search, quick development
- **Technology research**
  - A series of quickly developed planners that focus on specific POMDP techniques, tested in simulation
- **Atacama 2005**
  - Integrate lessons from earlier versions: optimize for stability, extensibility
  - POMDP technology may not be ready

Pushing POMDP Technology

- **Heuristic search**
  - Combine heuristic search with efficient value function representations
- **Factored state**
  - Do better when the state is **mostly observable**
  - Data structures that leverage **independence**
- **Continuous planning**
  - Fast replanning

Contributions

- Lessons learned from developing and field testing one of the first science autonomy systems
  - Distinguished by “keep moving” operational strategy and novel intelligent site survey operational mode
- Extensions to the POMDP state-of-art that apply both to SA and other domains

Schedule

- POMDP technology research **(ongoing)**
- Initial SA development **(Summer 04)**
- LITA Expedition 04 **(Fall 04)**
- POMDP SA model development **(Spring 05)**
- LITA Expedition 05 **(Fall 05)**
- Thesis writing **(Spring-Summer 06)**

Questions?

NASA Mars Agenda

2005 Mars Reconnaissance Orbiter
2007 Mars Scout 1: Phoenix Lander
2009 Mobile Lab
2011 Mars Scout 2
2015 Mars Sample Return

...

~2025 Human Landing

OASIS Project

- Broad agreement on goals and approach
  - For instance, ideas about preferences (key signature, anomaly, etc.)
- Distinctions of the proposed work
  - Operational strategy: keep moving
  - Kilometer-scale traverses, multiple sites per day
  - Novel intelligent site survey concept

LITA Science Autonomy 2004

- September 2004 Expedition
  - Early versions of all modules present
  - Test science observer/planner closed loop in very simple situations
  - Teleoperate rover to simulate SA system
  - Gain experience with architecture and domain modeling
  - Get feedback from science team

LITA Science Autonomy 2005

- September 2005 Expedition
  - Test full-up SA operational modes
  - Learn about integration into overall operations strategy
  - Characterize performance

Sensing Uncertainty

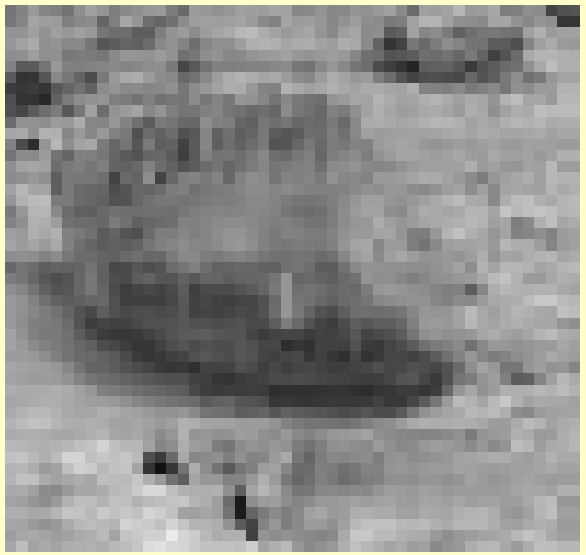

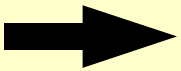

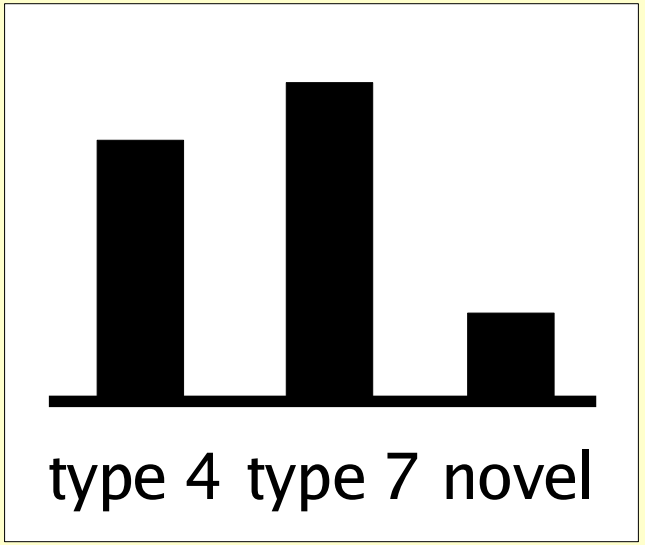

Sensing Uncertainty

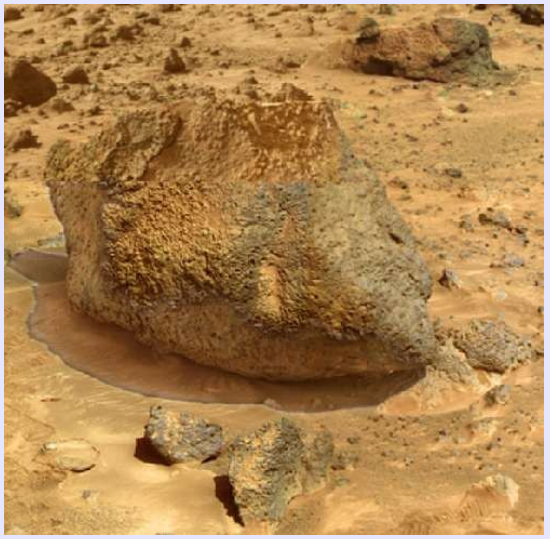

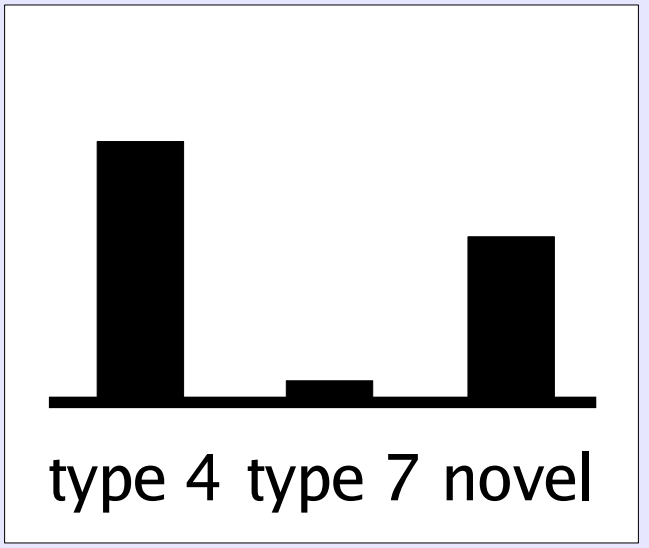

Sensing Uncertainty

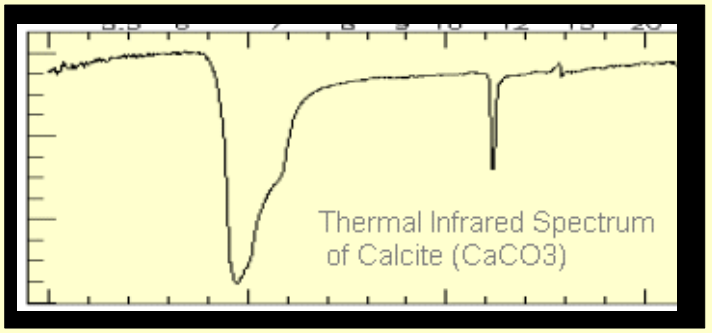

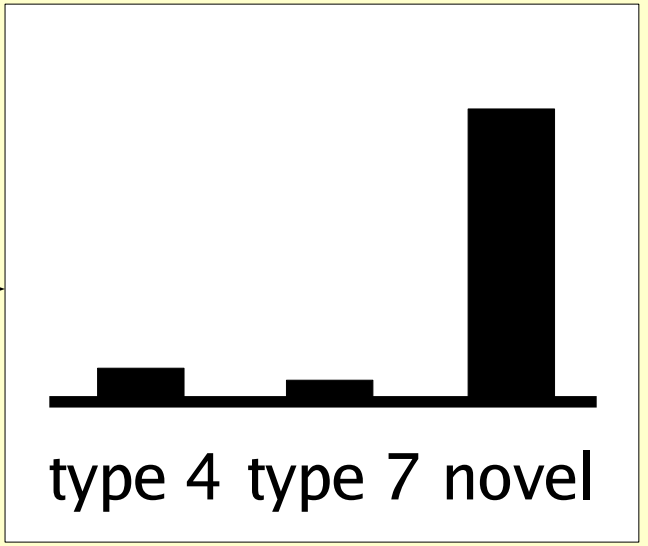

Sensing Uncertainty

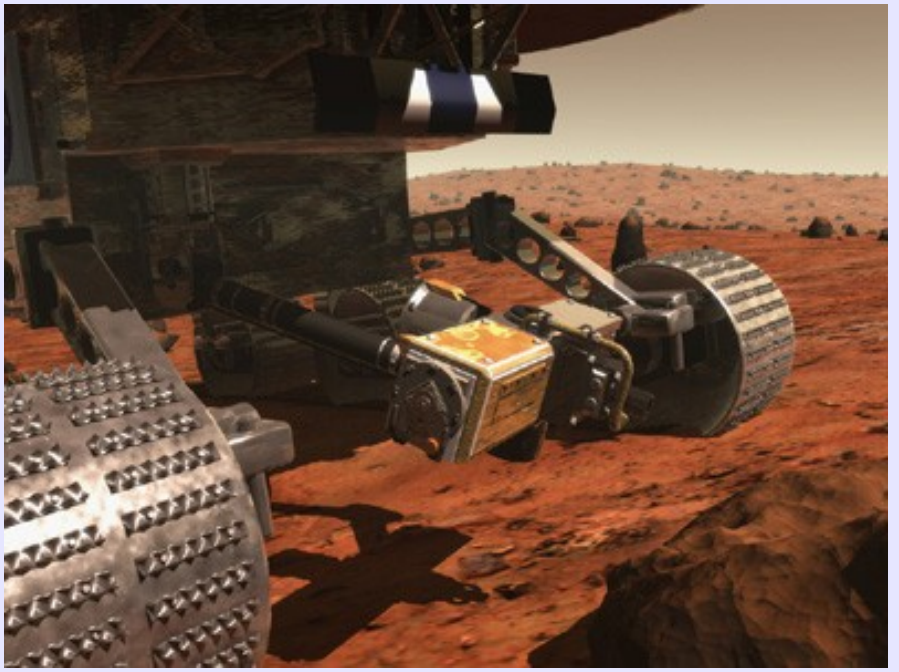

Science on the Fly

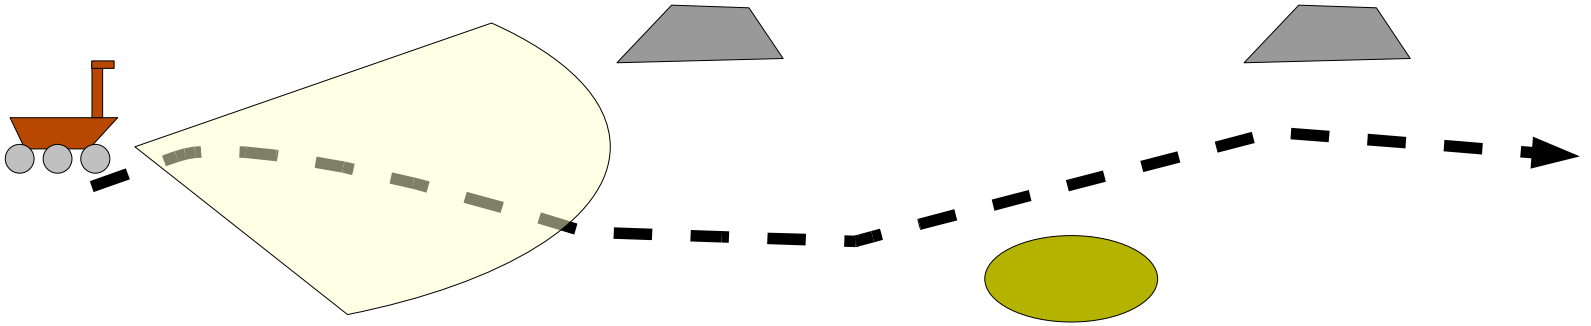

Science on the Fly

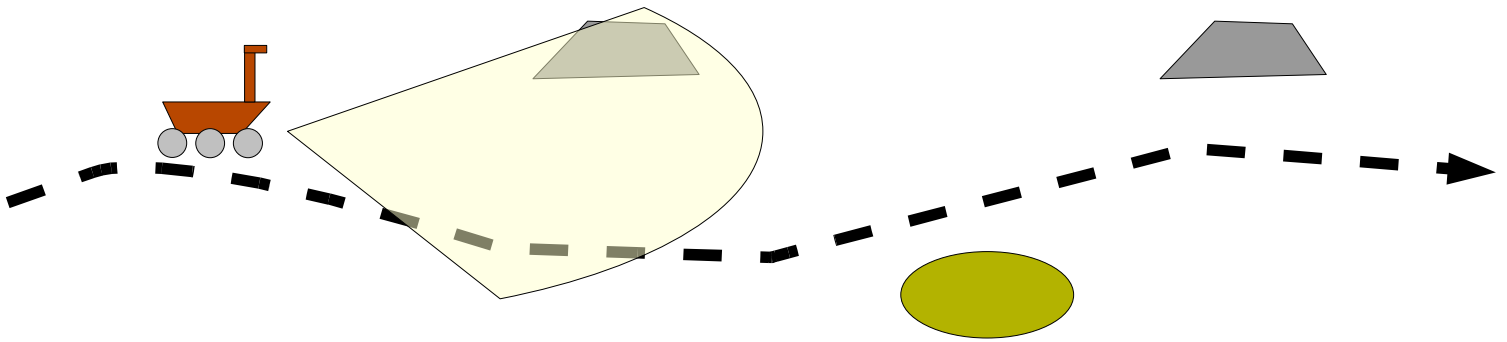

Science on the Fly

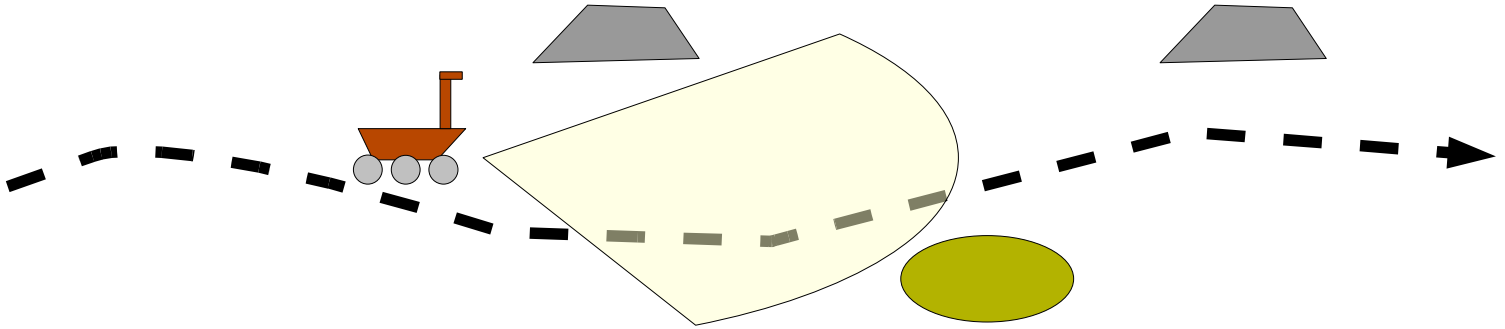

Science on the Fly

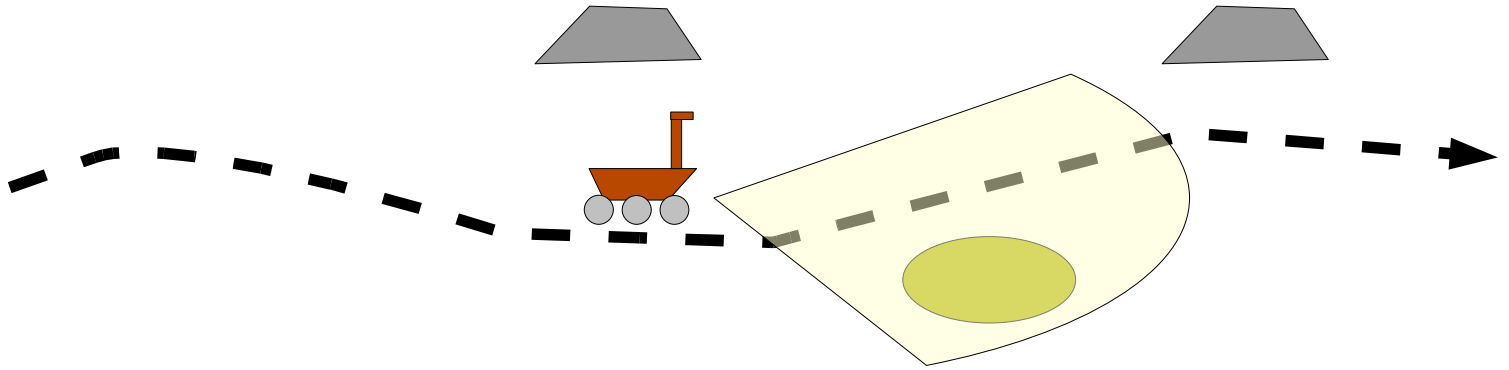

Science on the Fly

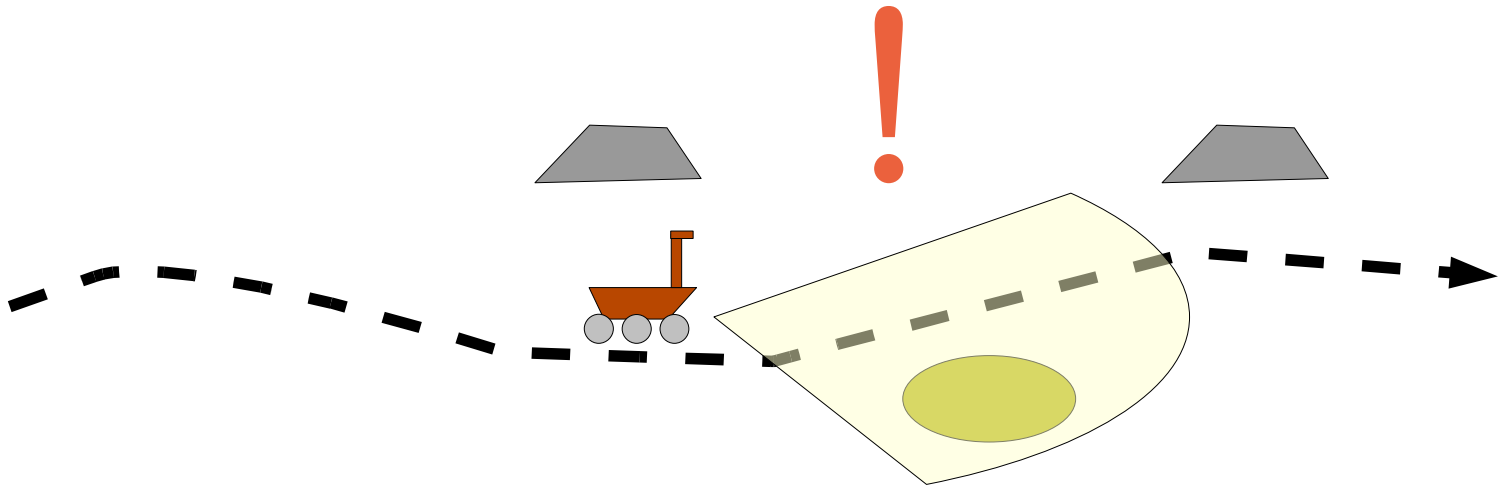

Science on the Fly

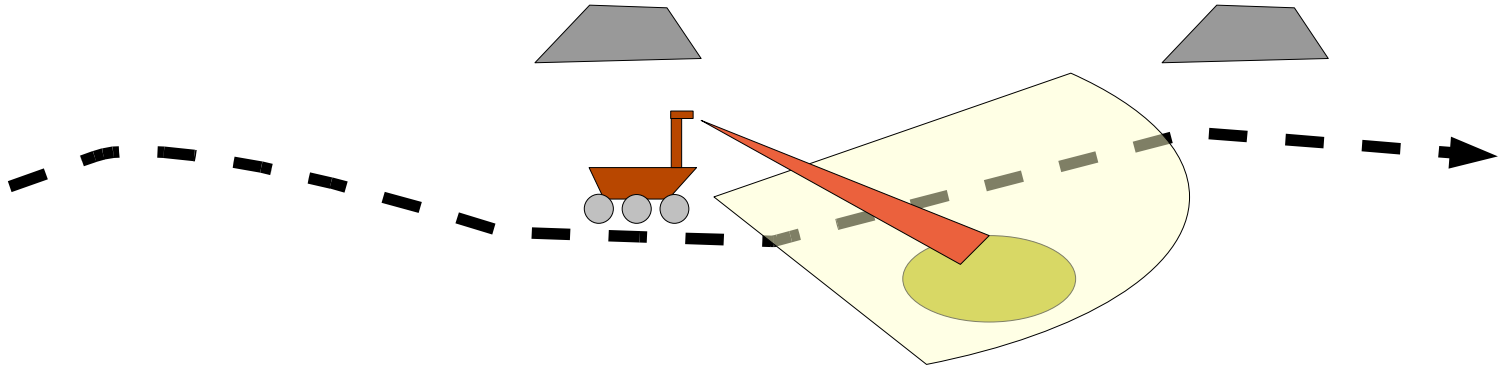

Science on the Fly

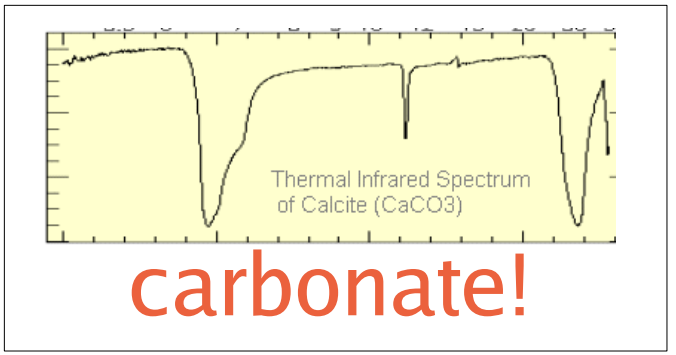

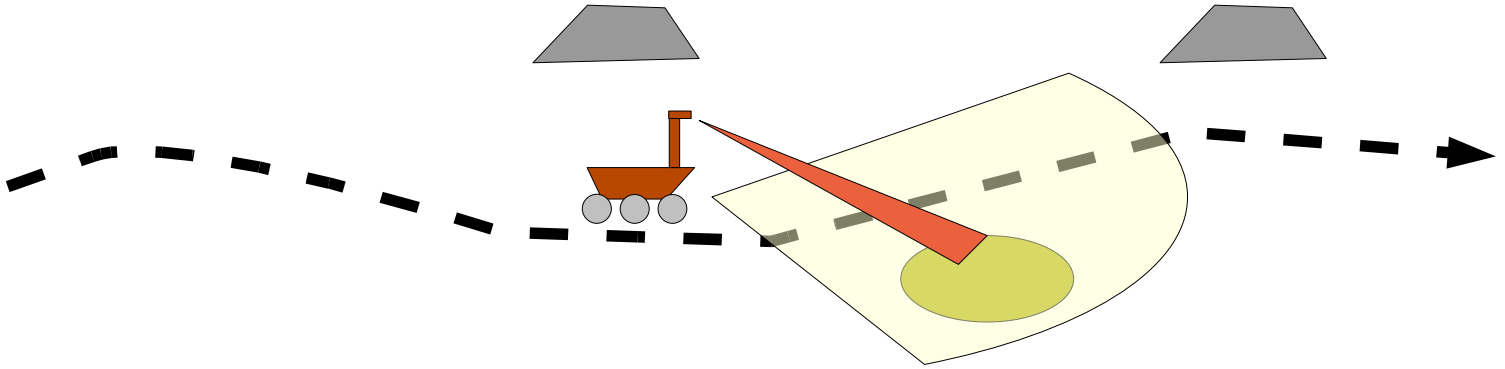

Science on the Fly

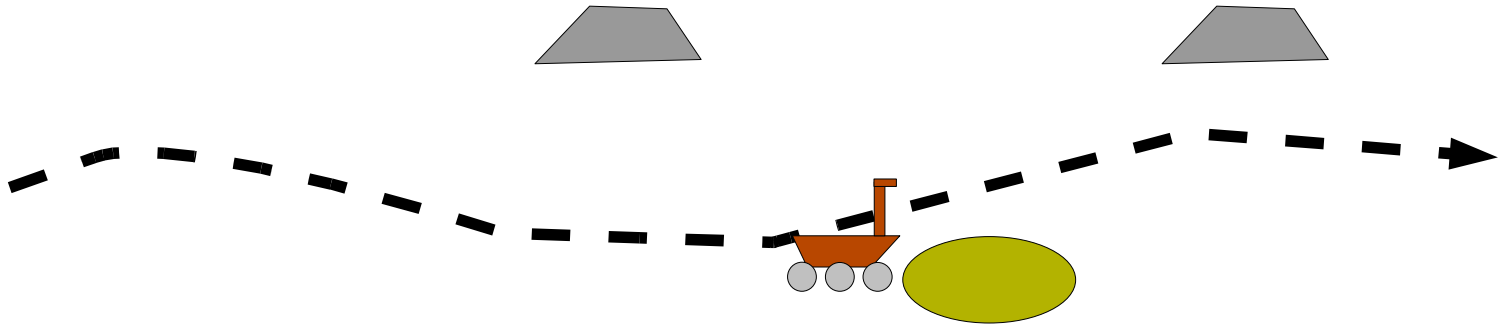

Science on the Fly

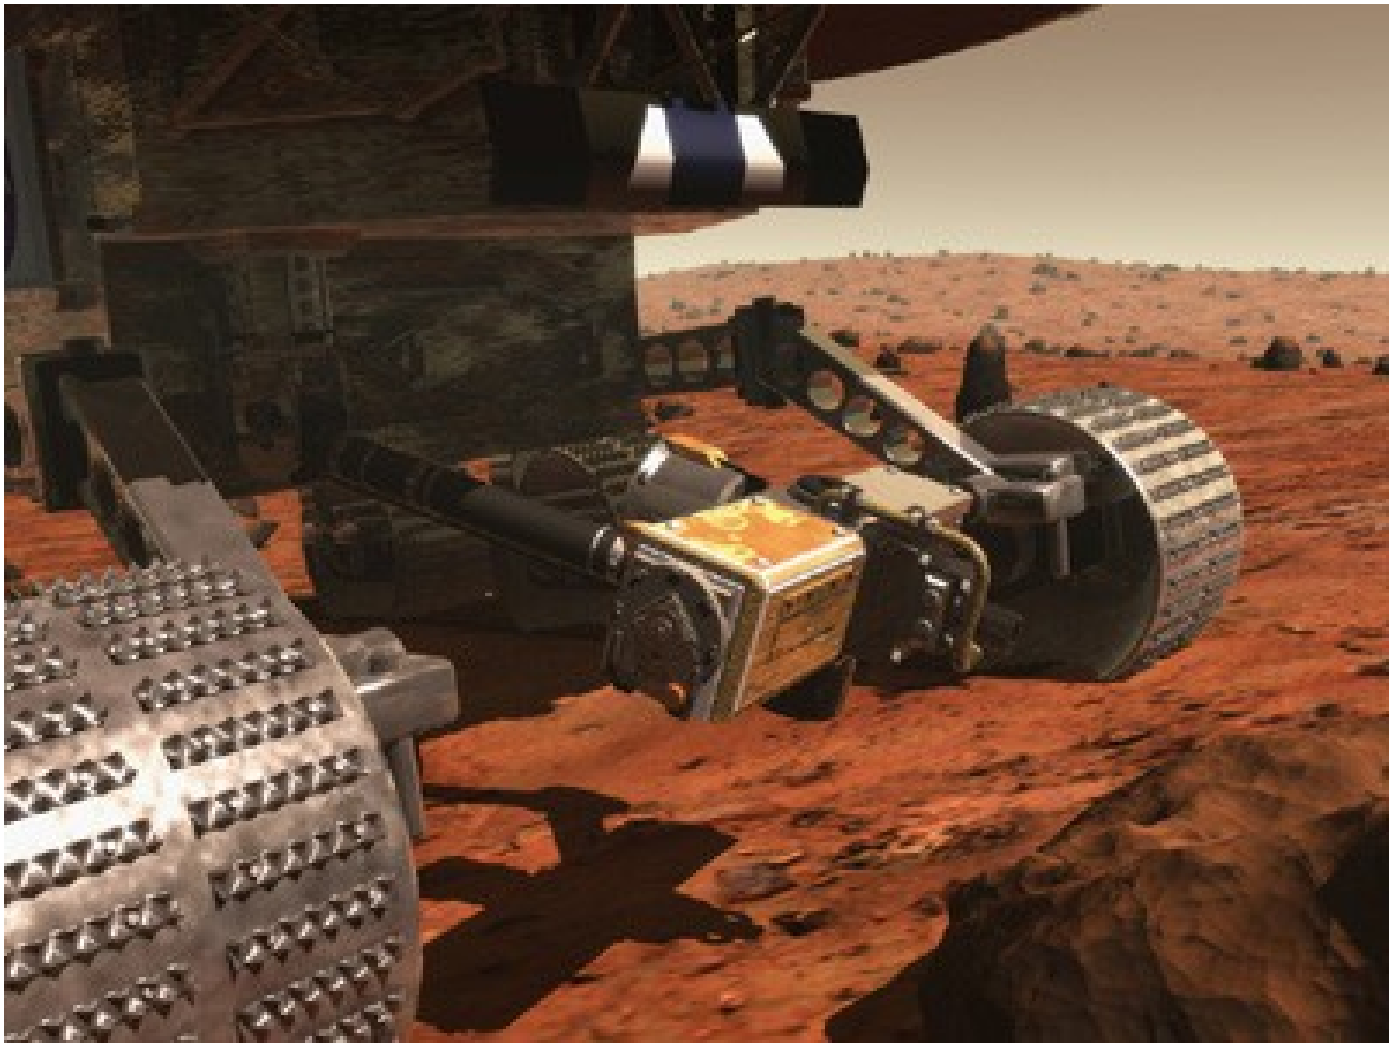

Science on the Fly

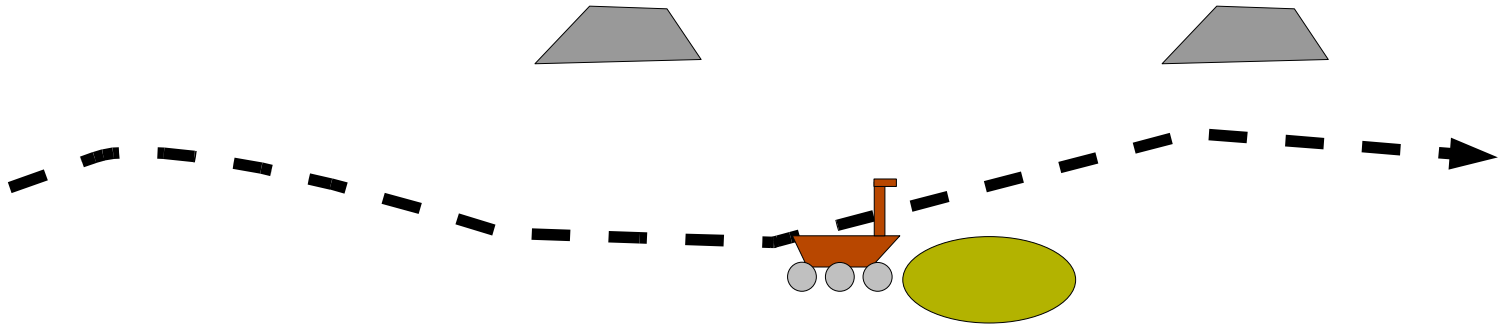

Science on the Fly

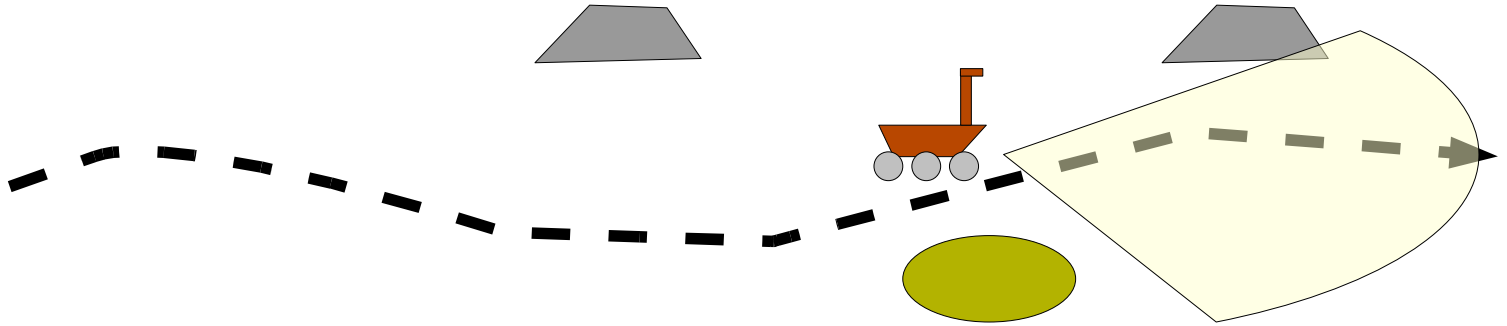

Science on the Fly