New Perspectives on the Piano Movers' Problem

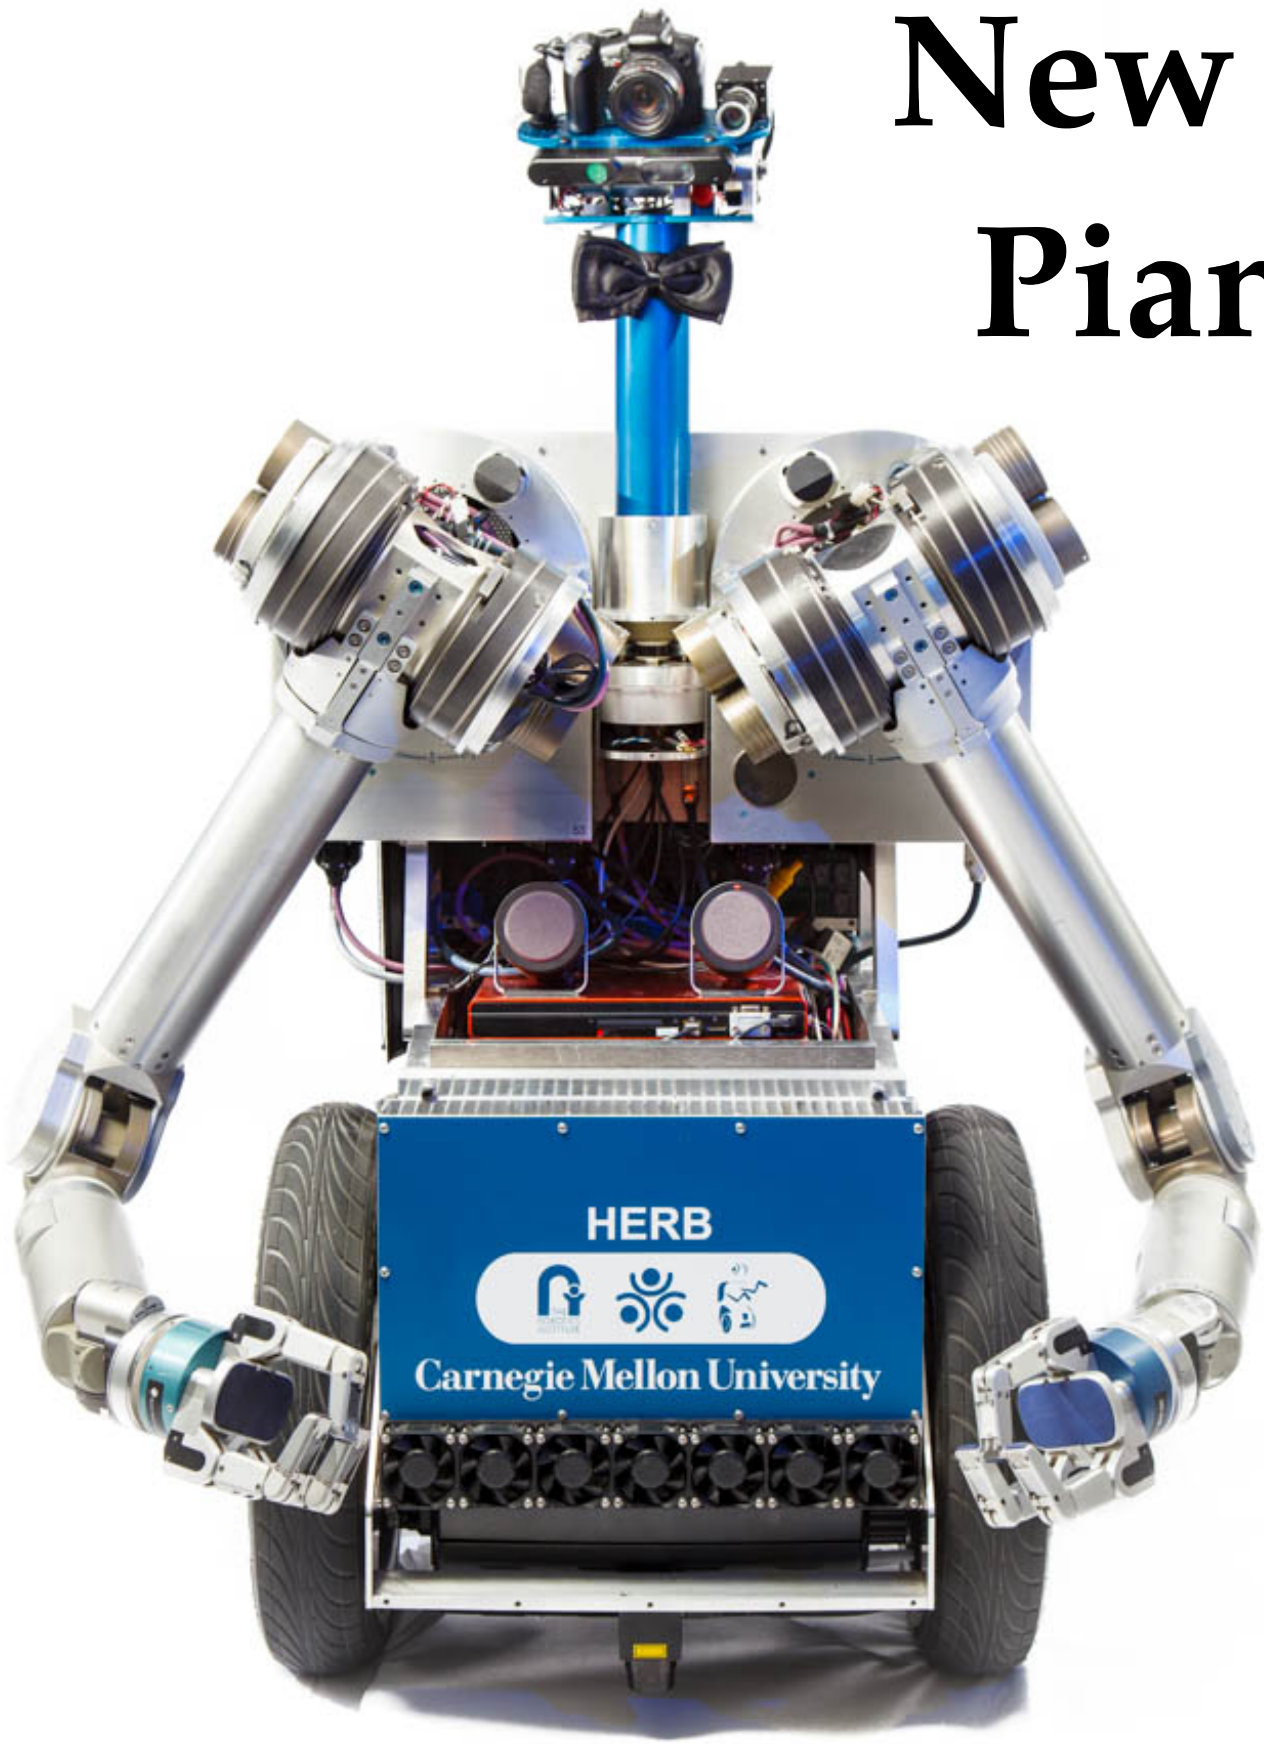

Siddhartha Srinivasa
Boeing Endowed Professor
Personal Robotics Lab
University of Washington

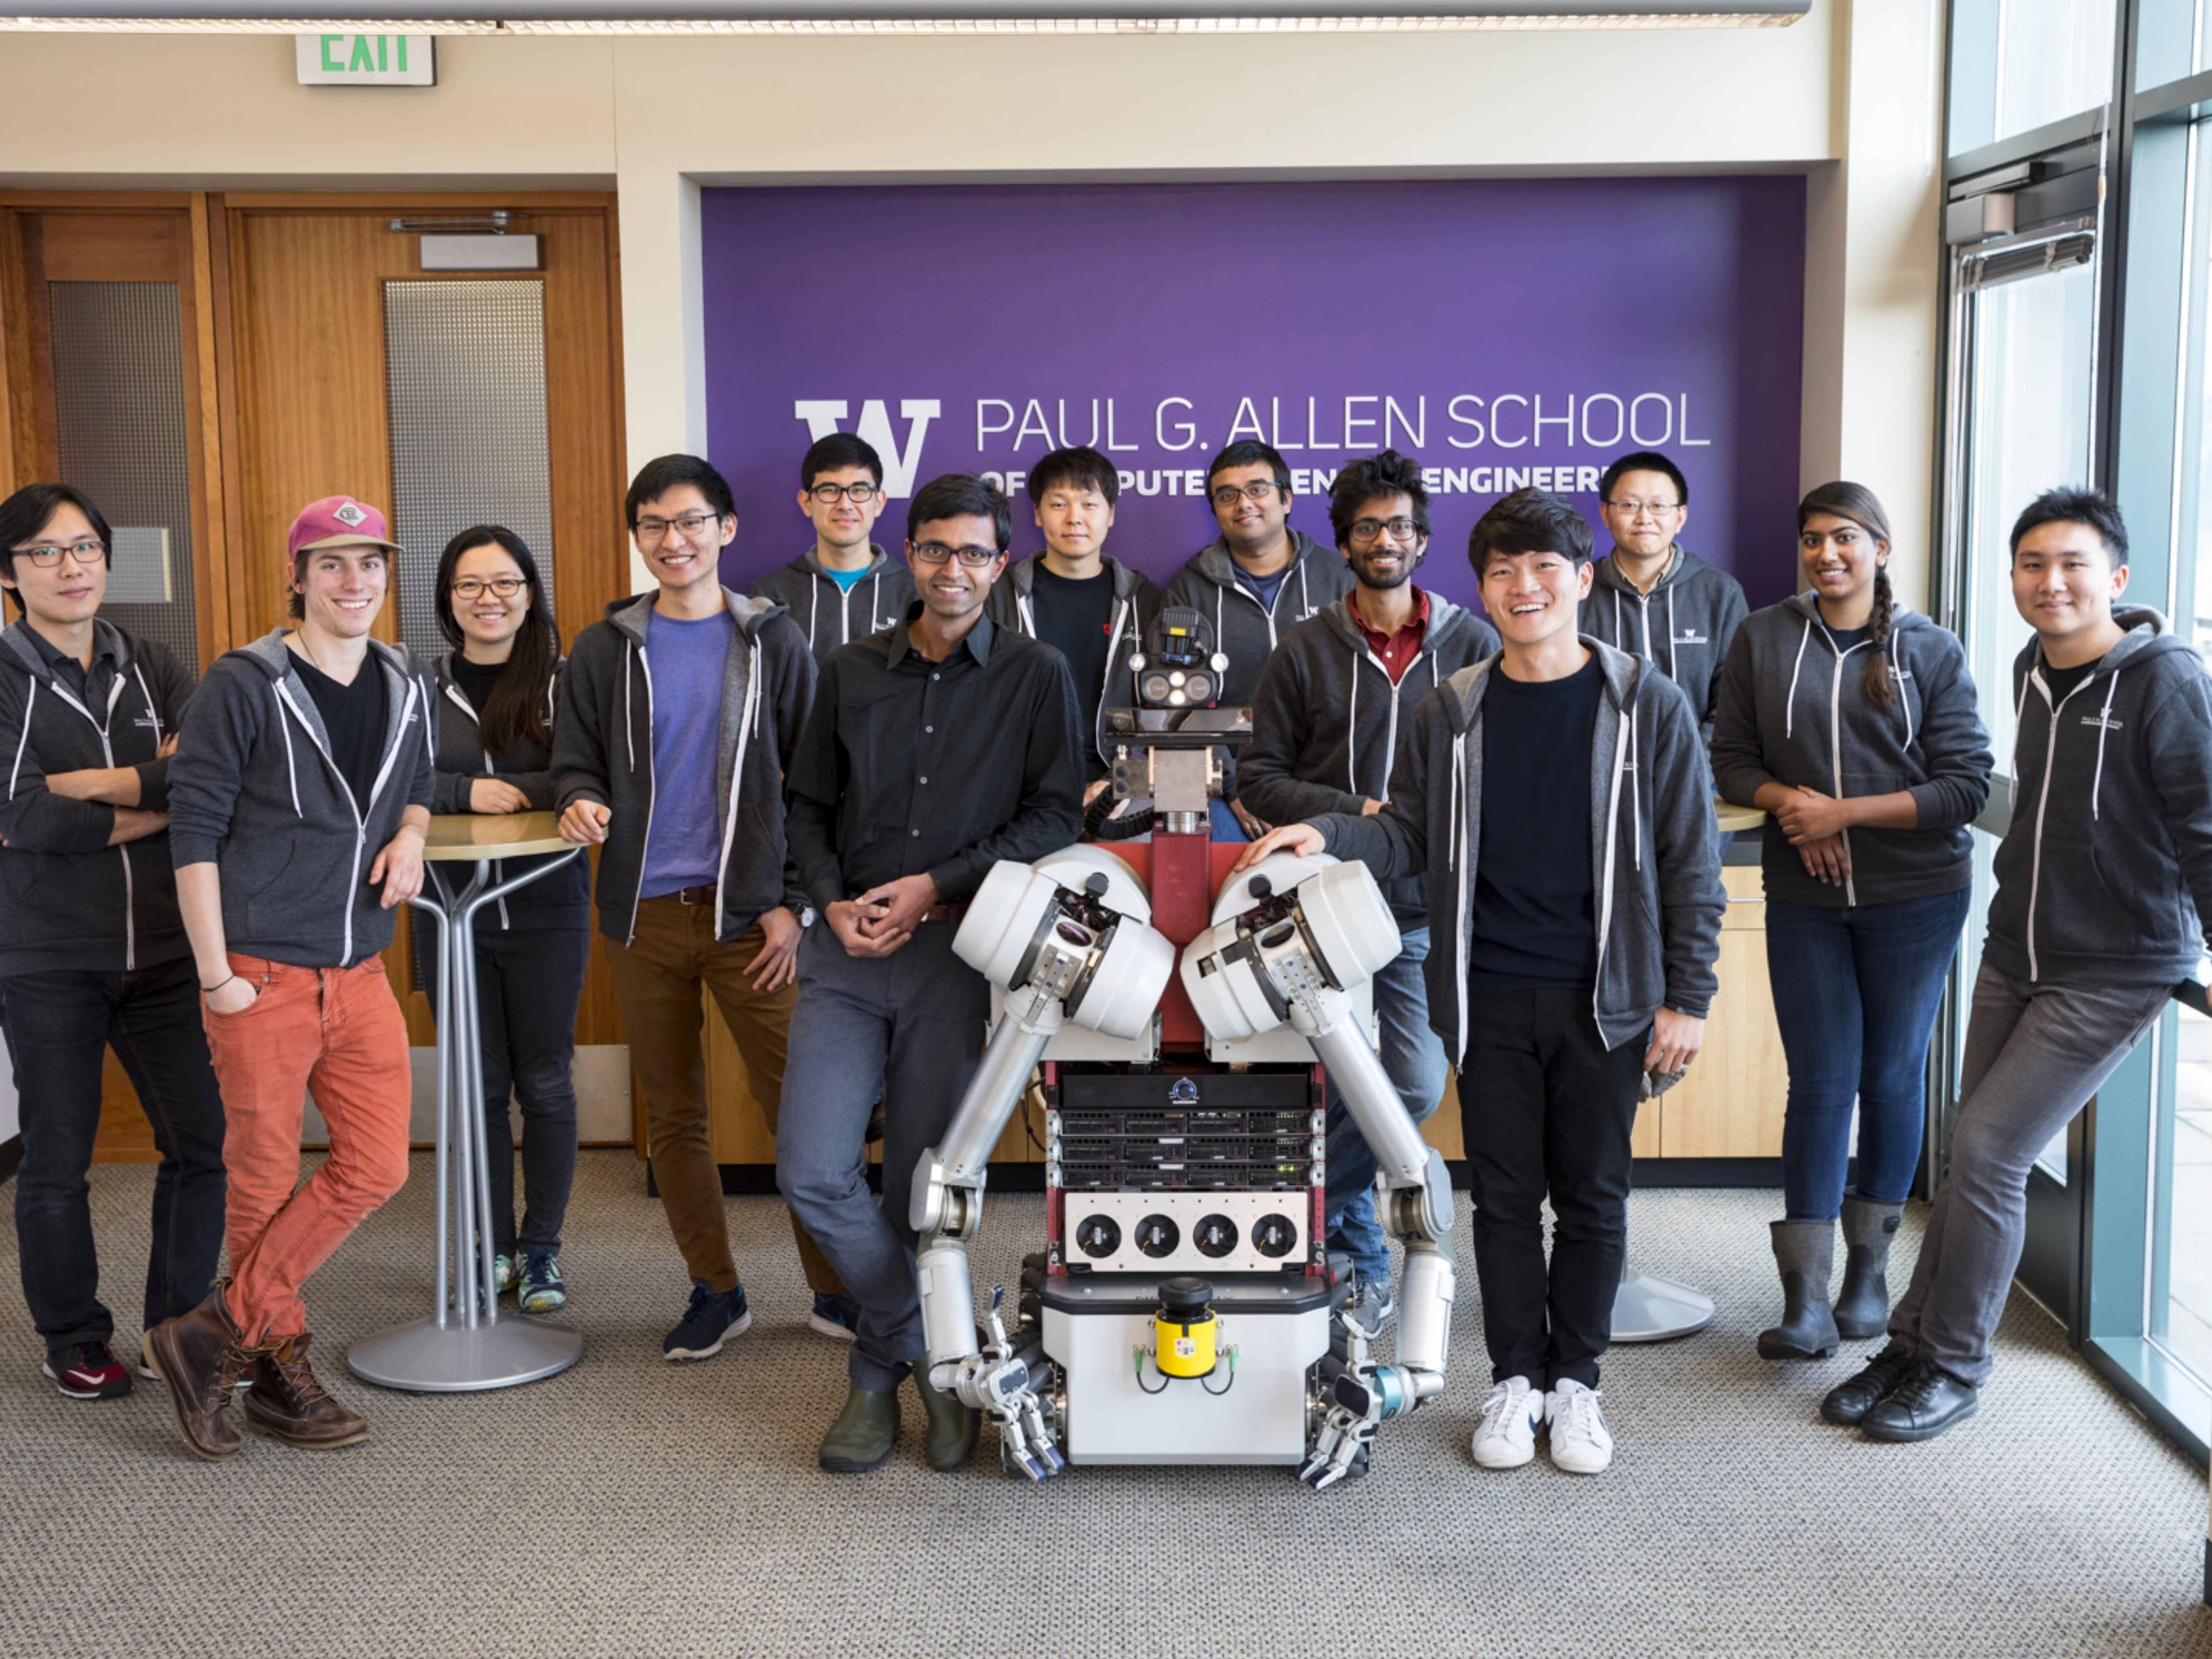
EXIT
PAUL G. ALLEN SCHOOL

Manipulation

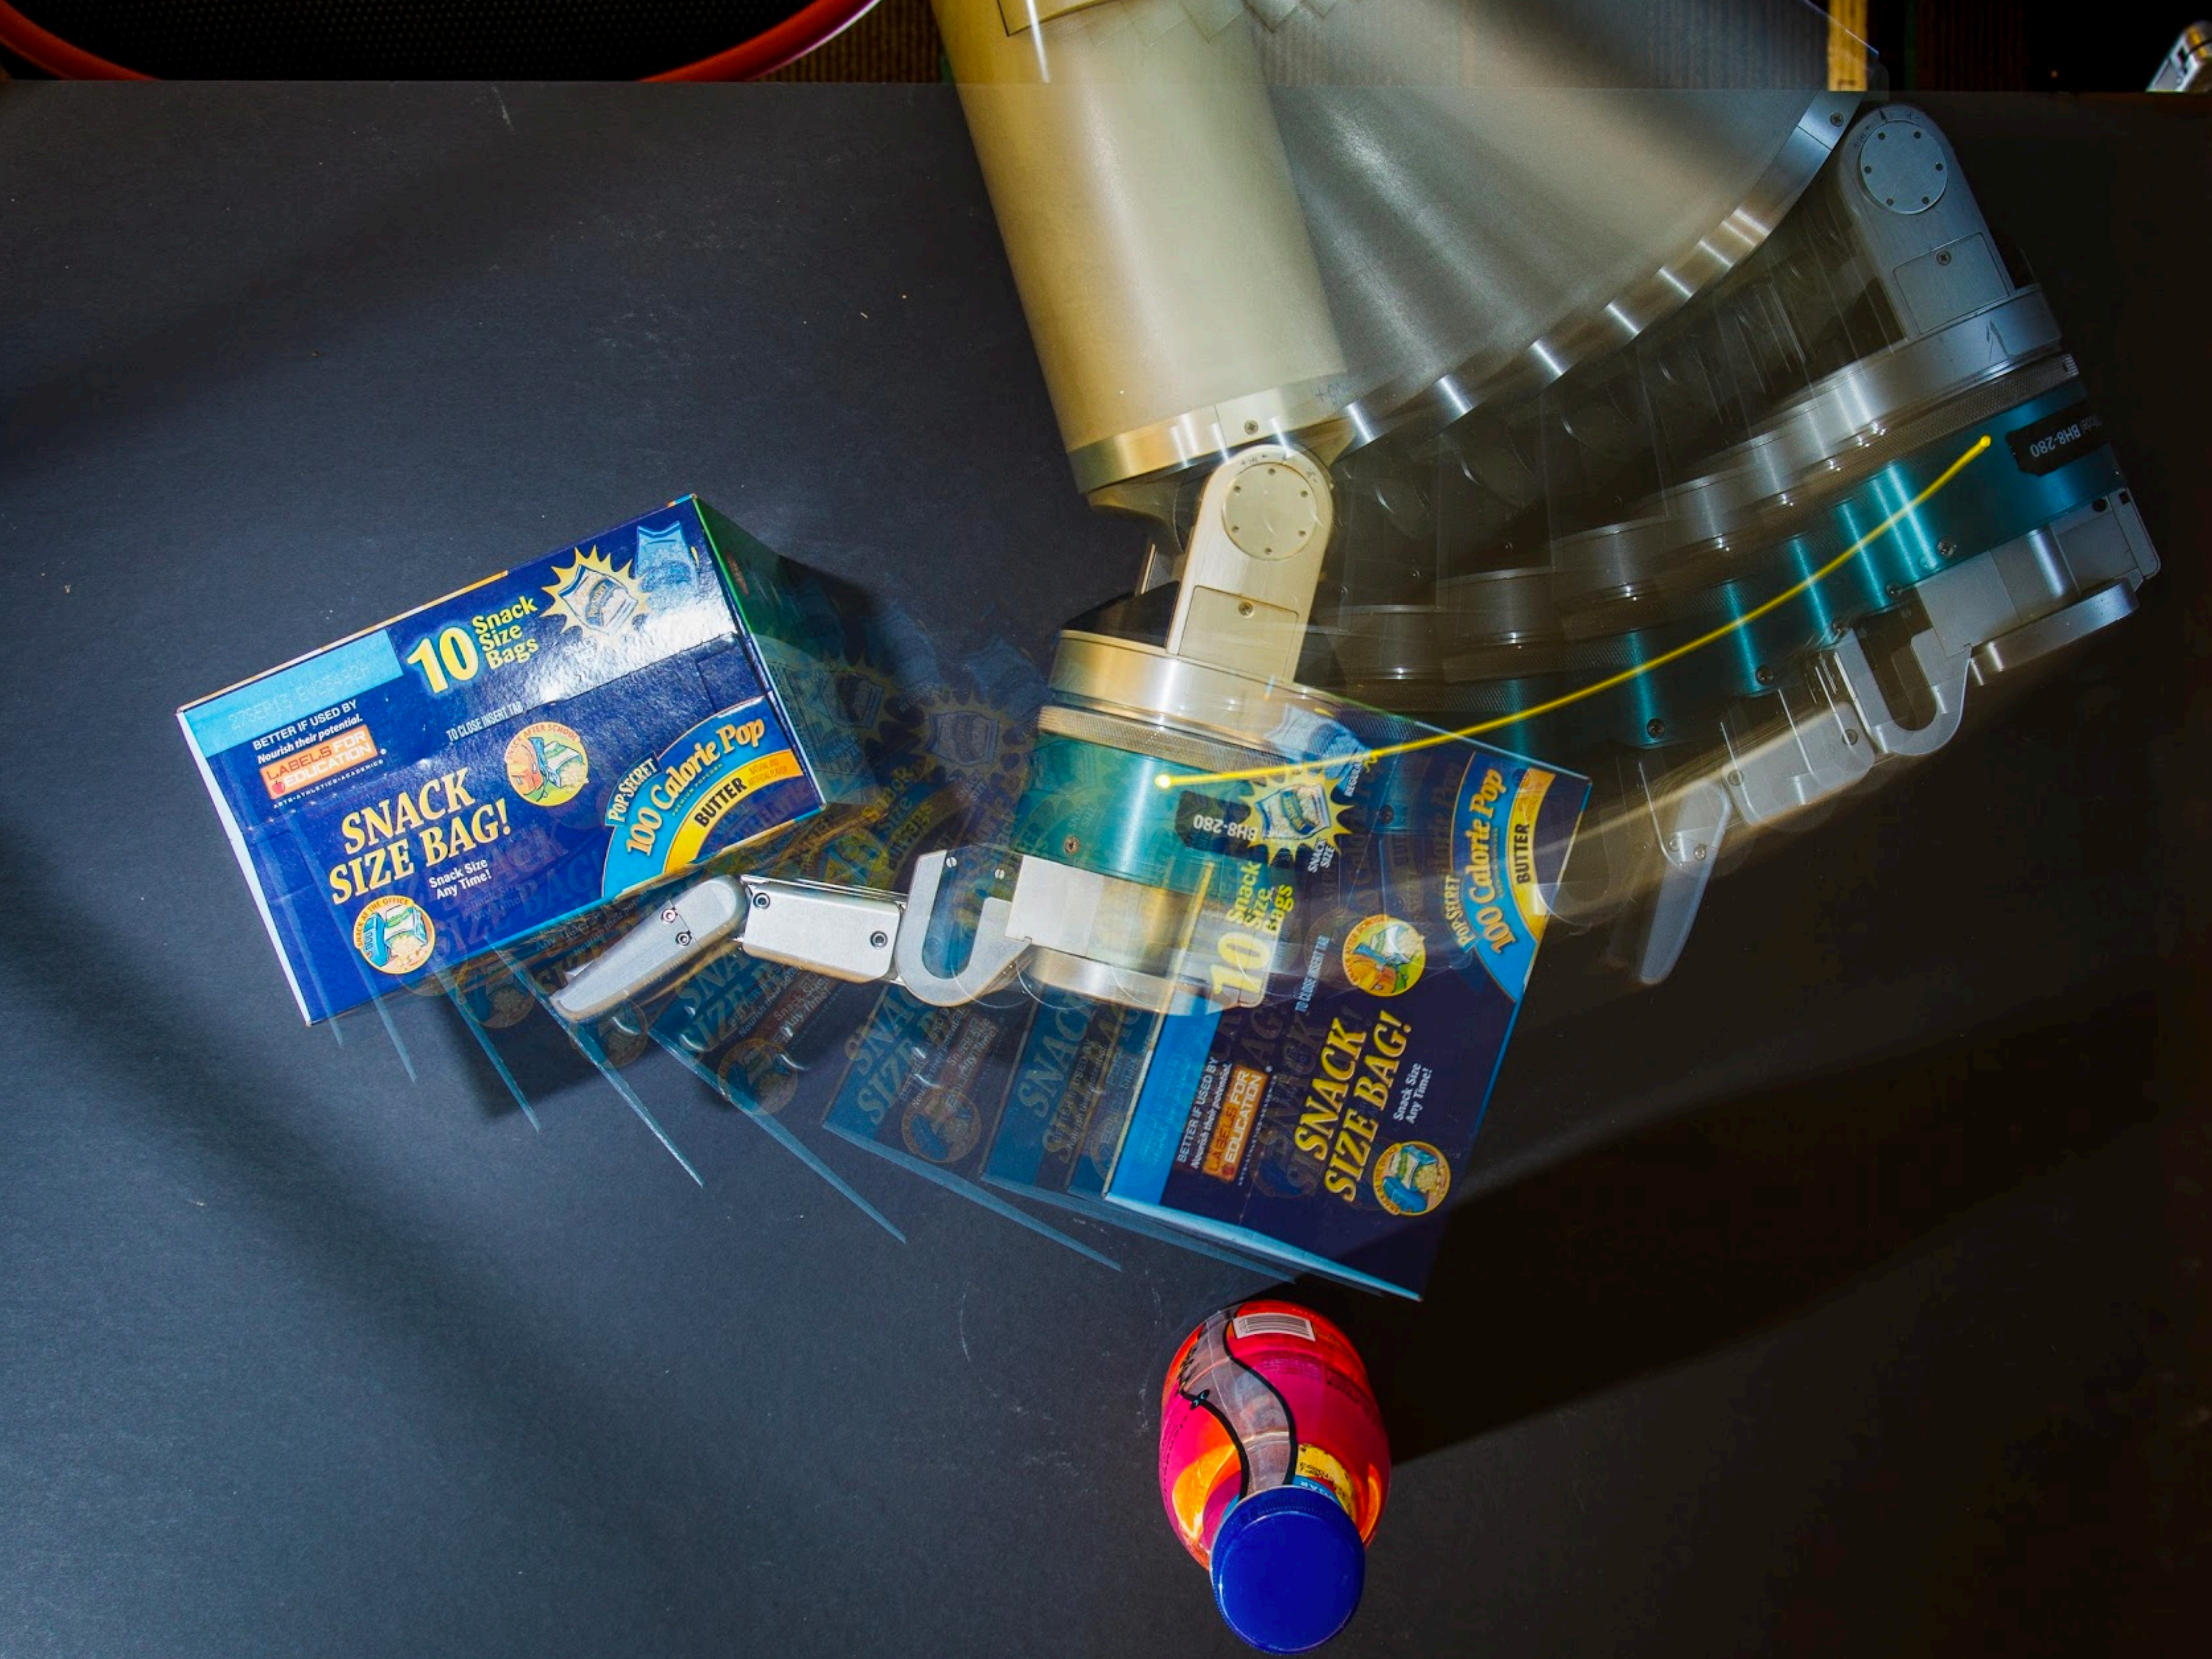
10 Snack Size Bags
BETTER IF USED BY
Nourish their potential.
LABELS FOR EDUCATION
TO CLOSE INSERT TAB
POP SECRET
100 Calorie Pop
BUTTER
SNACK SIZE BAG!
Snack Size Any Time!
BH8-280

Personal Robotics Lab
Carnegie Mellon University

Motion Planning is a technology

Formalizing the Core

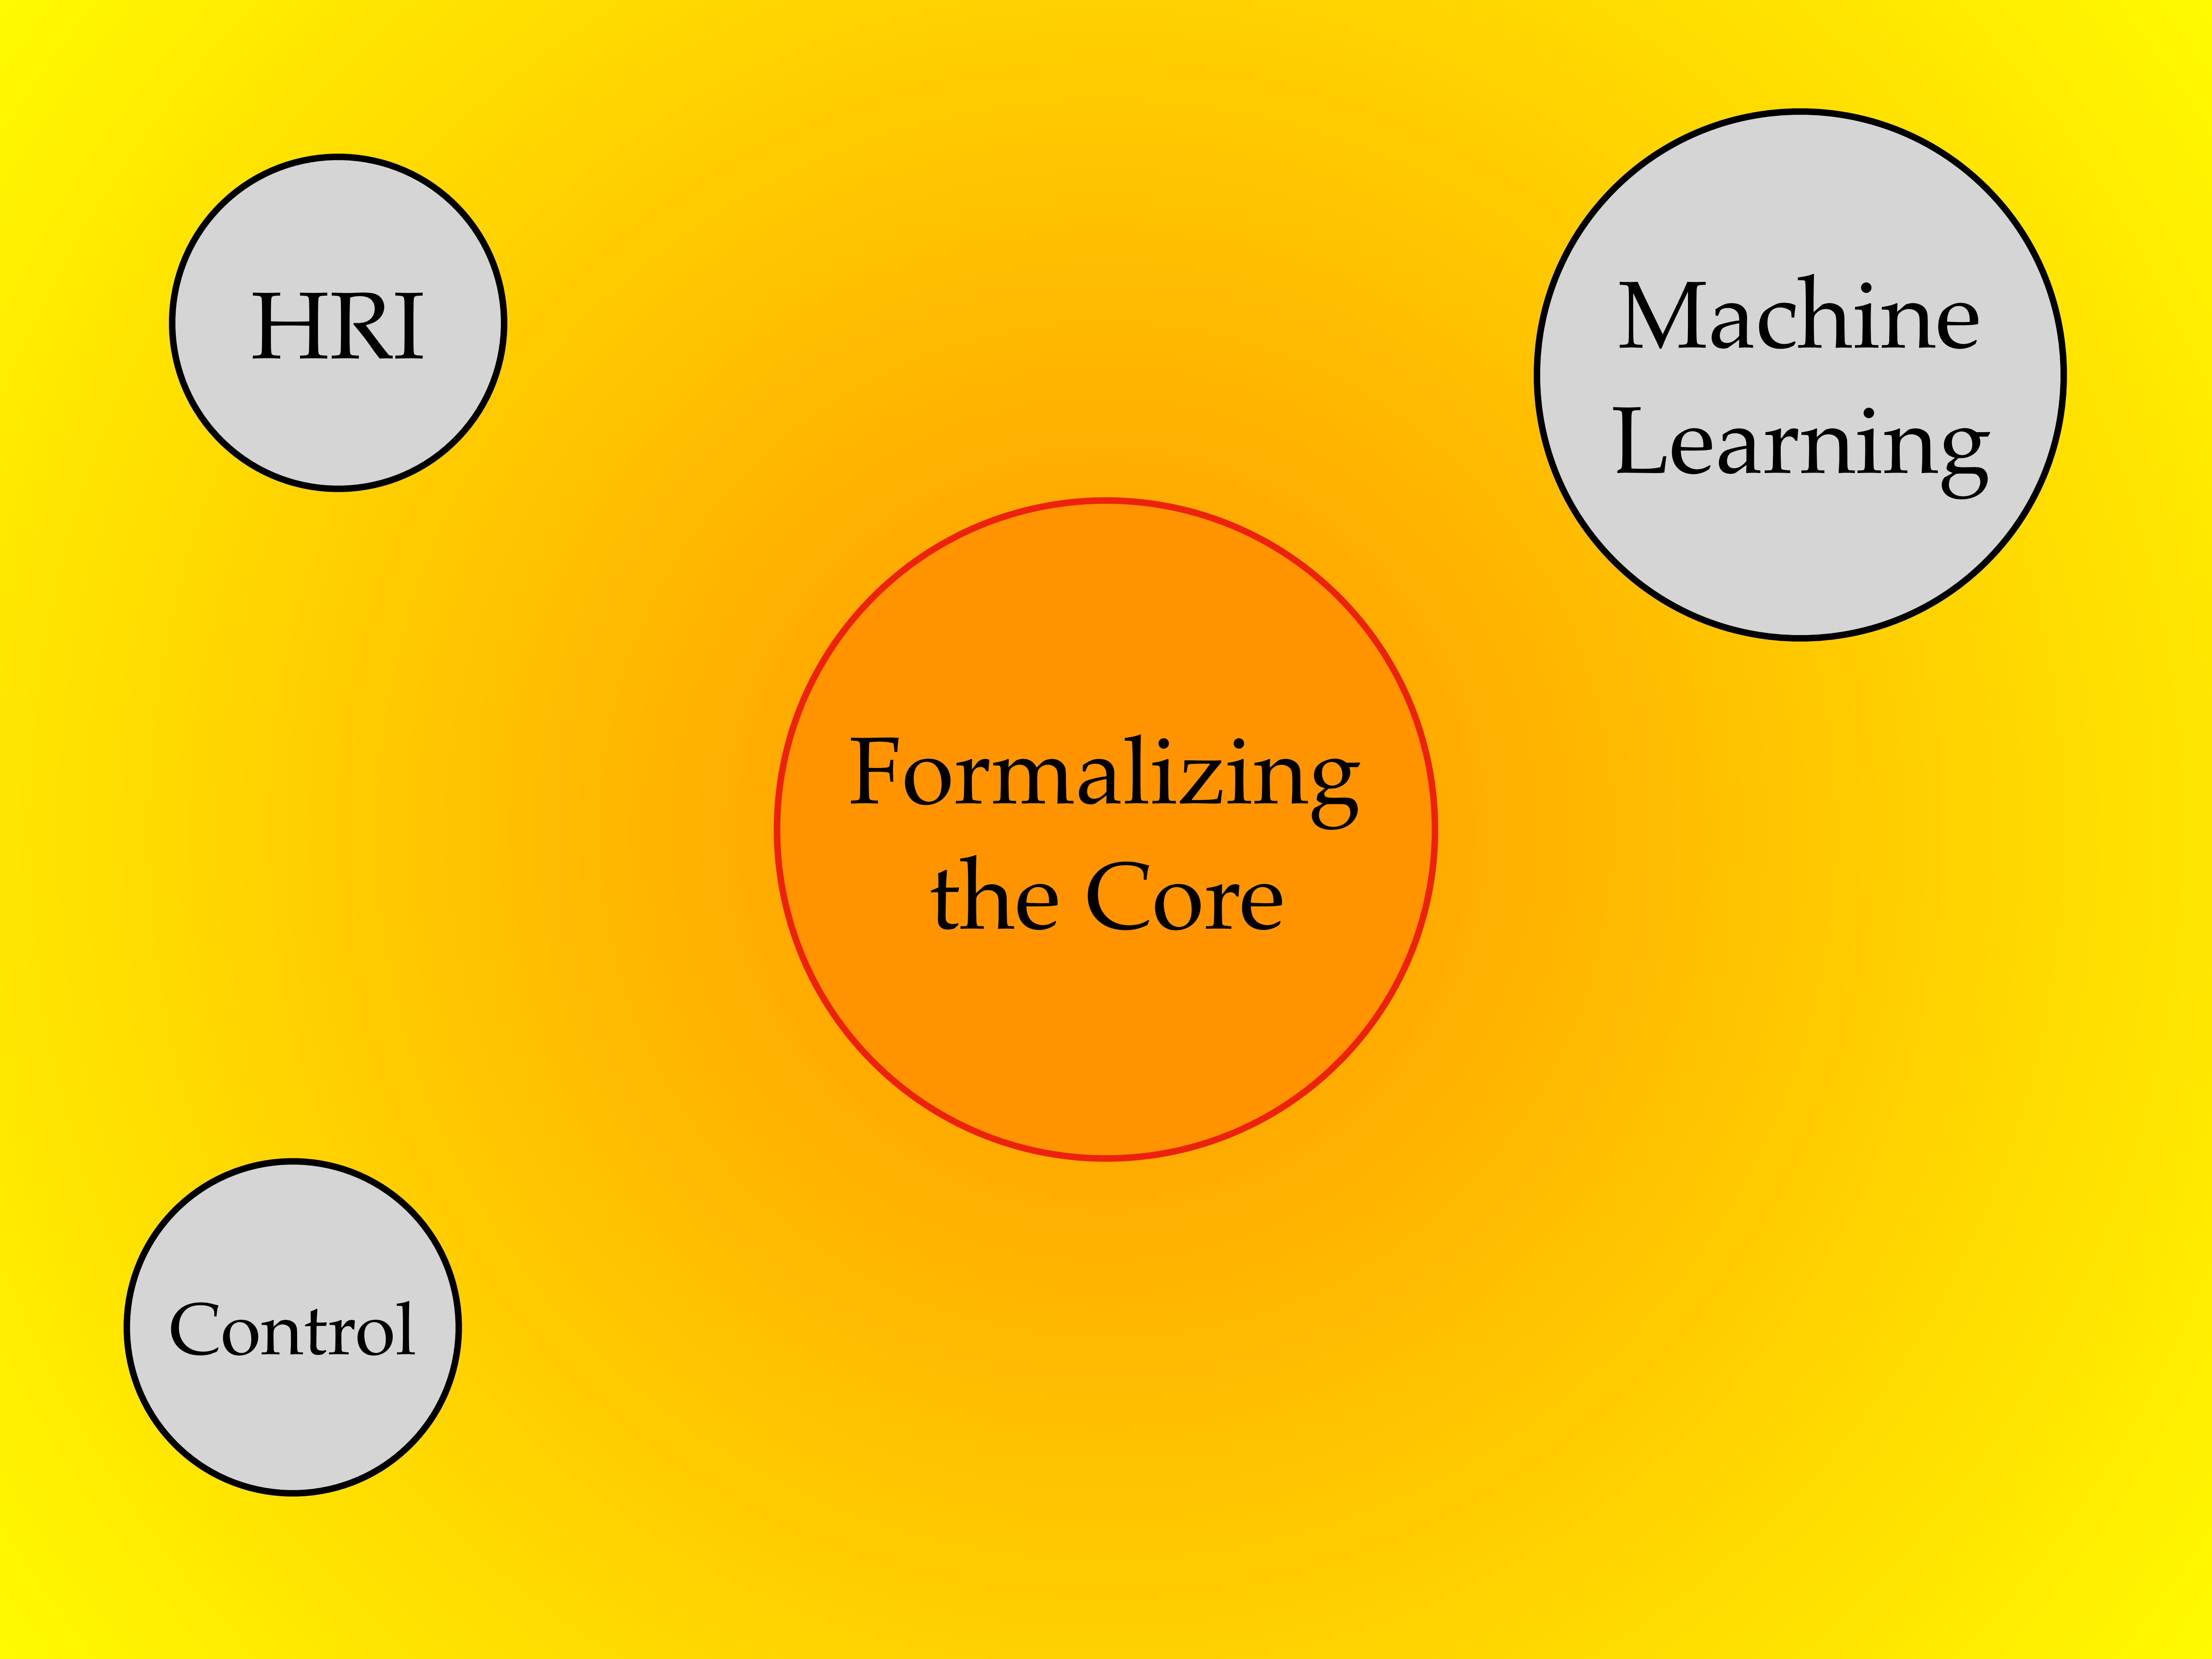
HRI
Machine Learning
Formalizing the Core
Control

The Piano Movers' Problem

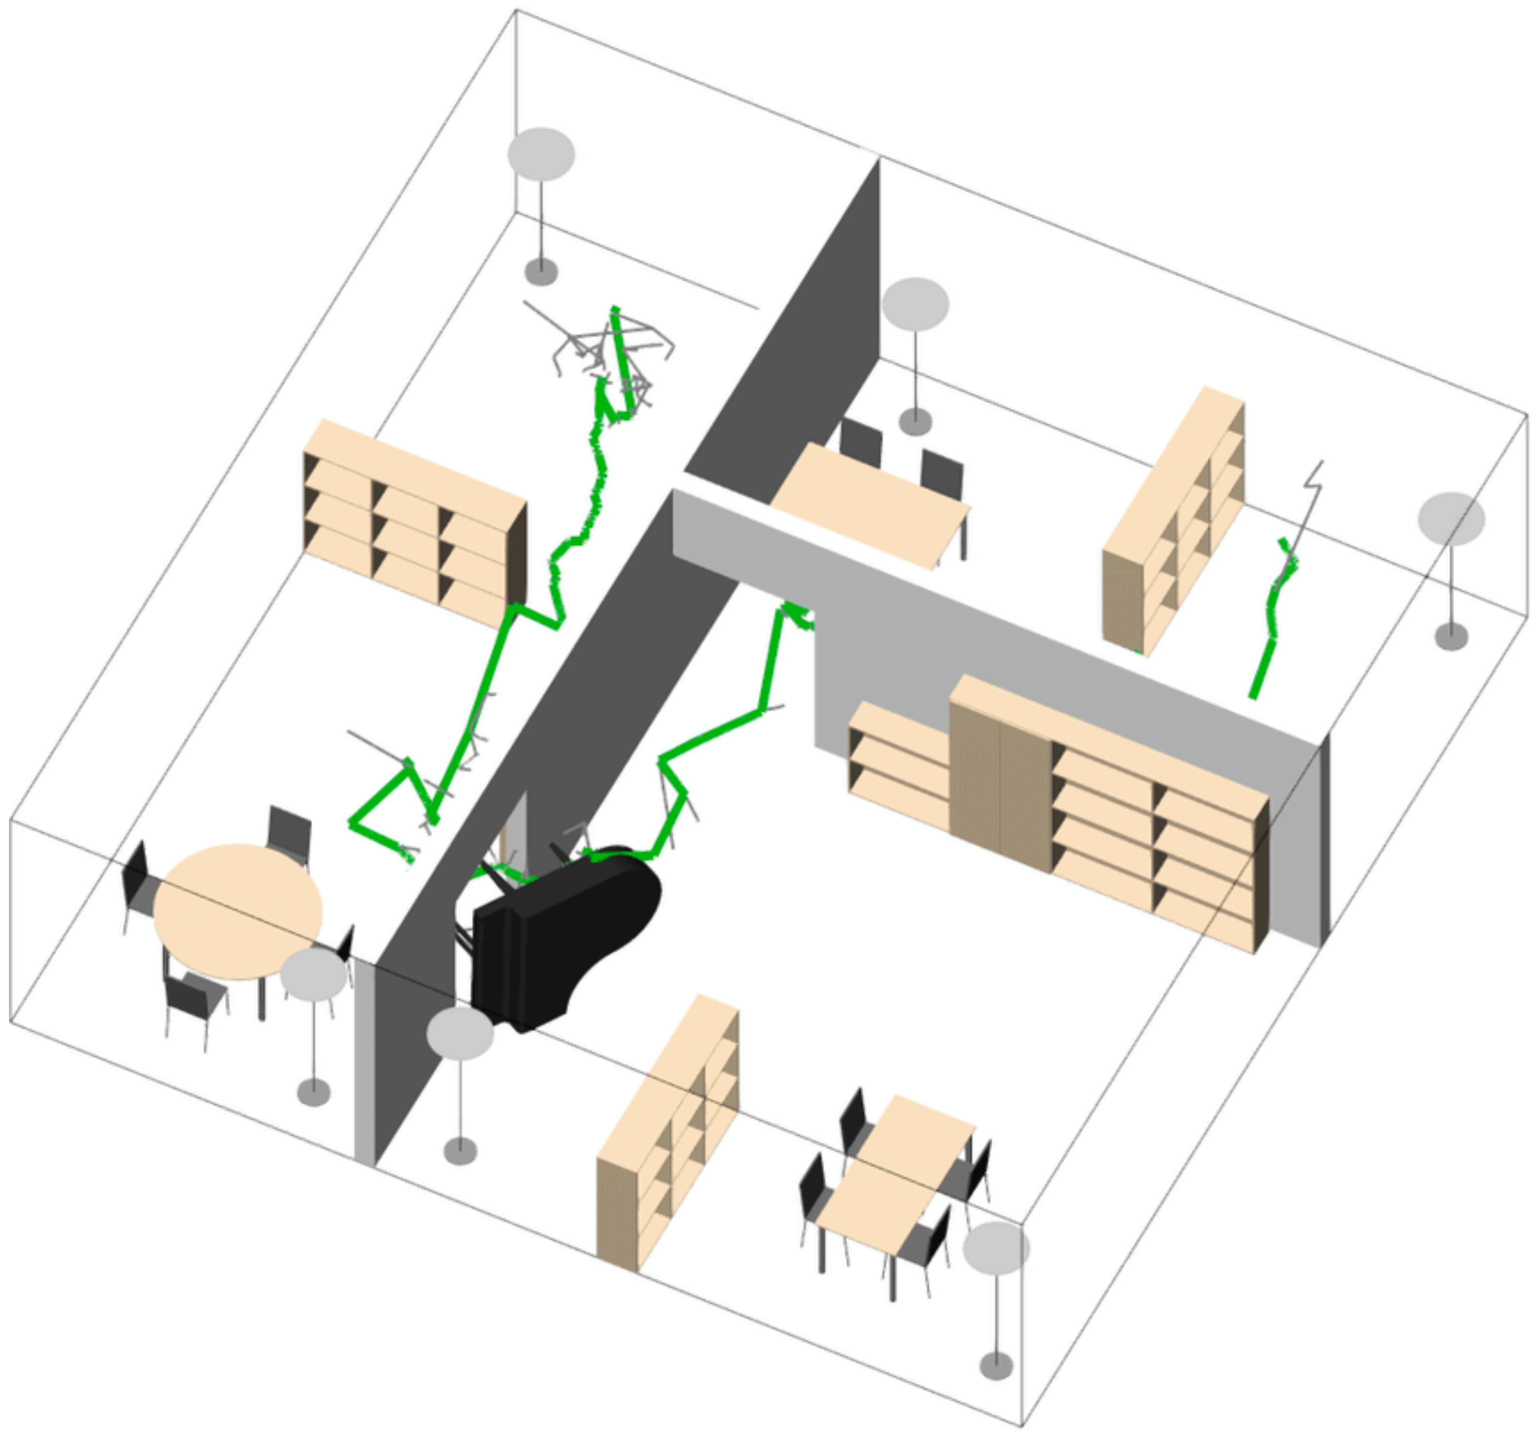

Roadmaps

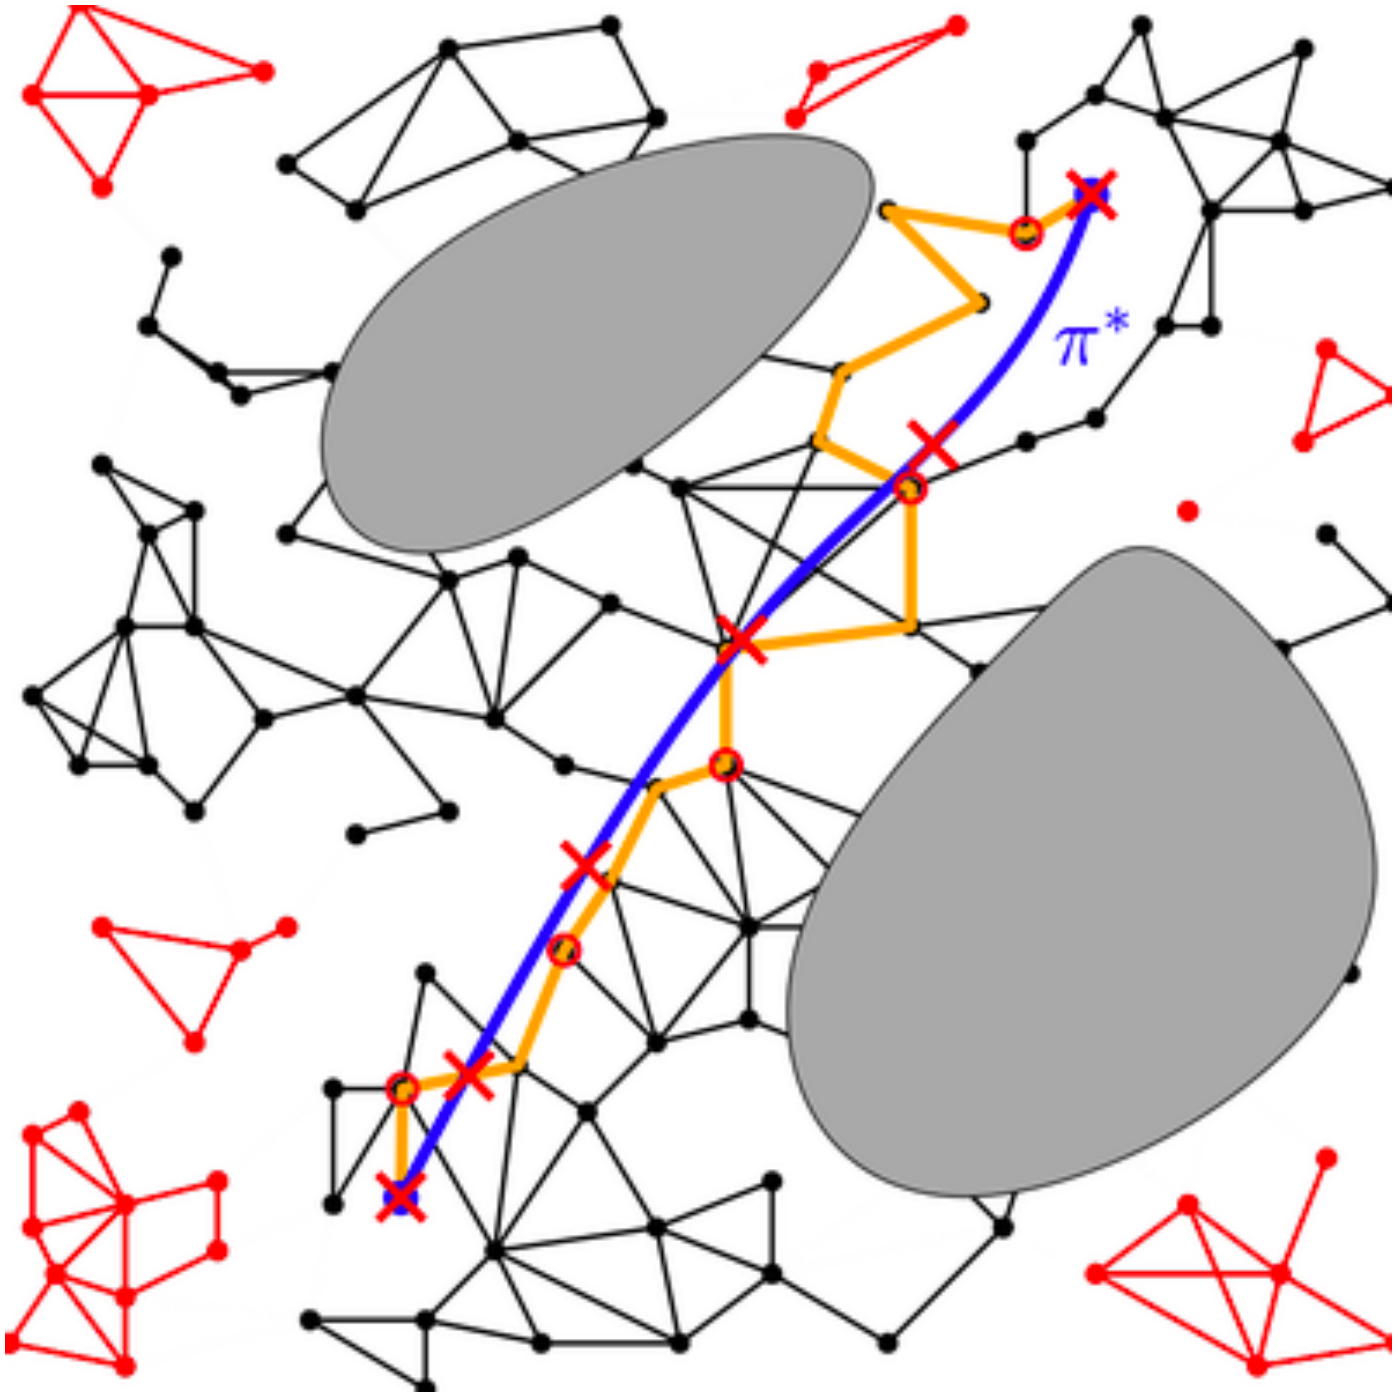

A* Search is Optimal …

Exploits Optimal Substructure
Best-first Search over Vertices

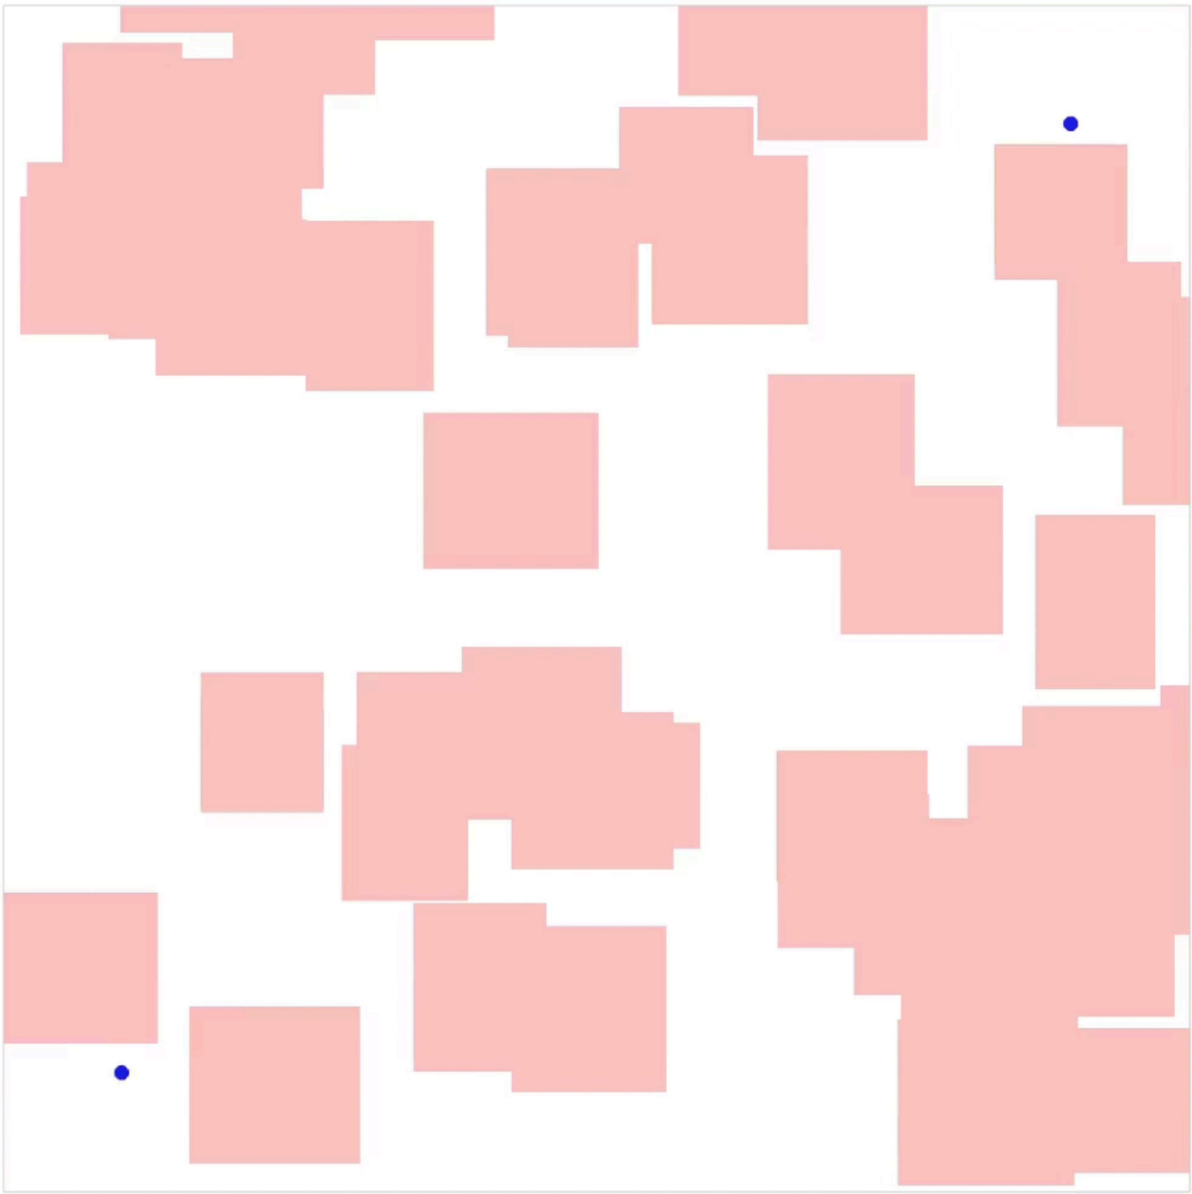

A* Search is Optimal …

Expands the Fewest Number of Vertices

But is this what we *really* want in Motion Planning?

Edge Evaluation Dominates Planning Time

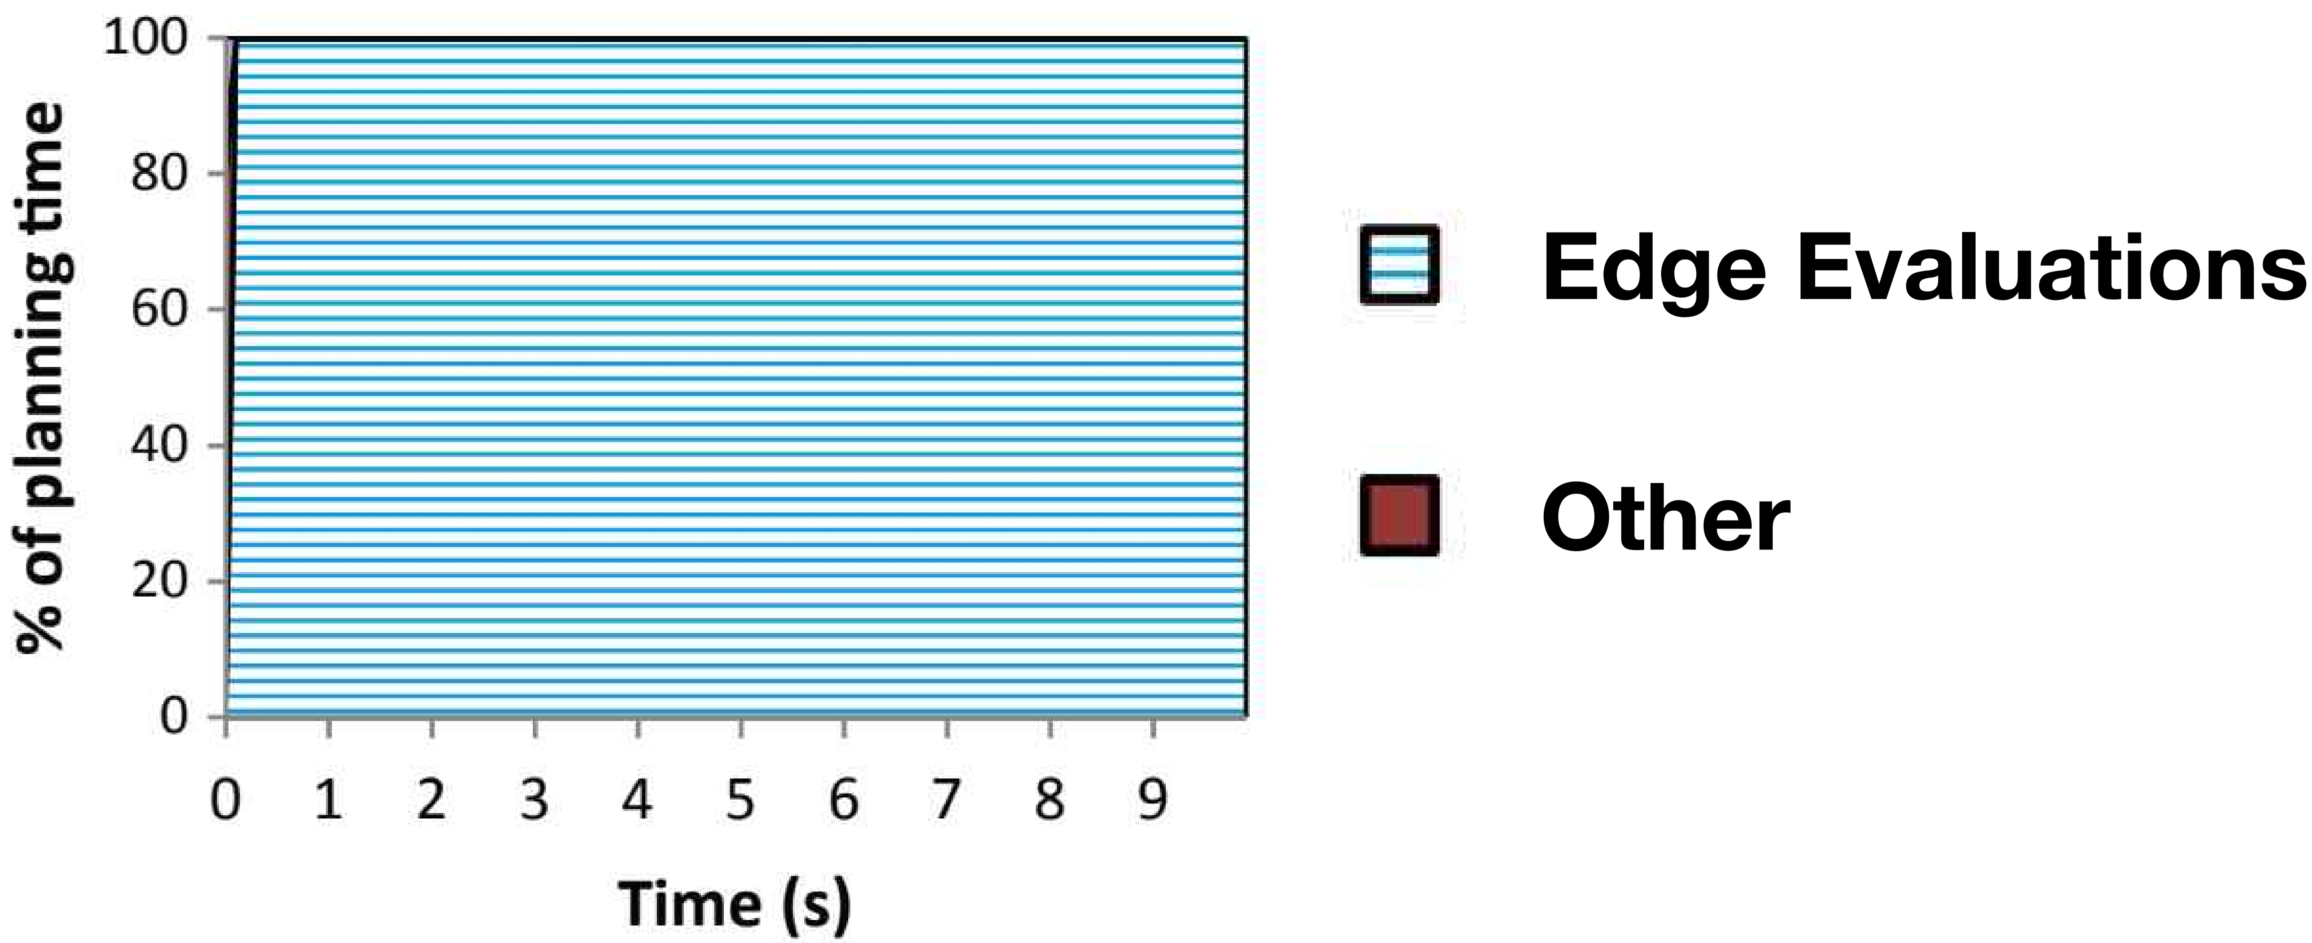

Hauser, Kris., Lazy collision checking in asymptotically-optimal motion planning. *ICRA* 2015

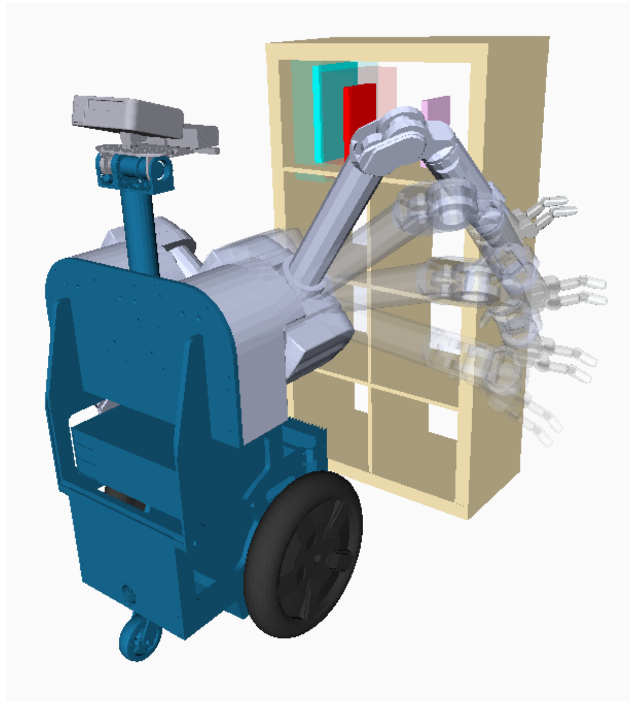

VS

Edge Evaluation vs. Search

Explicit vs. Implicit Graphs

Is there a Search Algorithm
that Minimizes
the Number of Edge Evaluations?

LazySP

ICAPS 2018 [Best Conference Paper Award Winner]

First Provably Edge-Optimal A*-like Search Algorithm

Number of Edges Evaluated


Number of Edges Evaluated
A*
LazySP

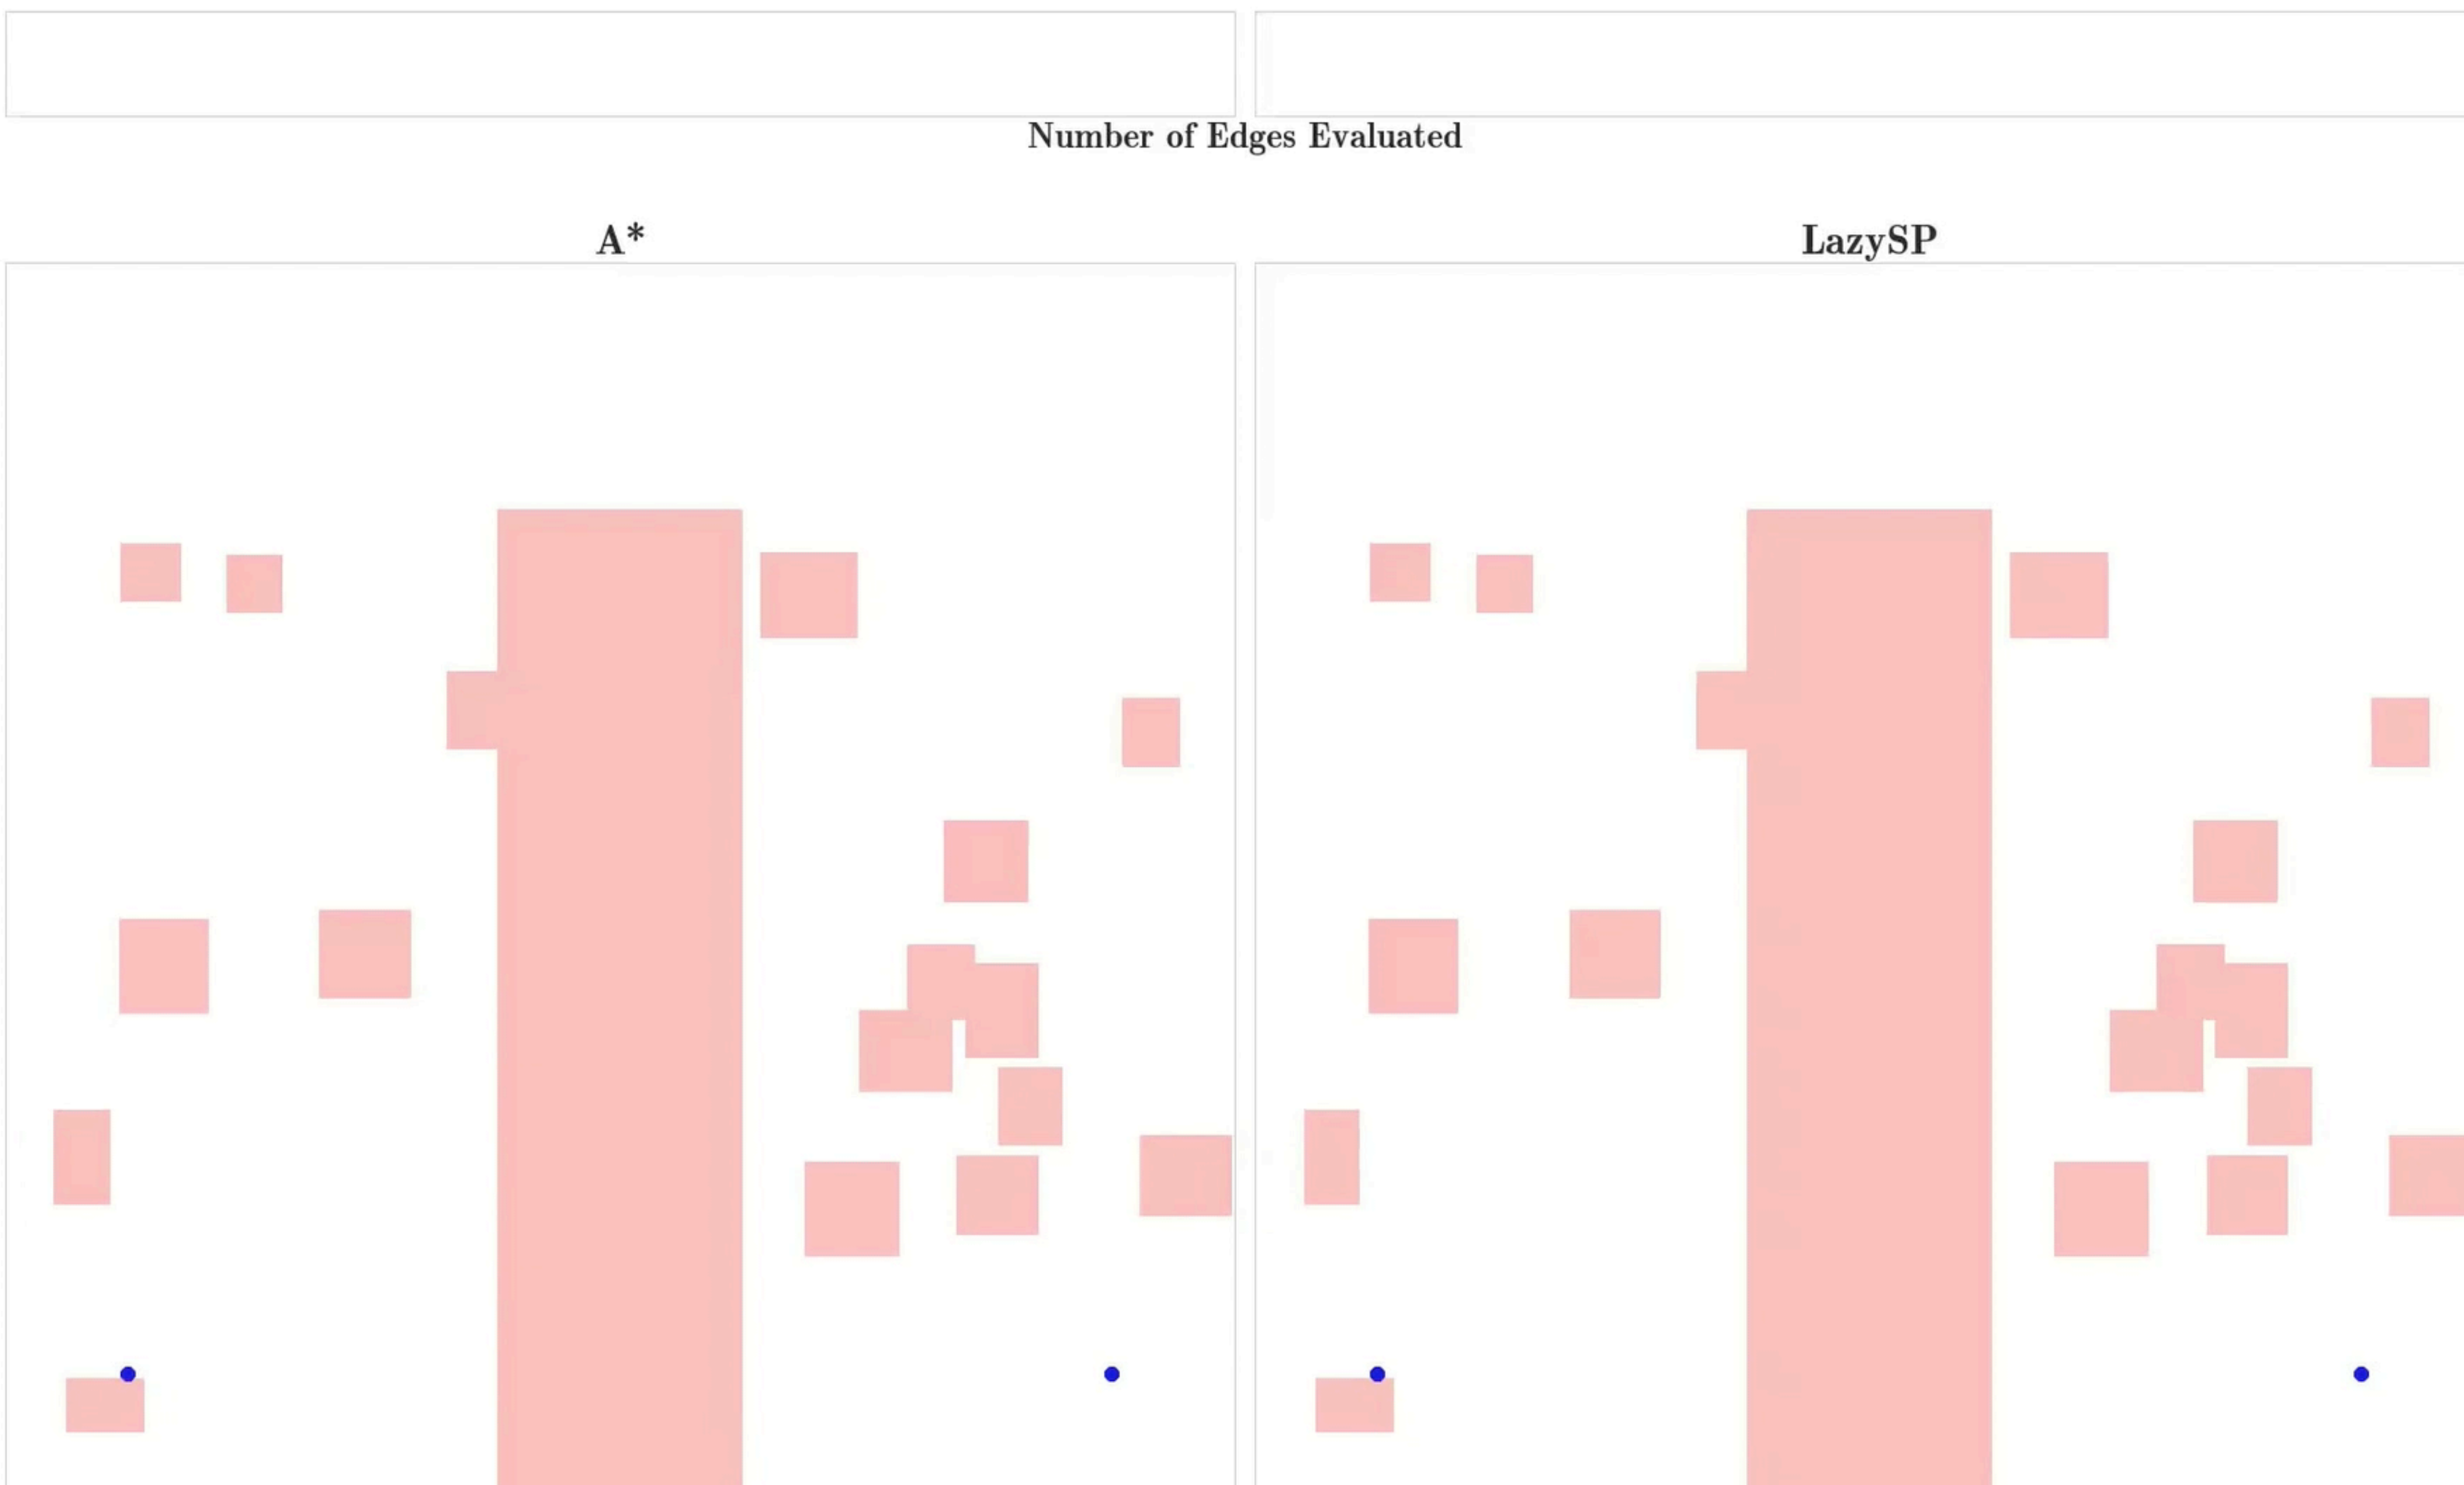
Number of Edges Evaluated
A*
LazySP

Number of Edges Evaluated

A* LazySP

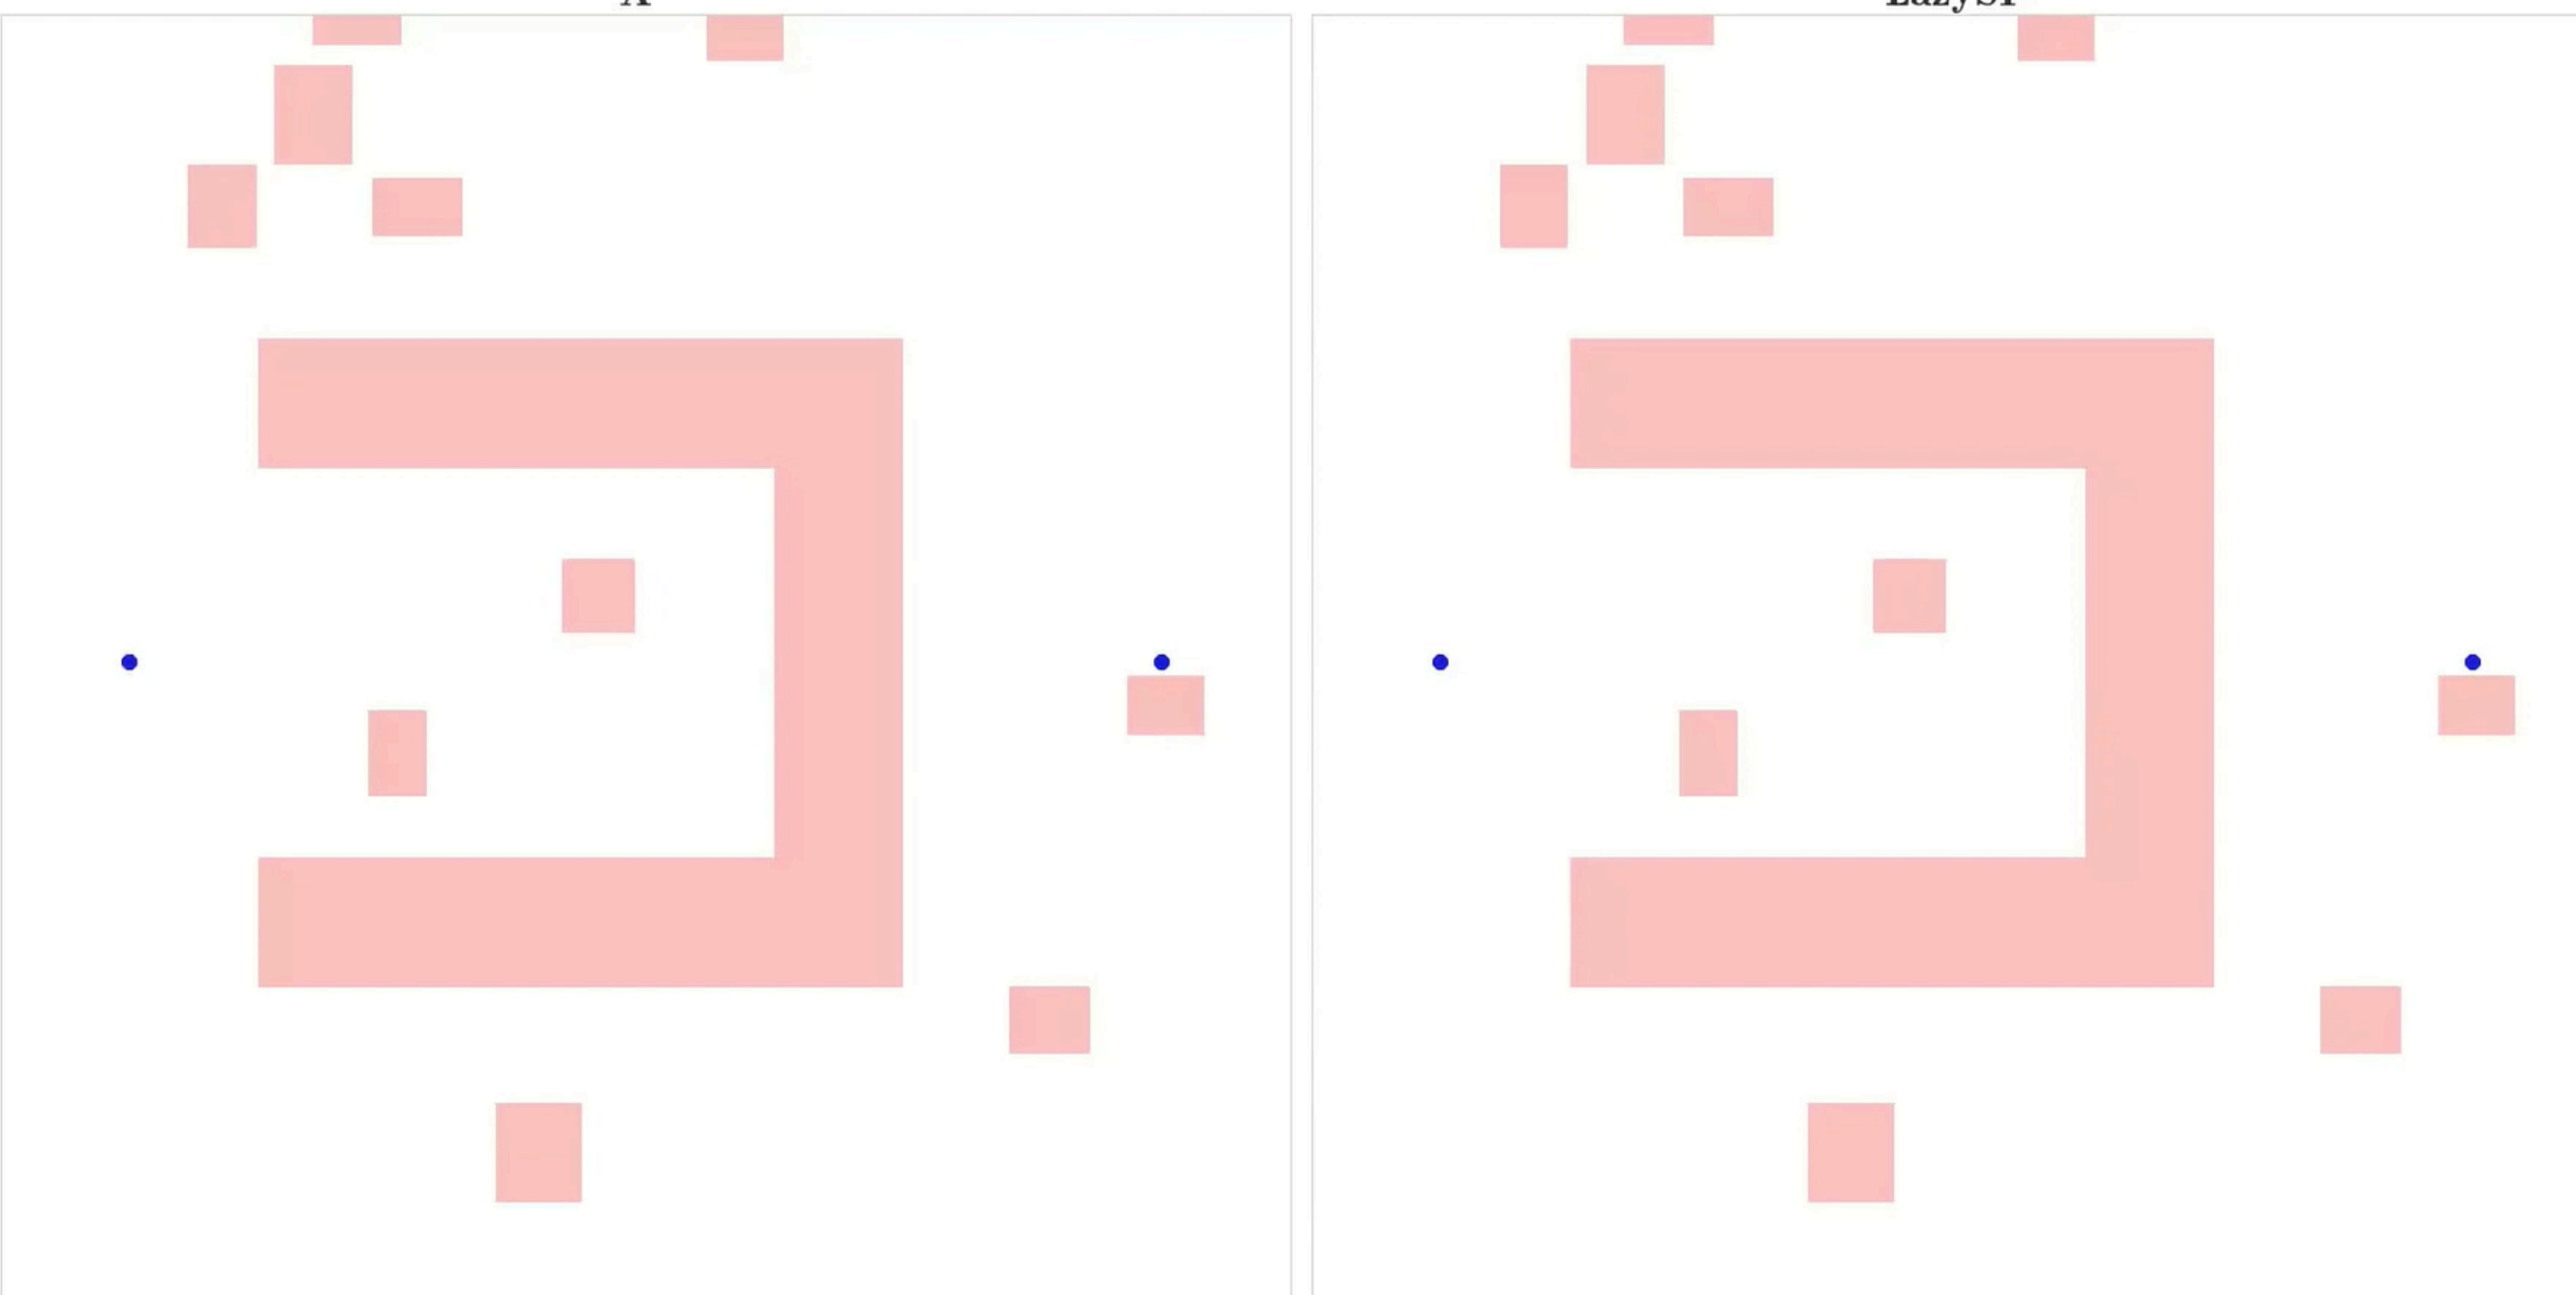

LazySP

Greedy Best-first Search over Paths

To find the shortest path, eliminate all shorter paths!

LazySP

Optimism Under Uncertainty

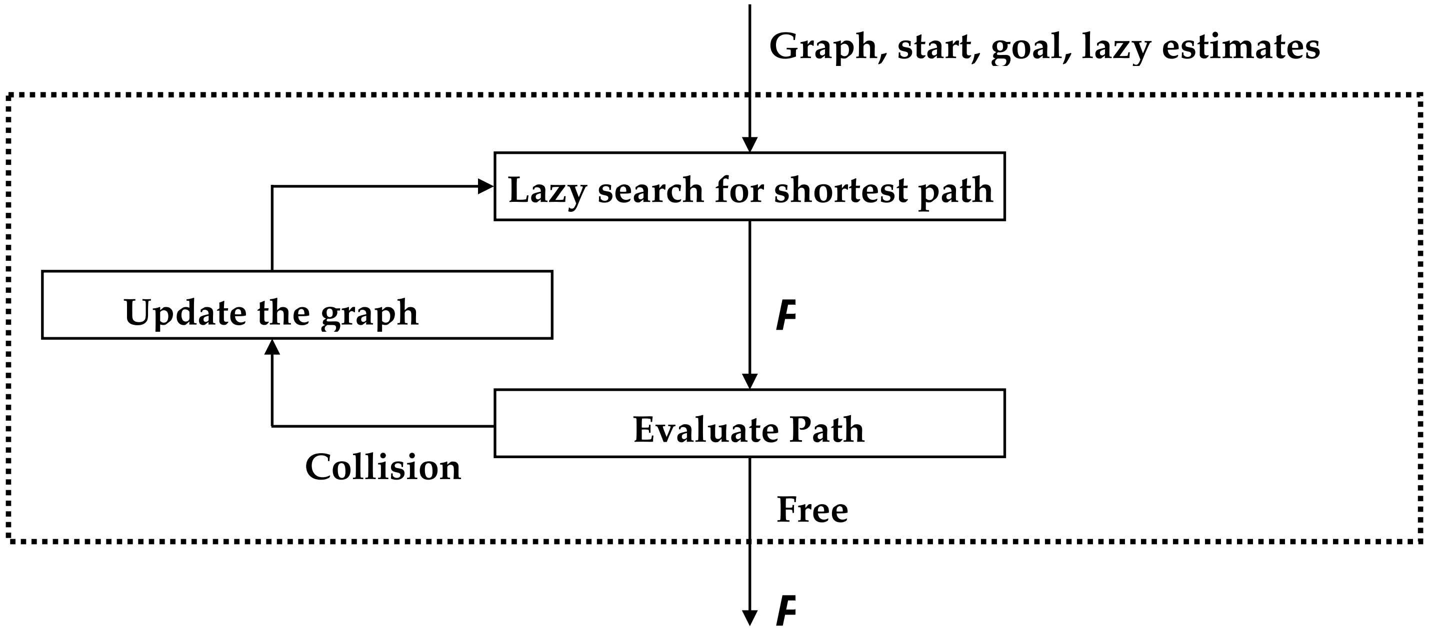

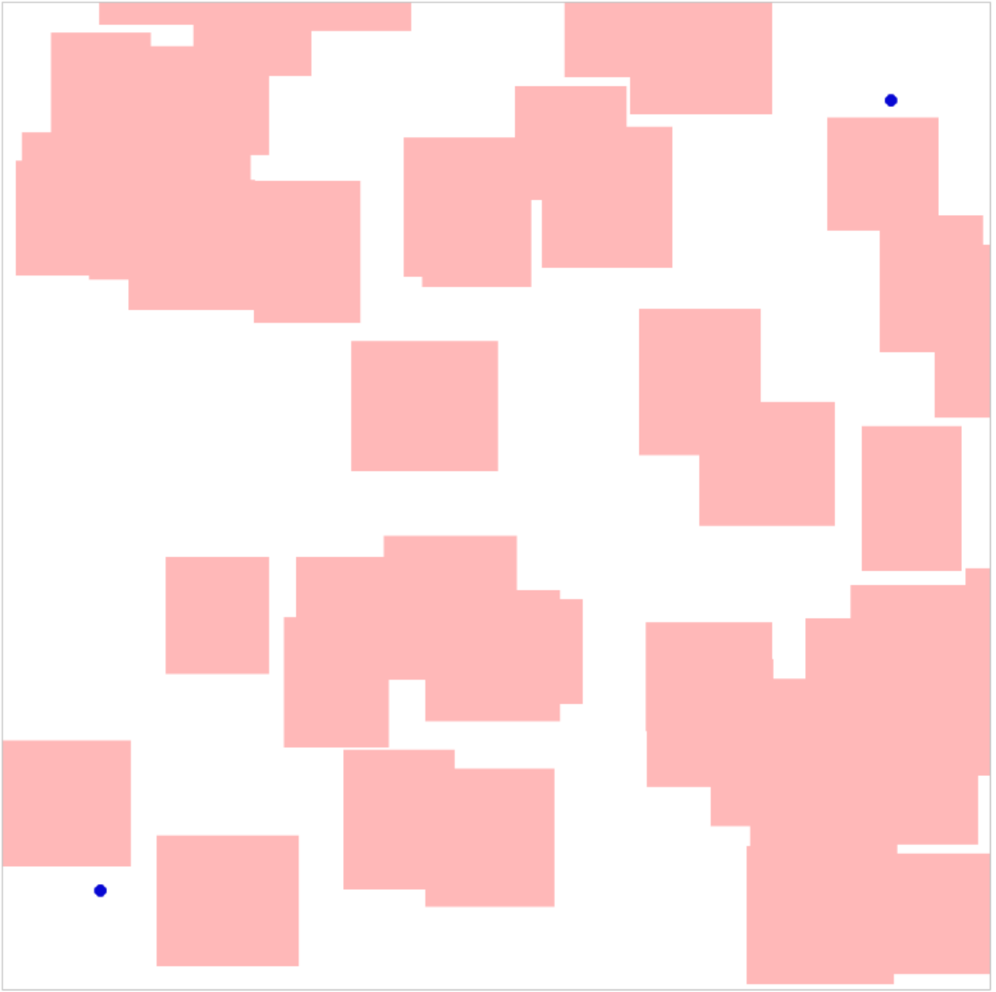

LazySP

Optimism Under Uncertainty

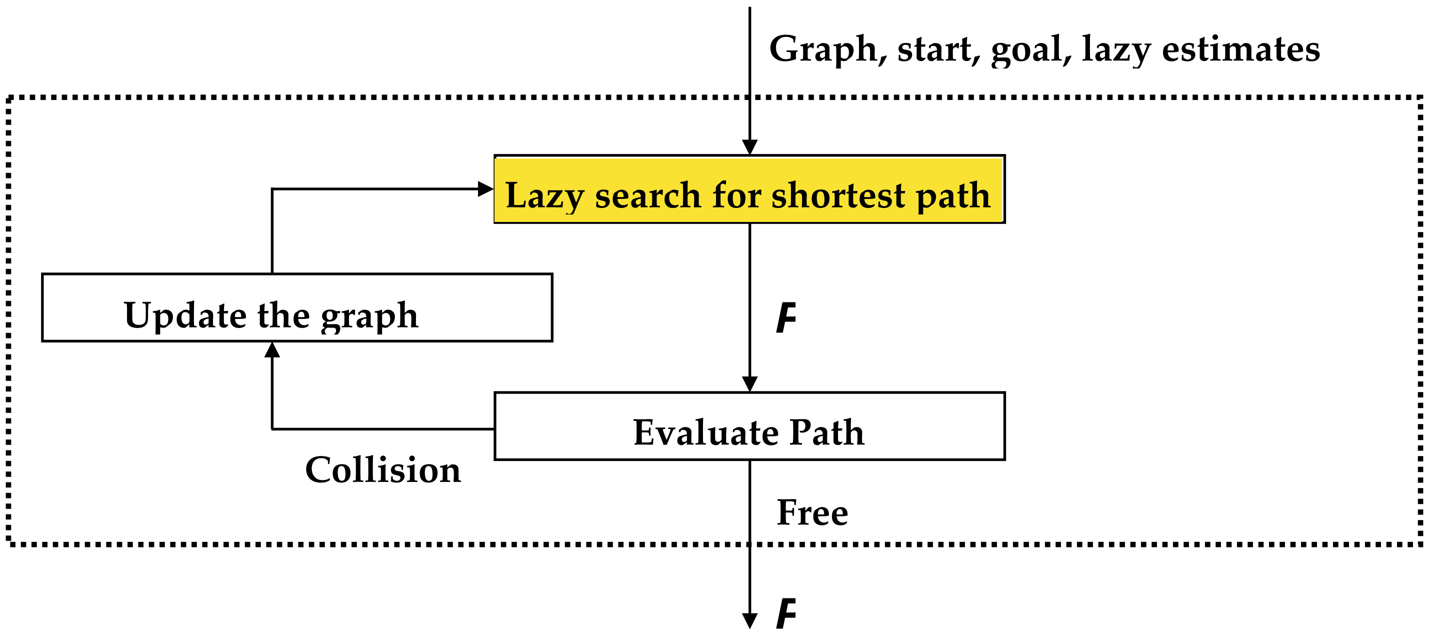

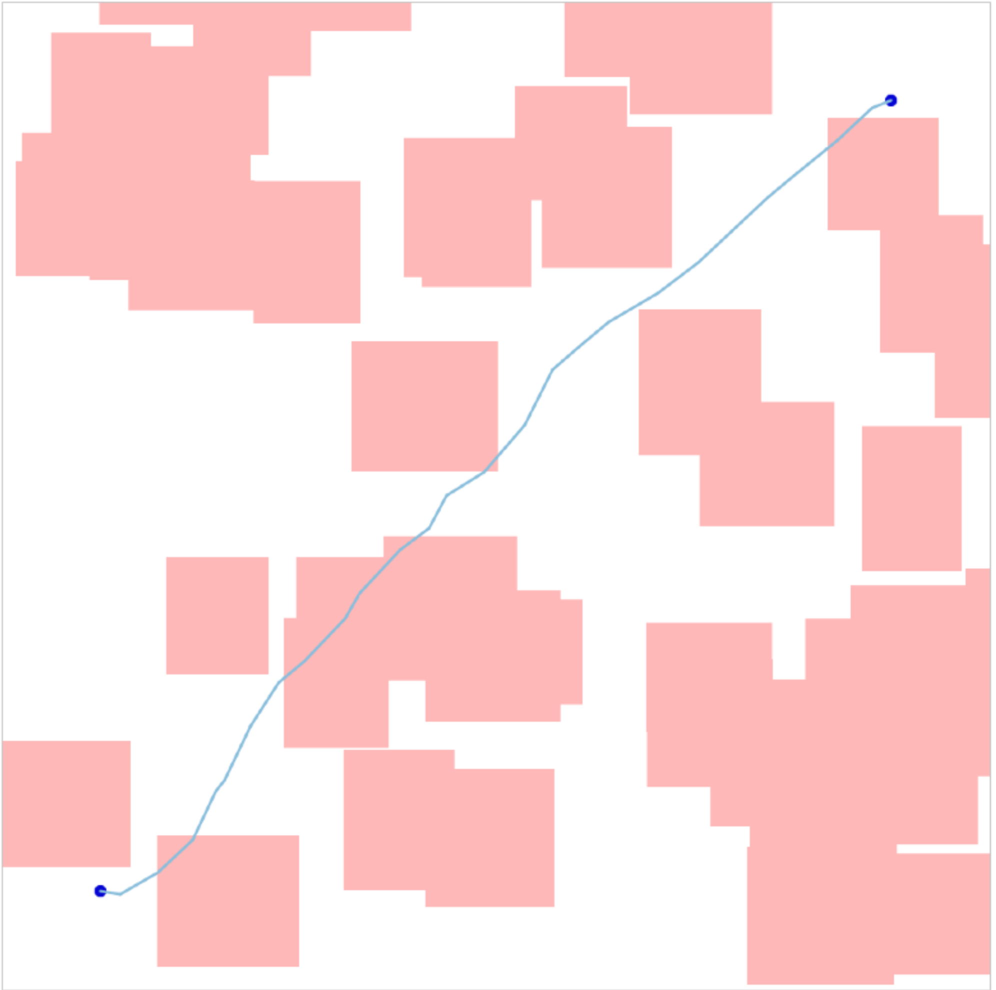

LazySP

Optimism Under Uncertainty

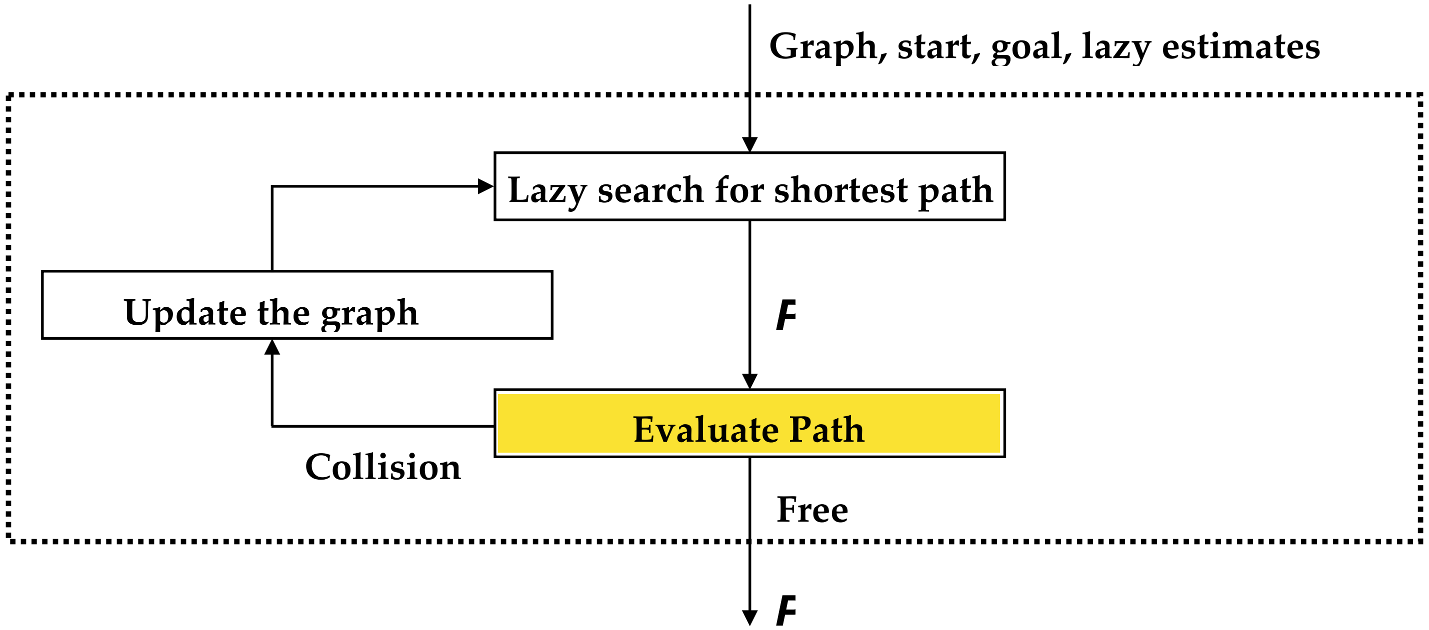

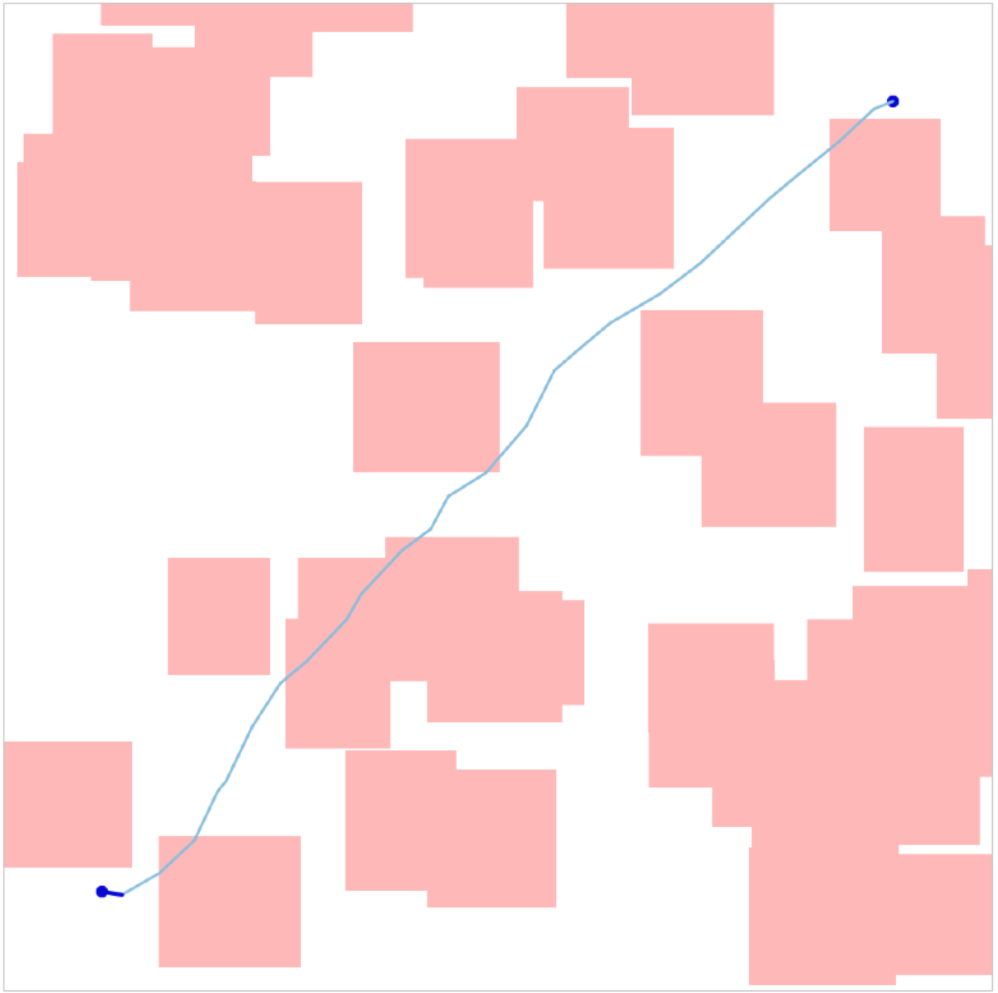

LazySP

Optimism Under Uncertainty

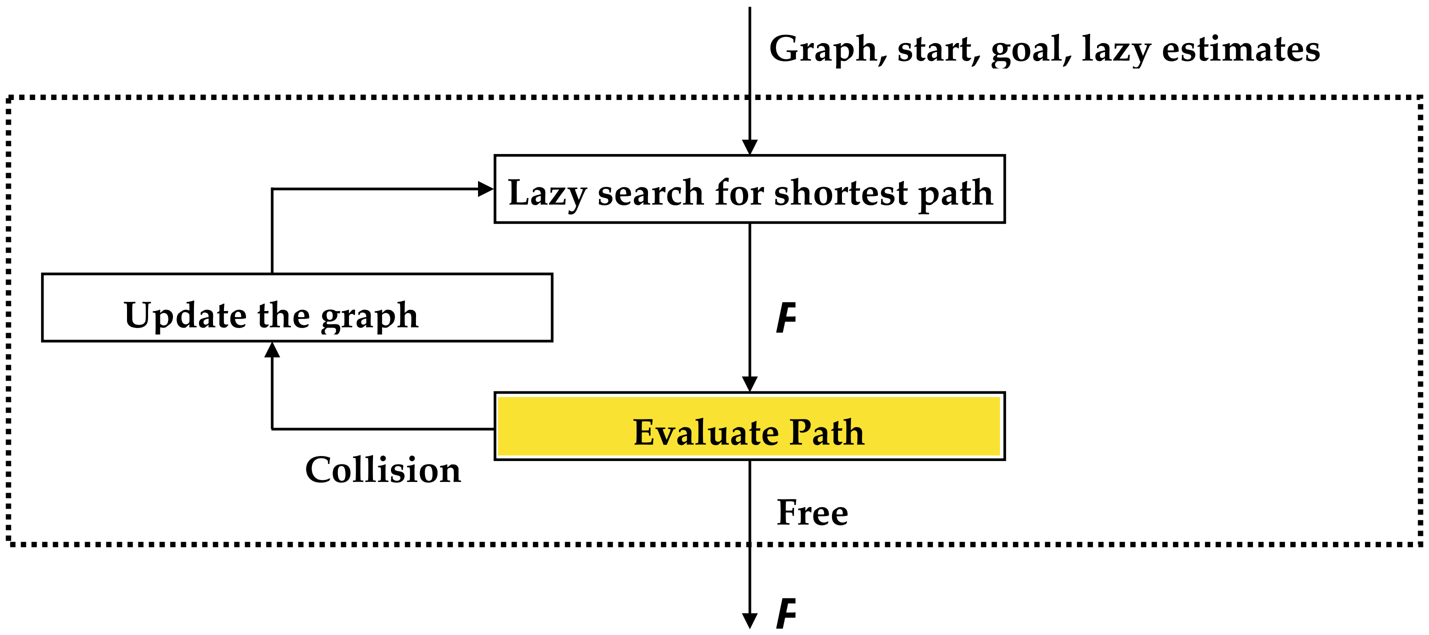

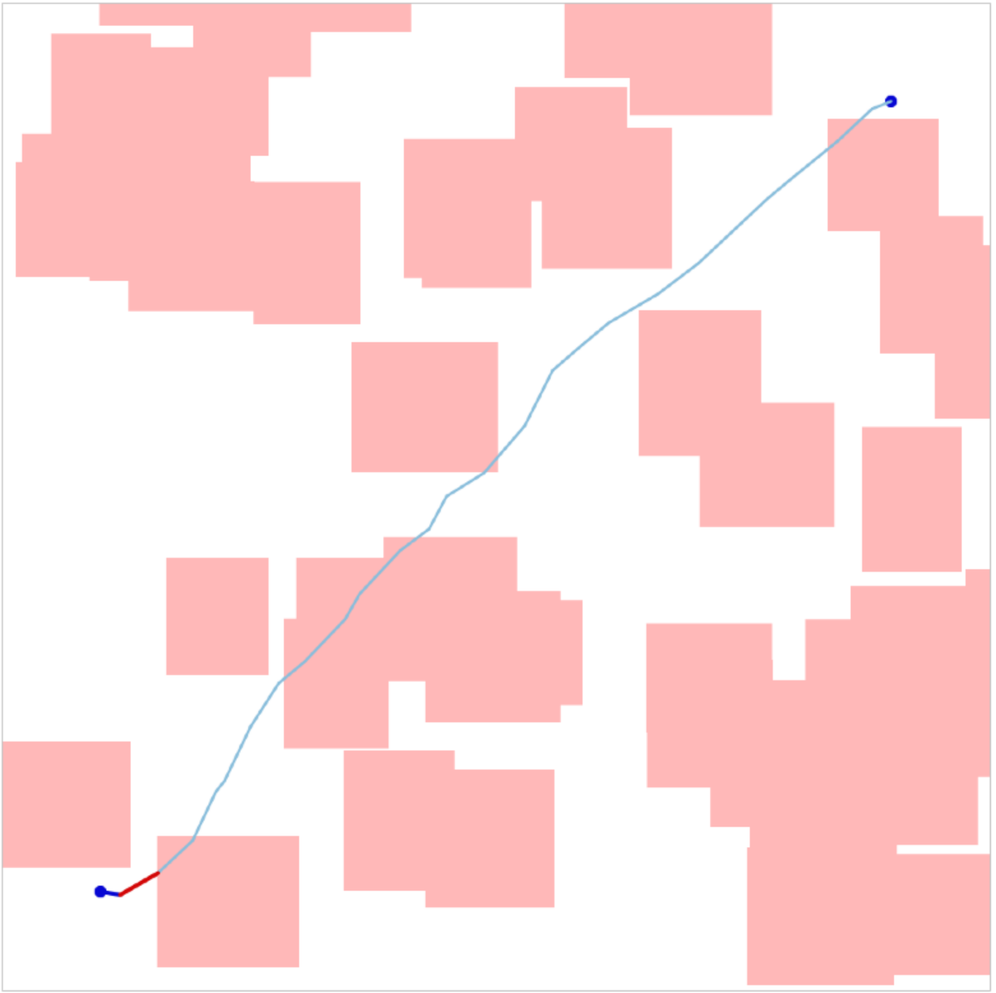

LazySP

Optimism Under Uncertainty

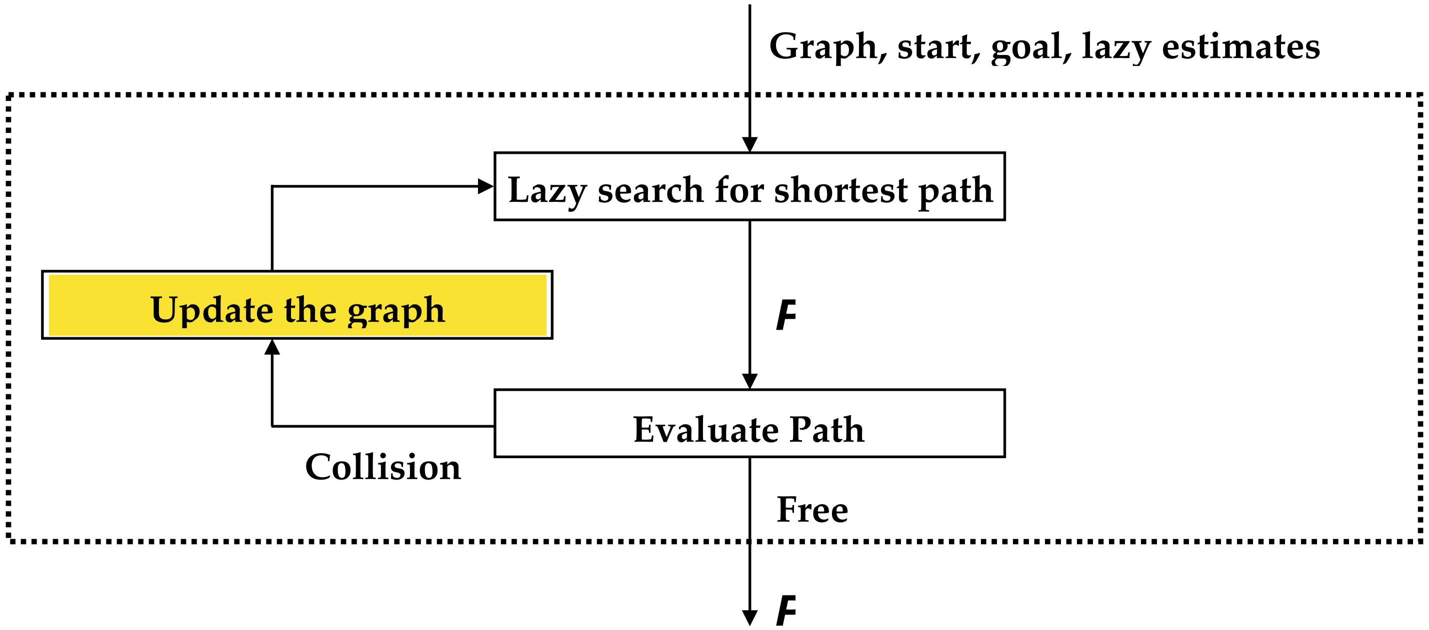

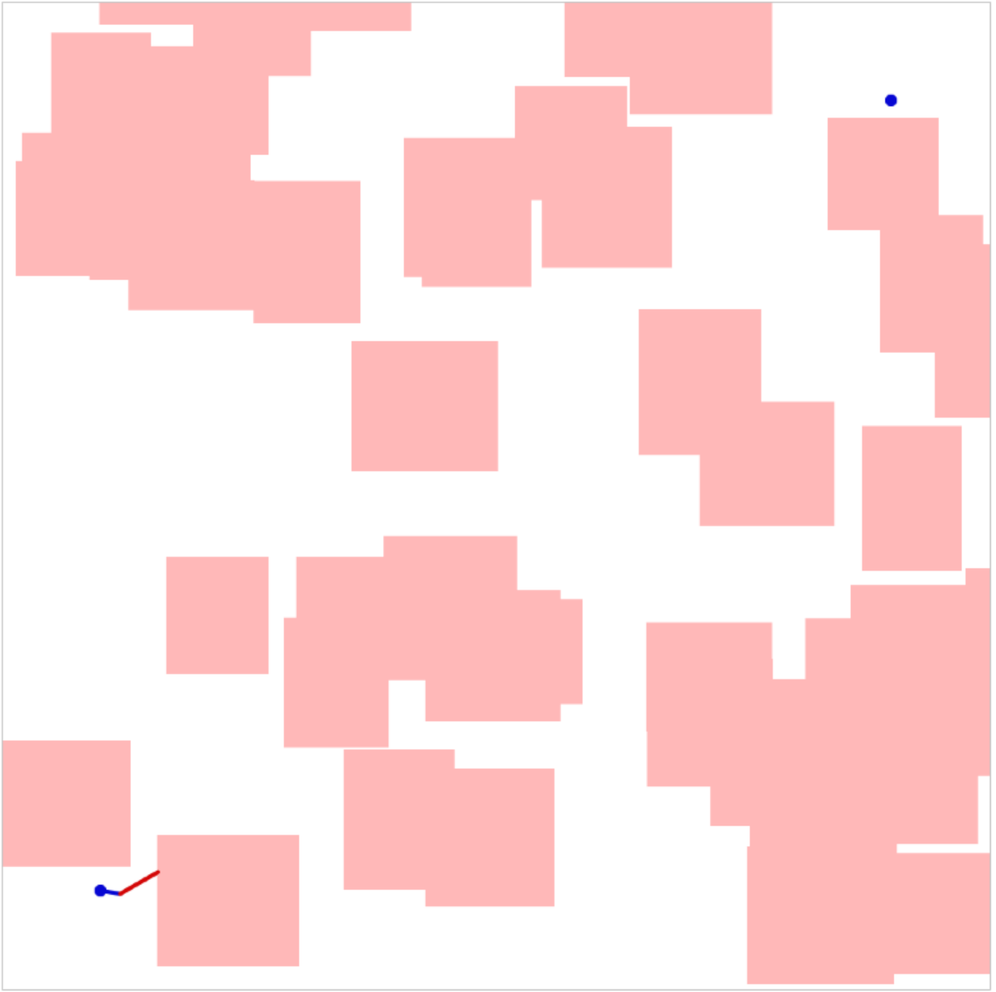

LazySP

Optimism Under Uncertainty

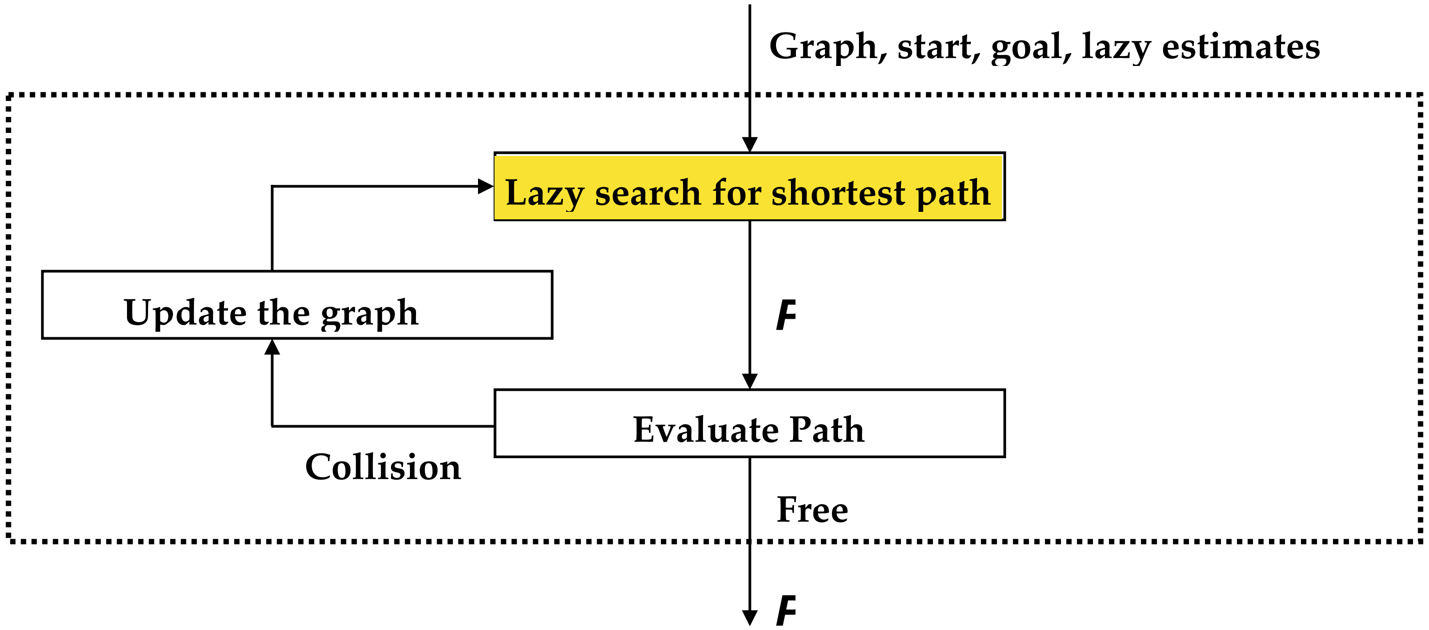

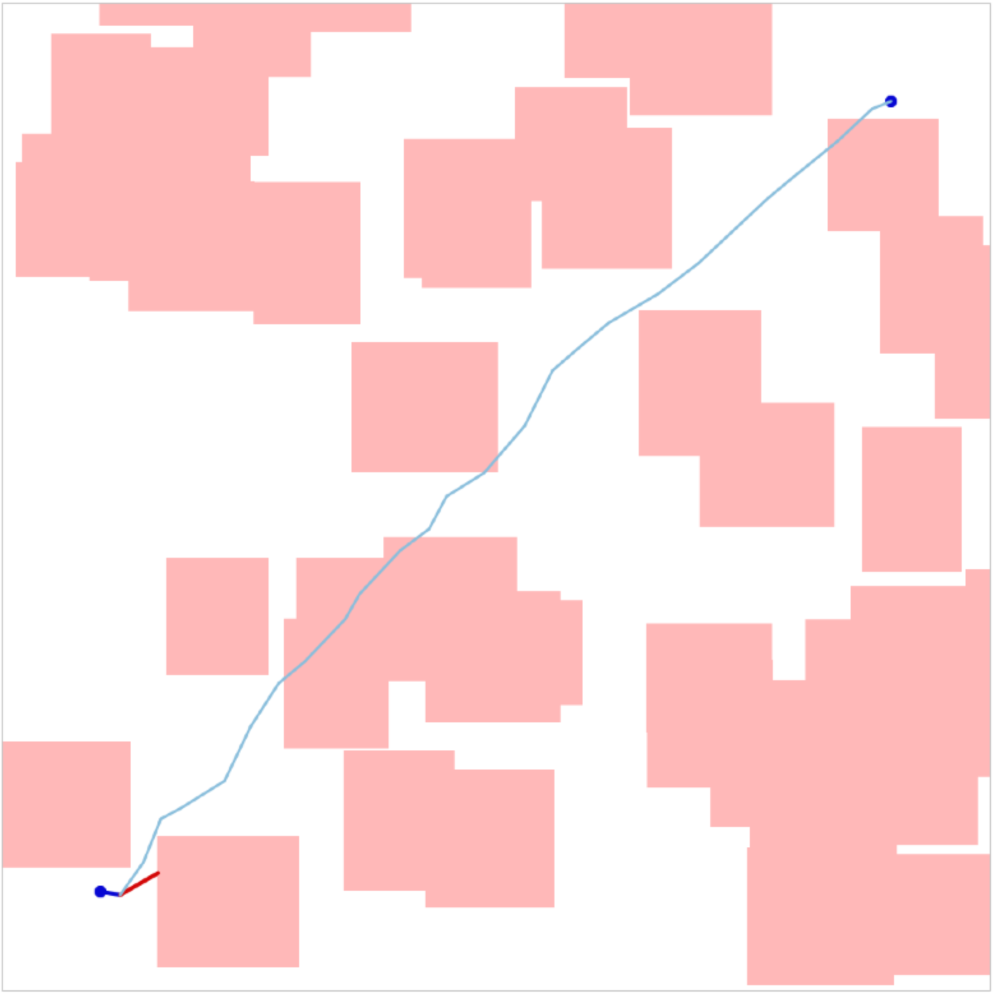

LazySP

Optimism Under Uncertainty

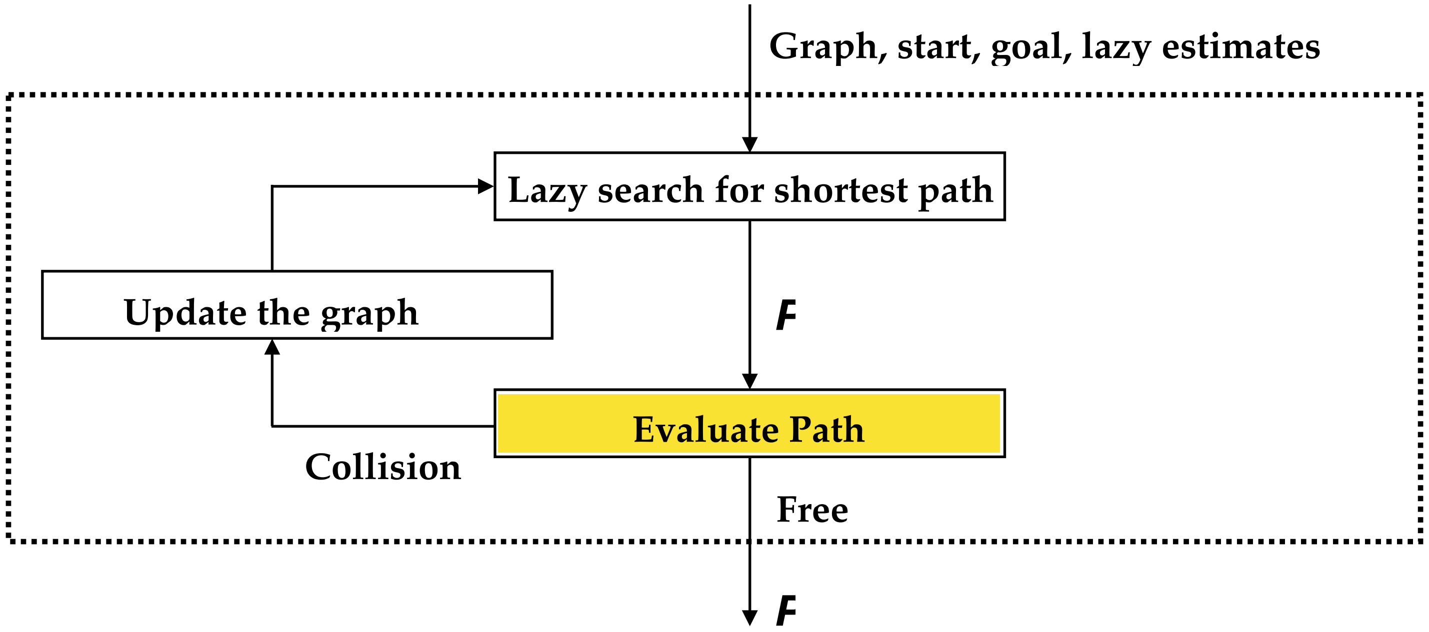

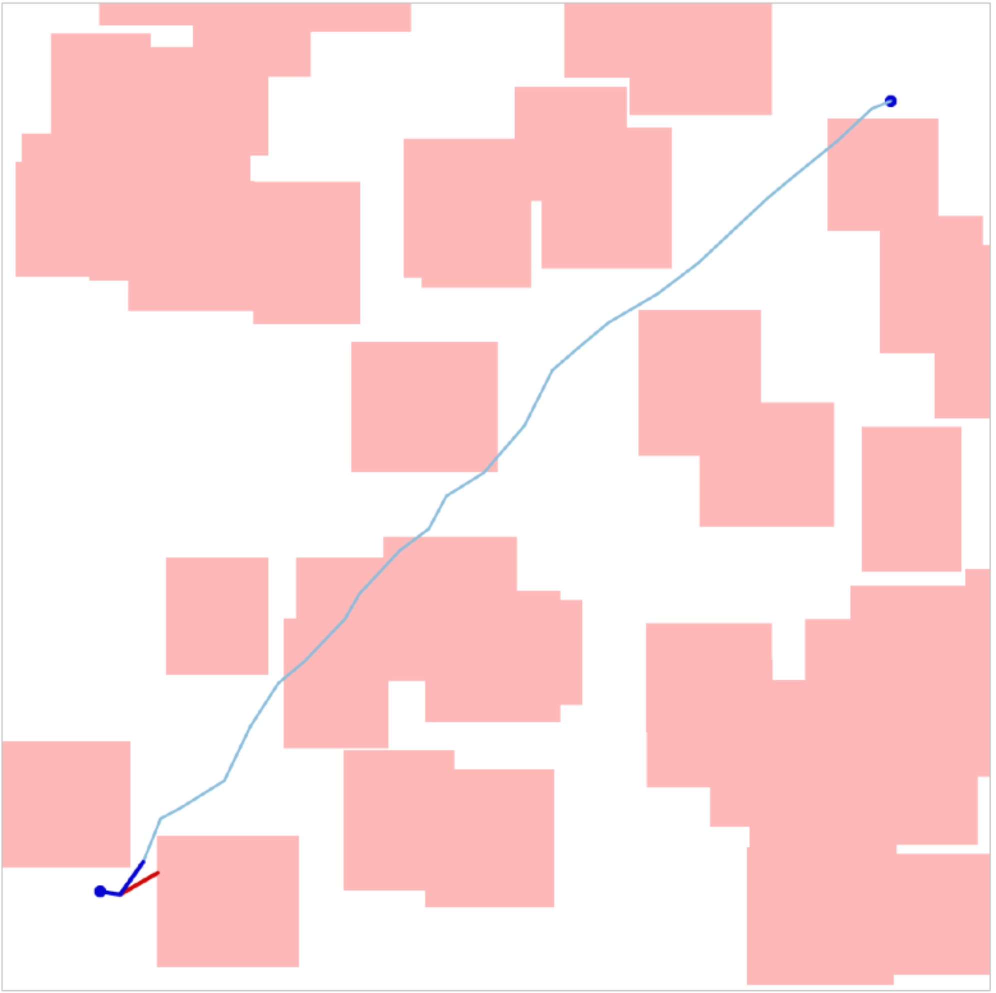

LazySP

Optimism Under Uncertainty

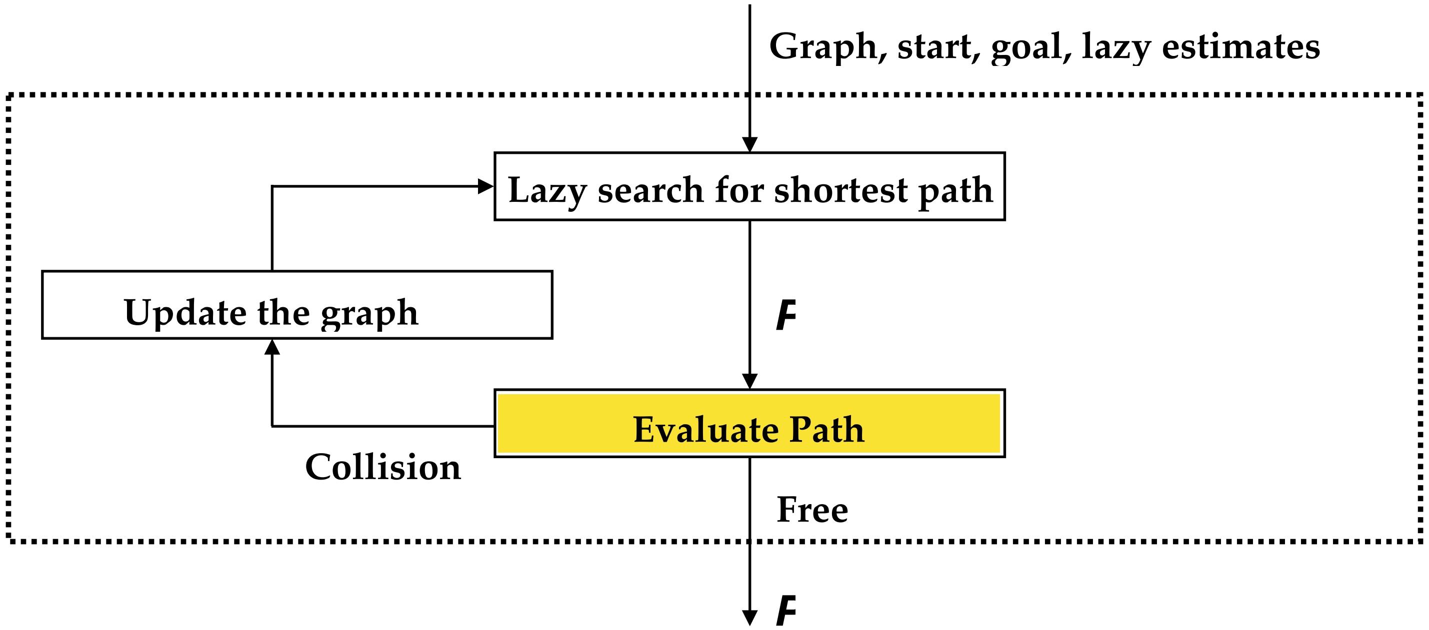

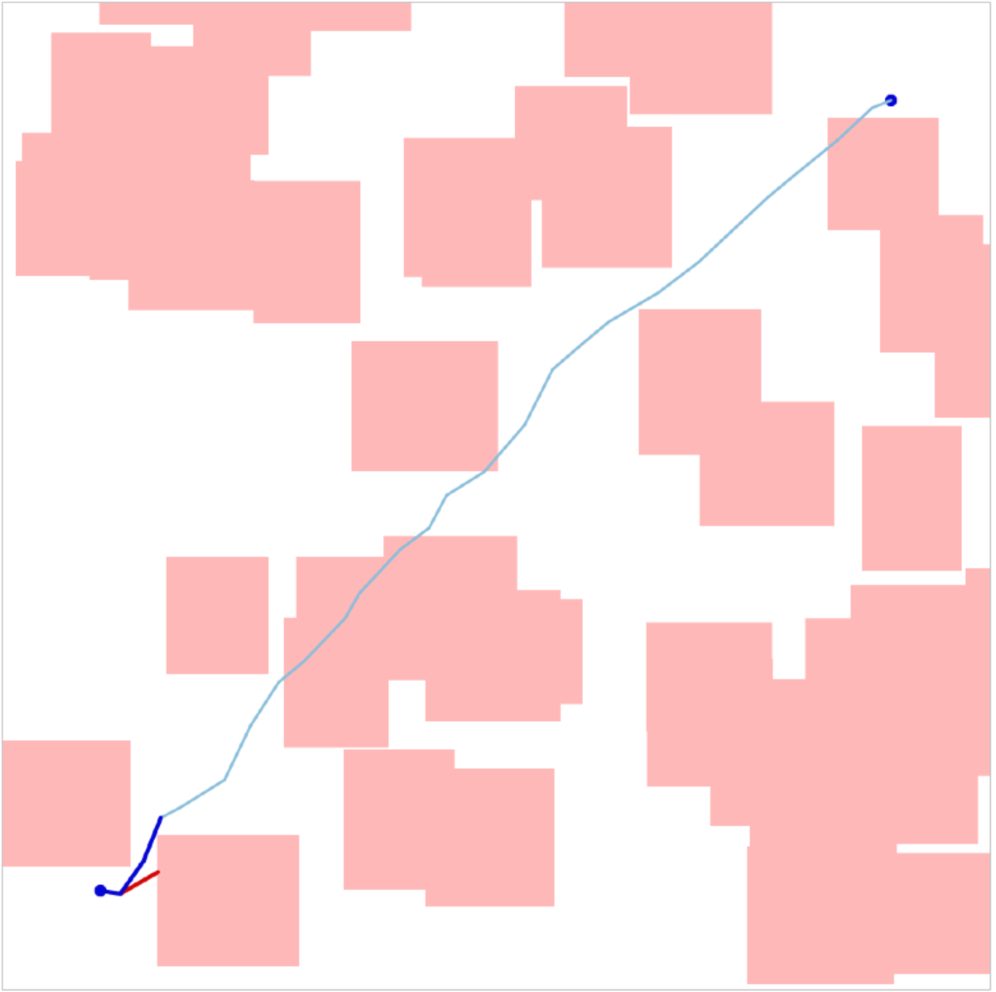

LazySP

Optimism Under Uncertainty

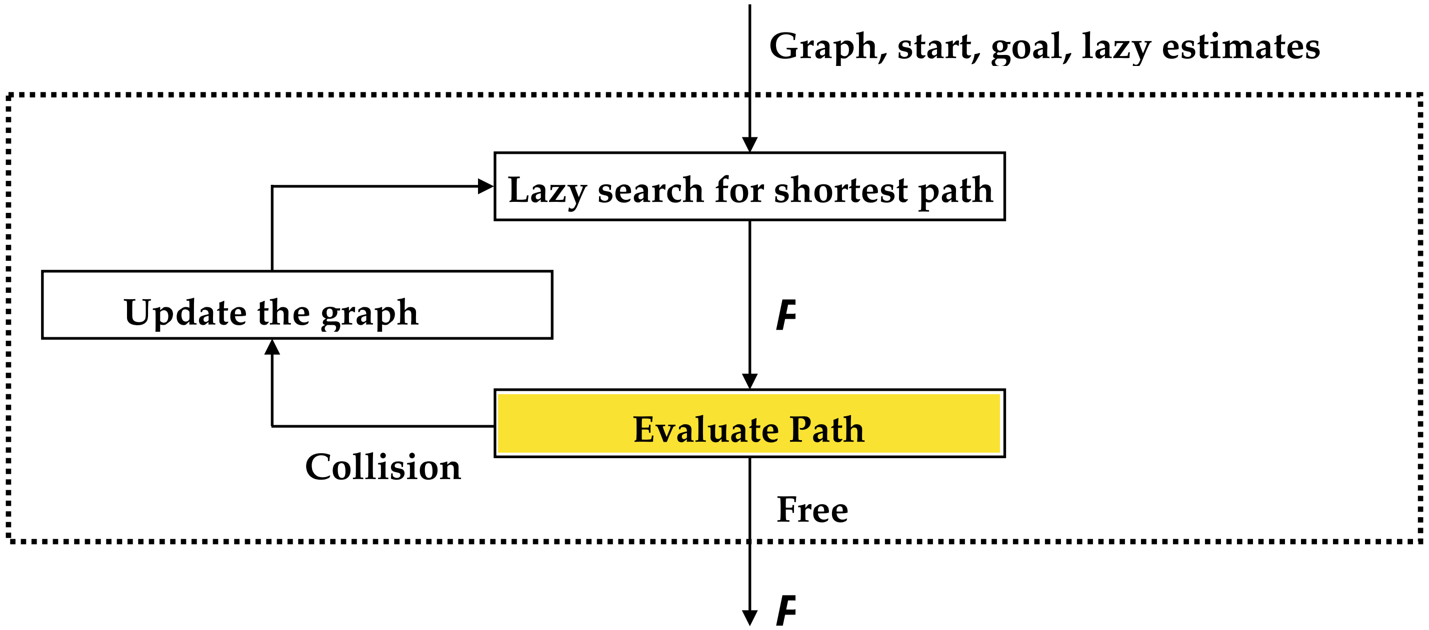

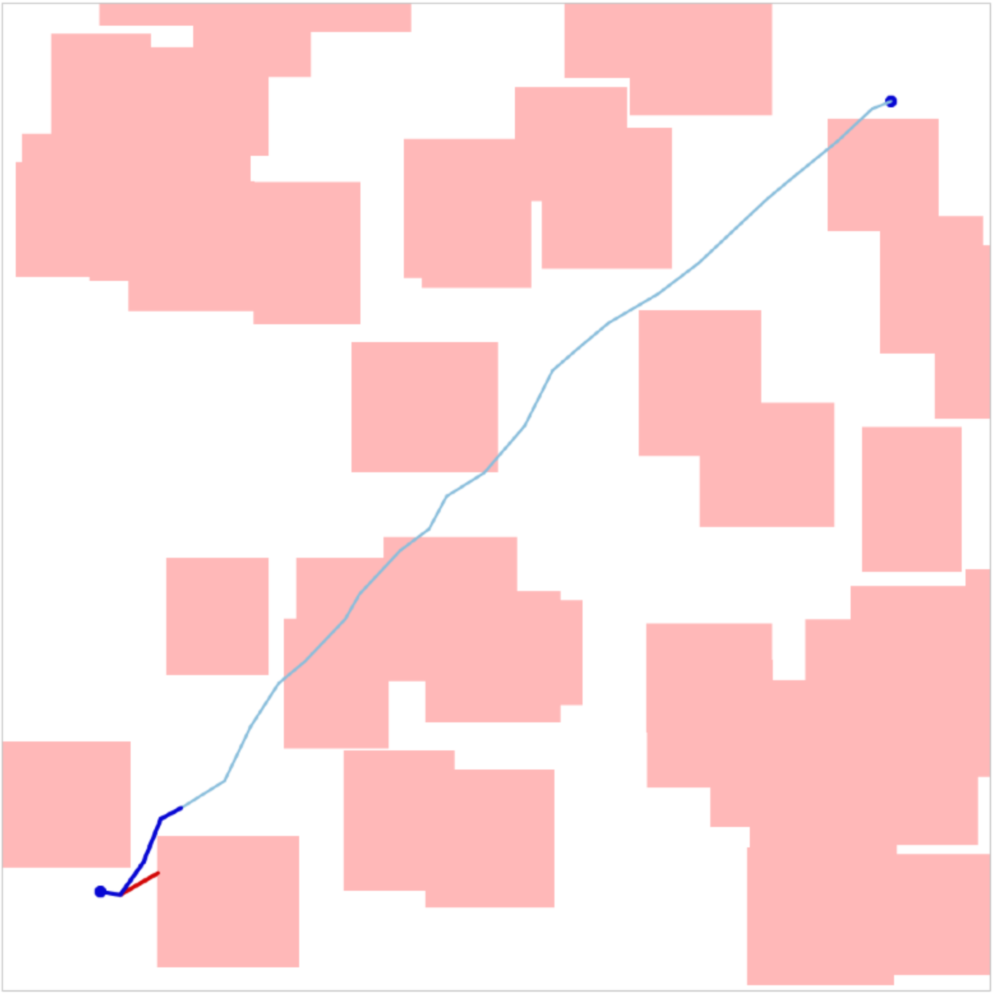

LazySP

Optimism Under Uncertainty

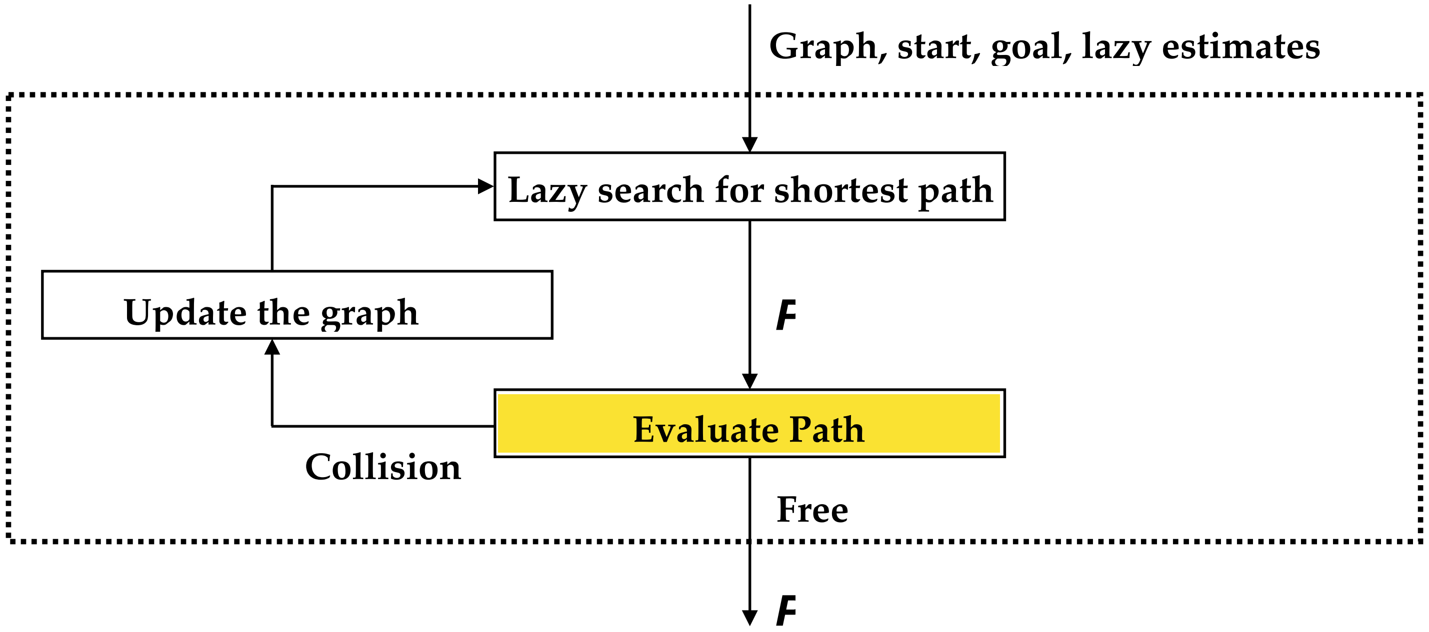

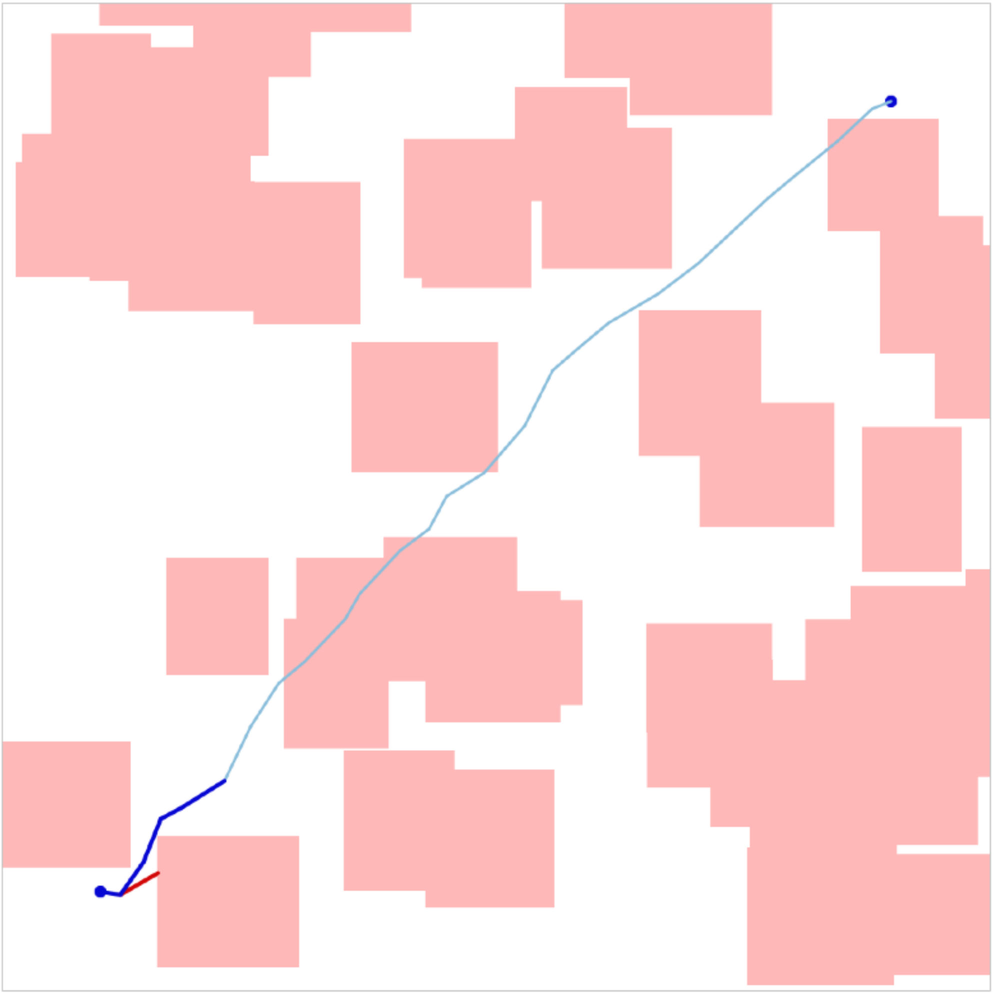

LazySP

Optimism Under Uncertainty

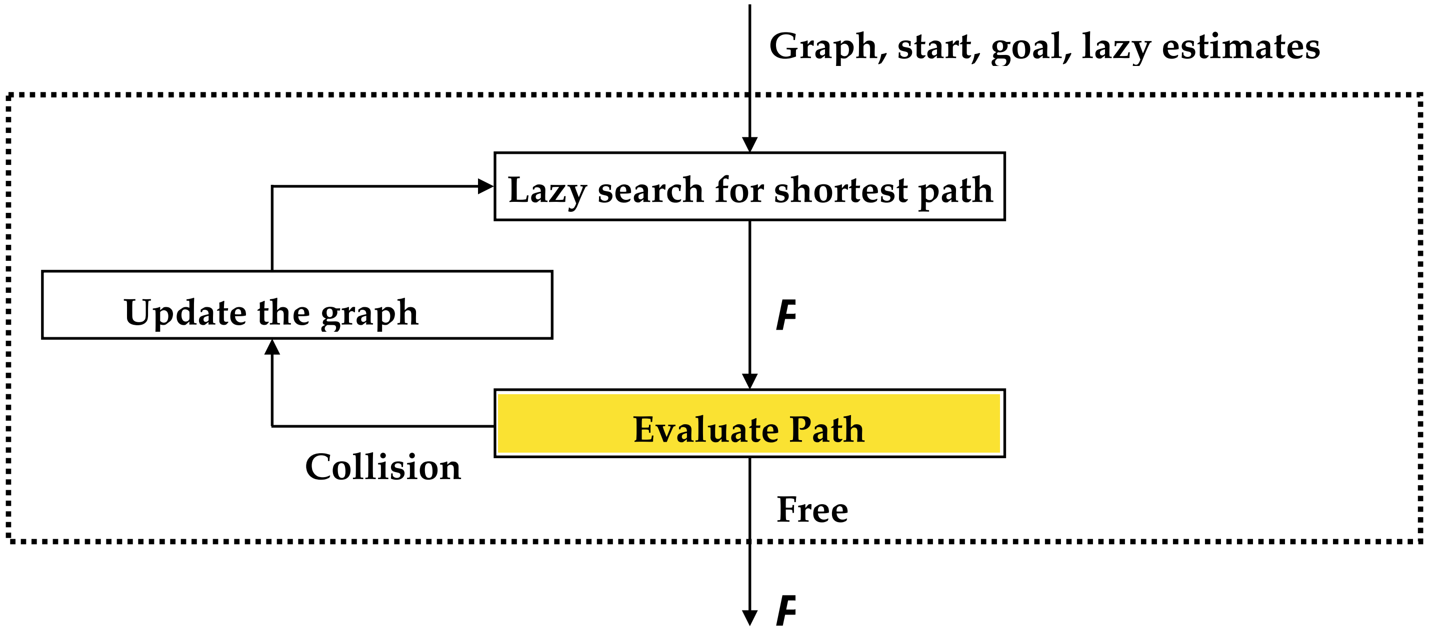

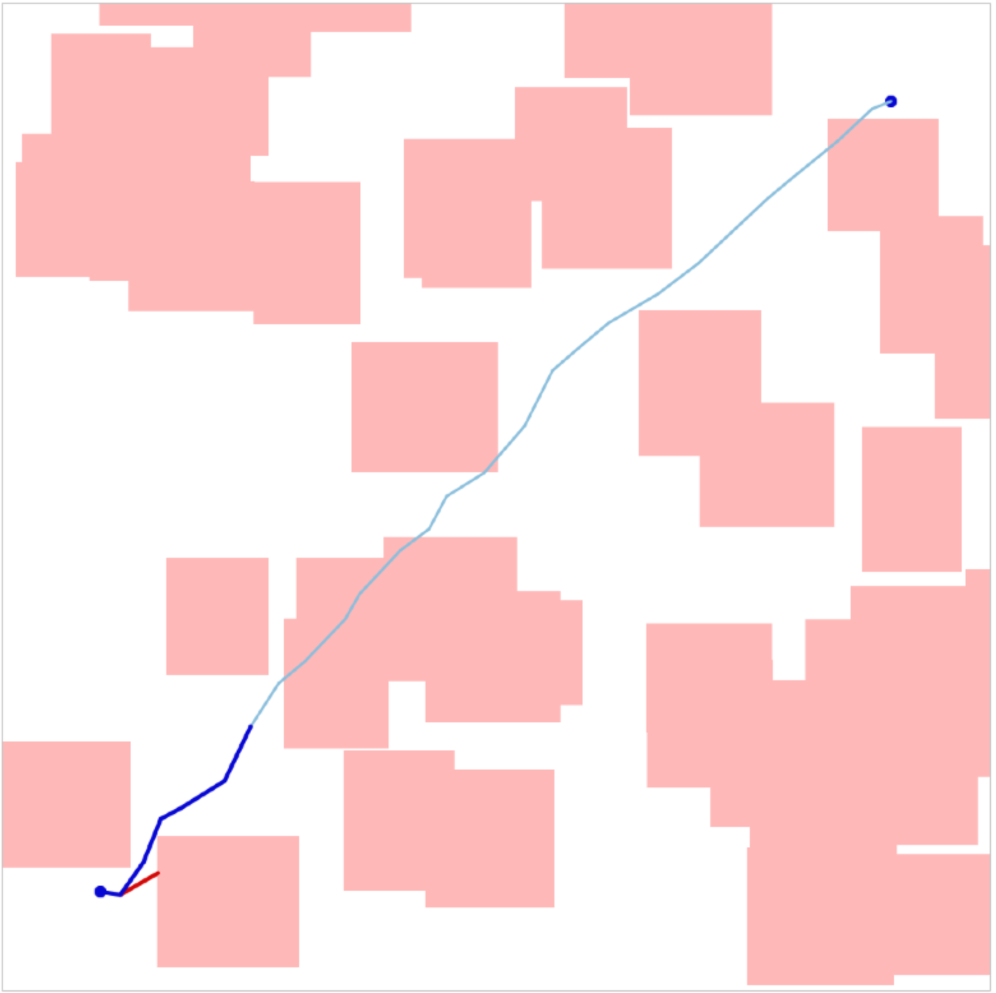

LazySP

Optimism Under Uncertainty

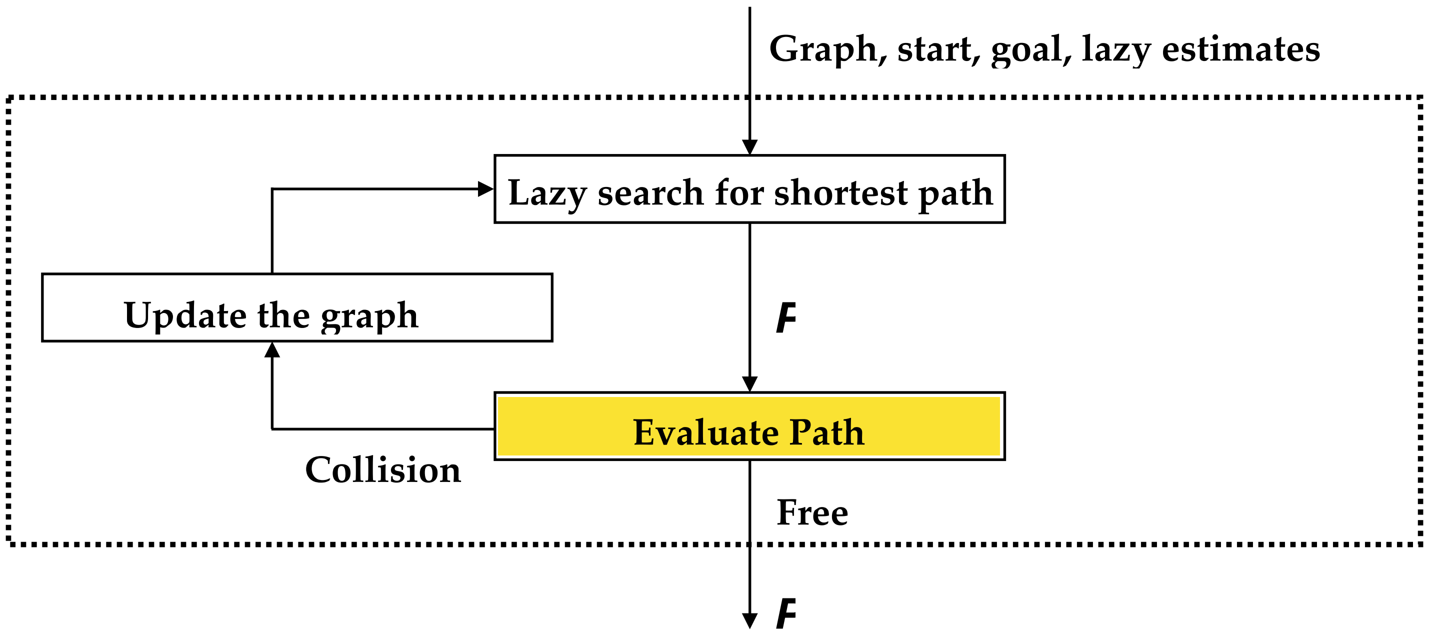

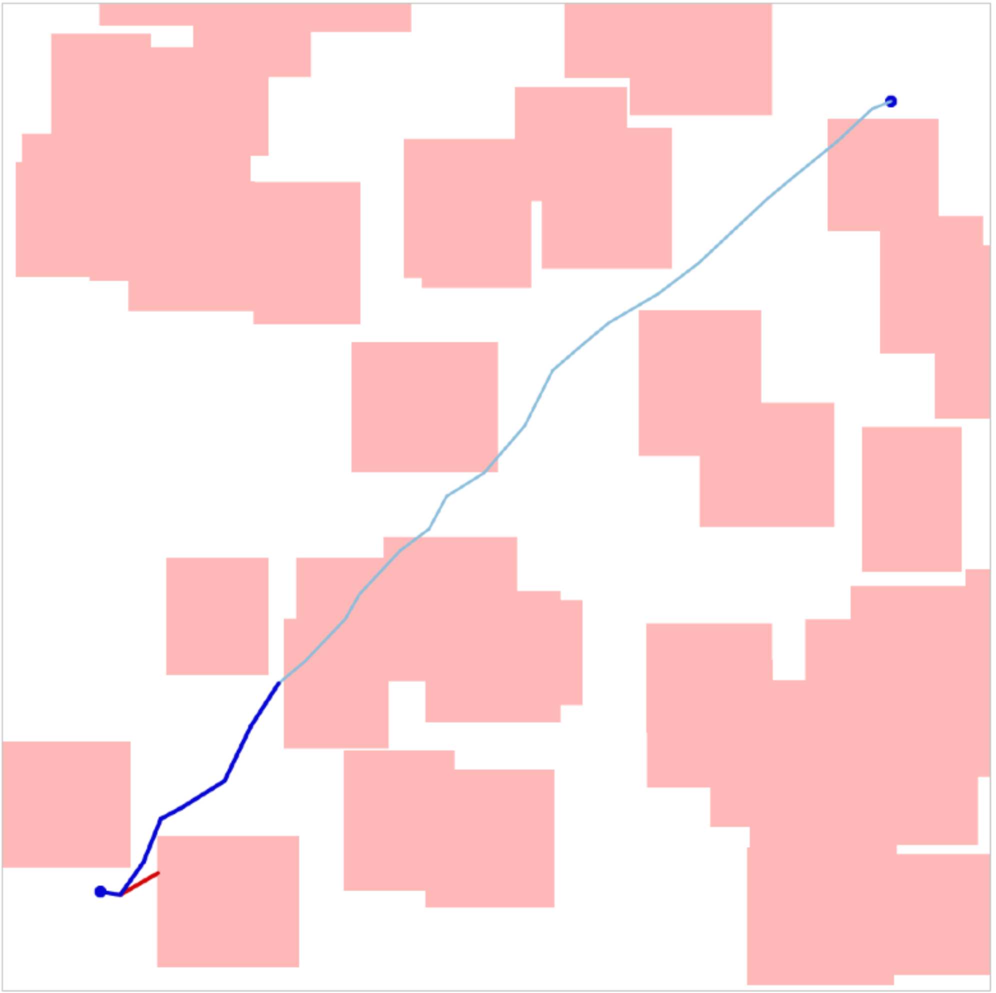

LazySP

Optimism Under Uncertainty

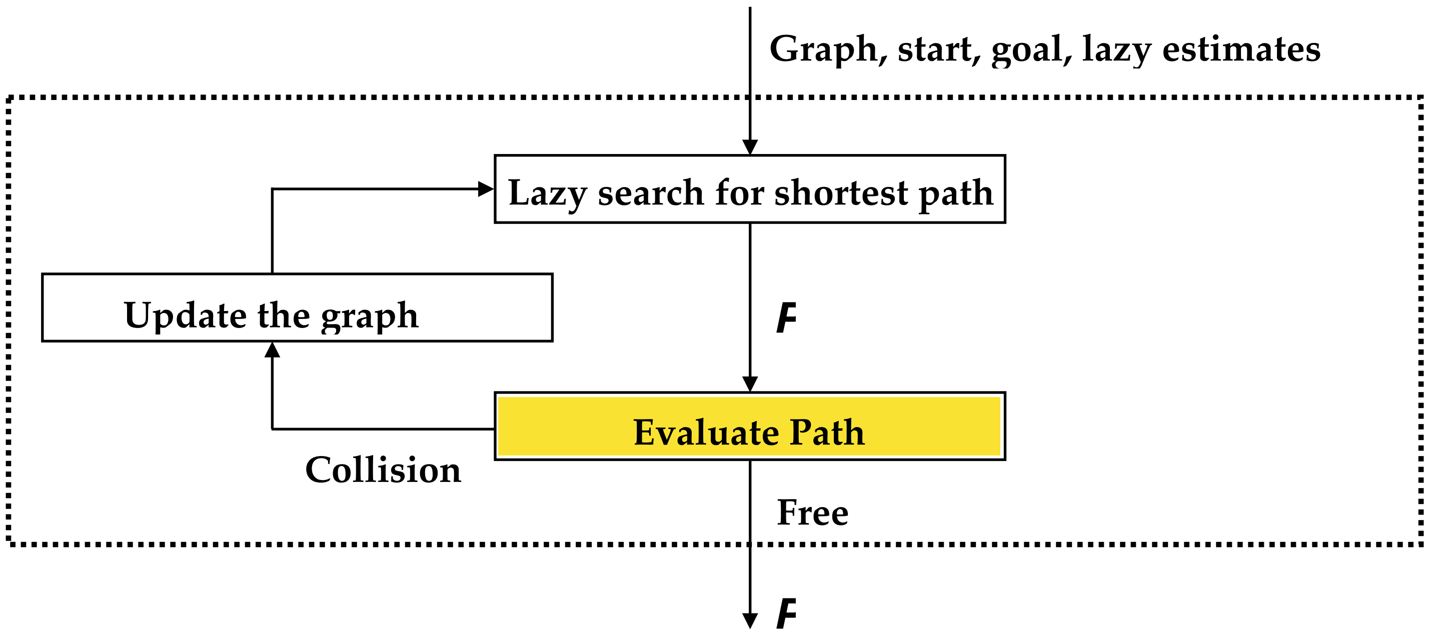

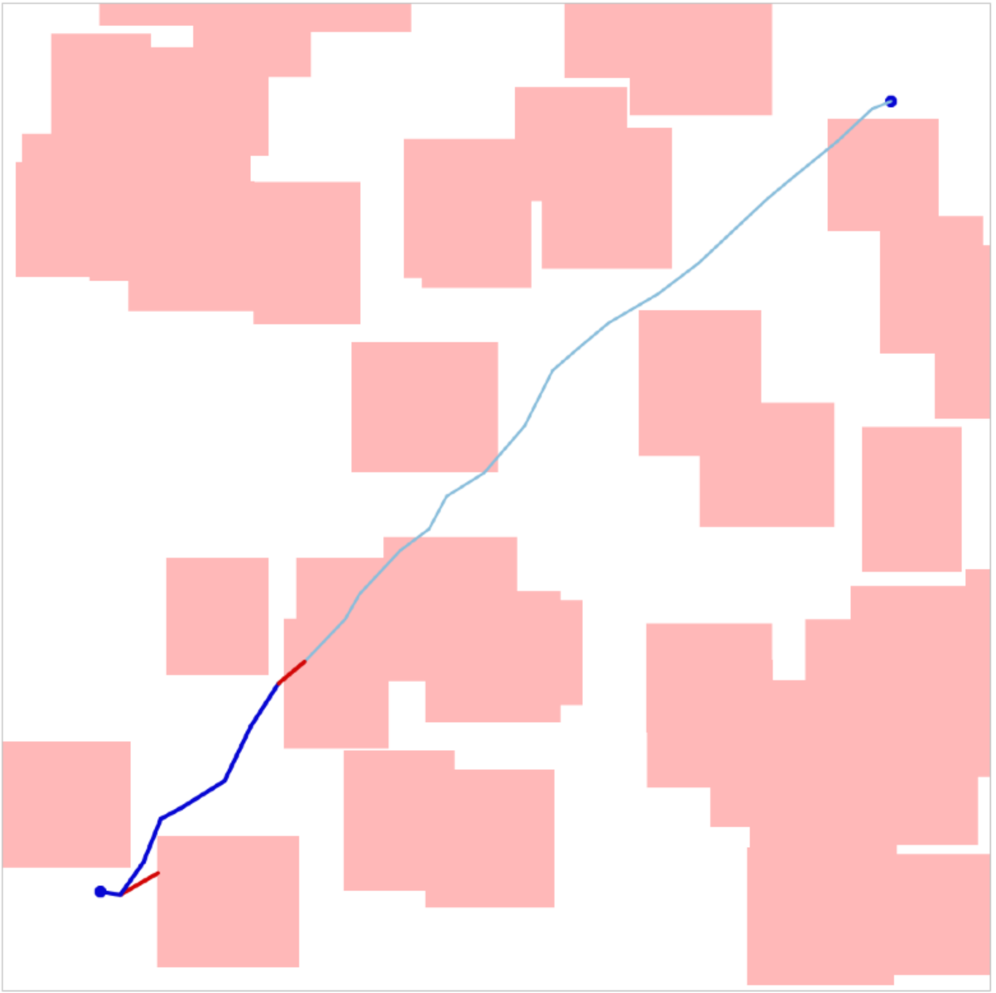

LazySP

Optimism Under Uncertainty

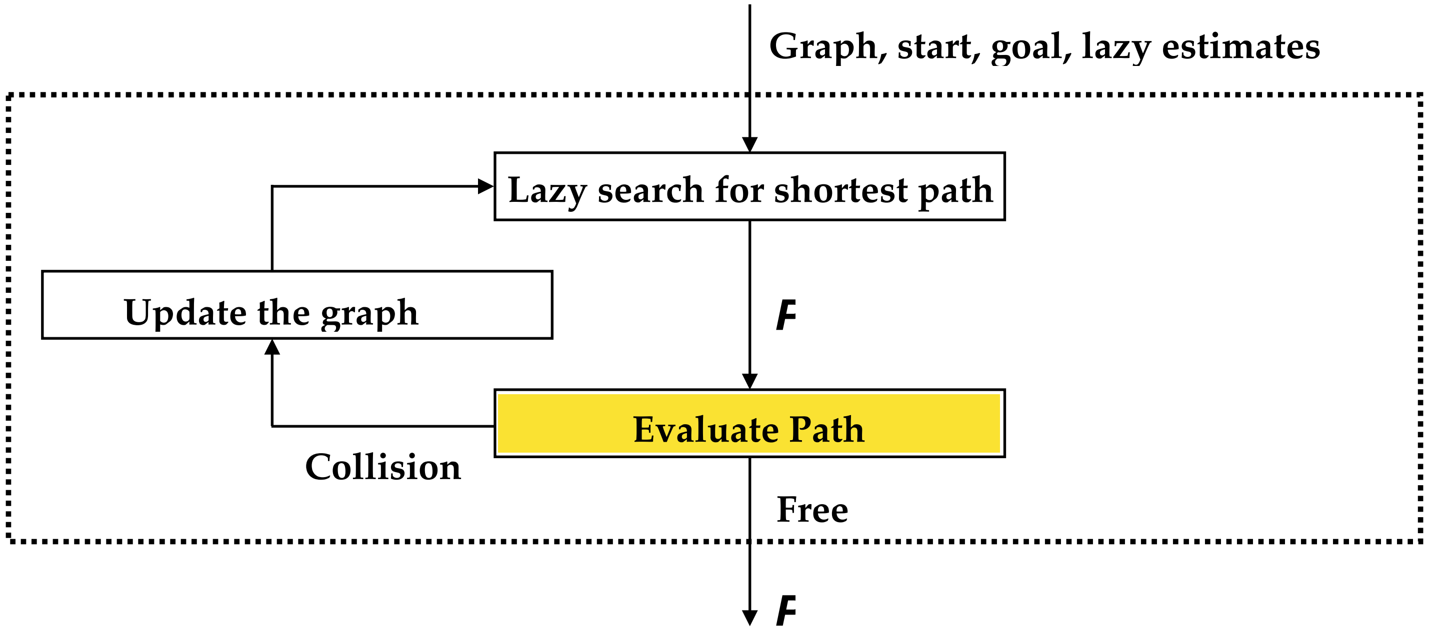

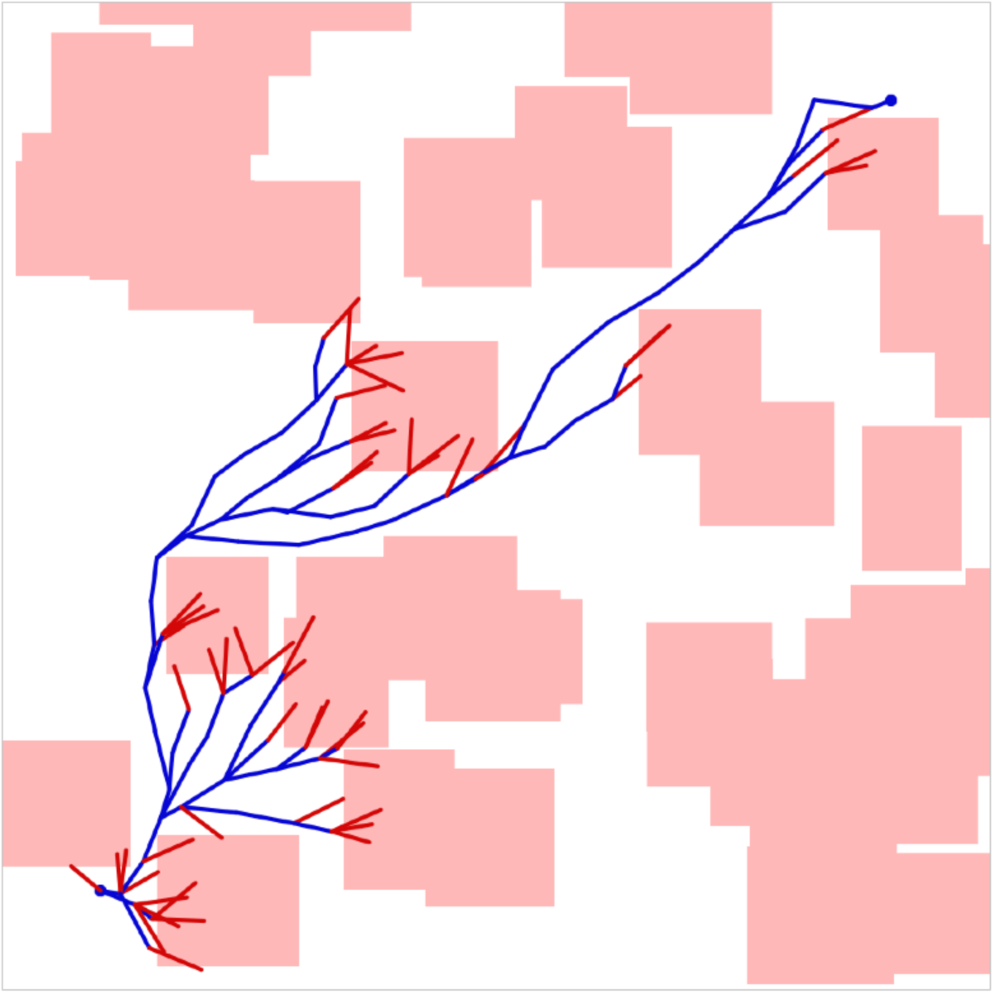

LazySP

Optimism Under Uncertainty

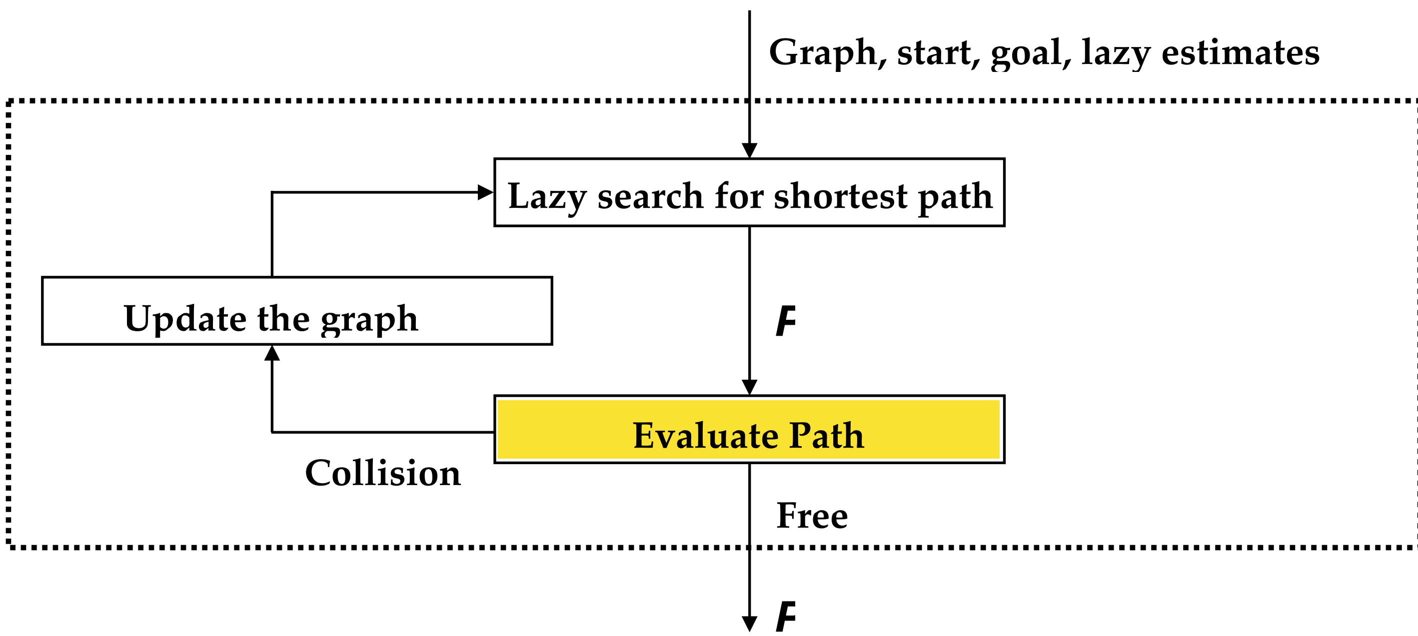

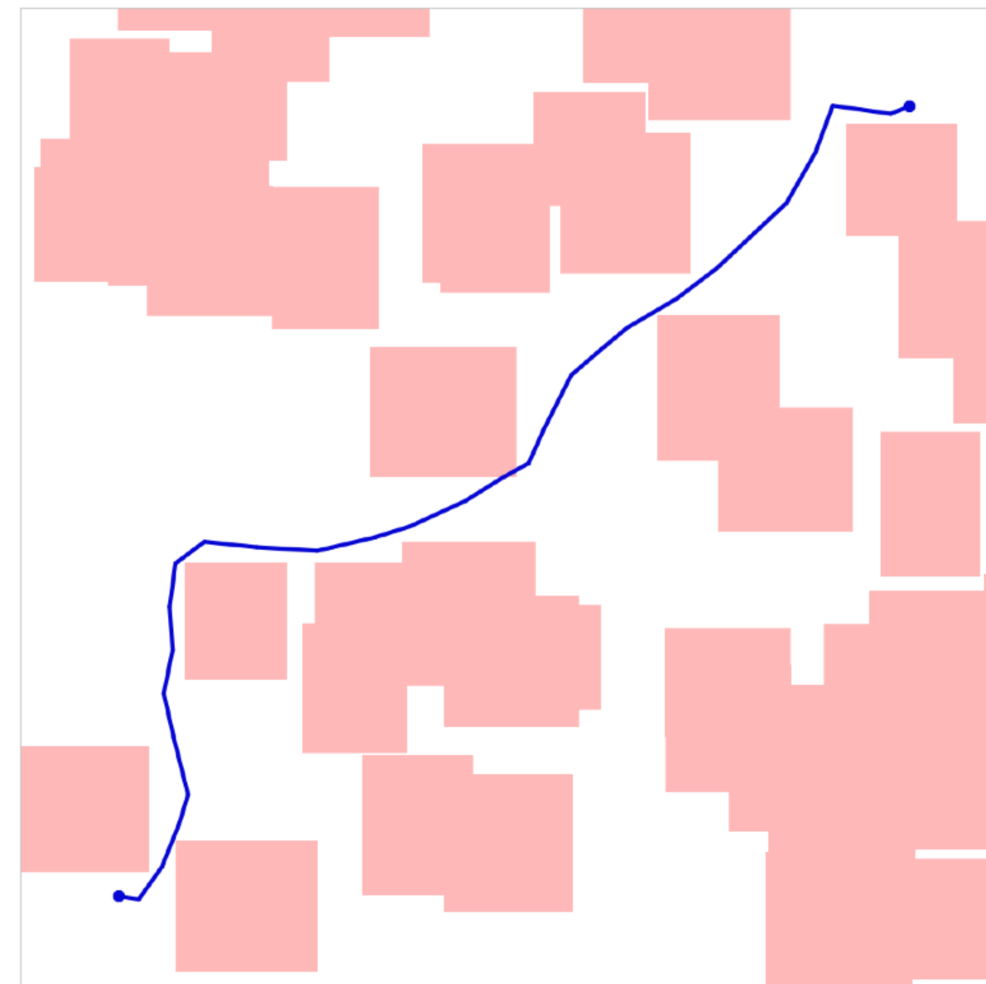

Edge Selectors

Forward
(first unevaluated edge)

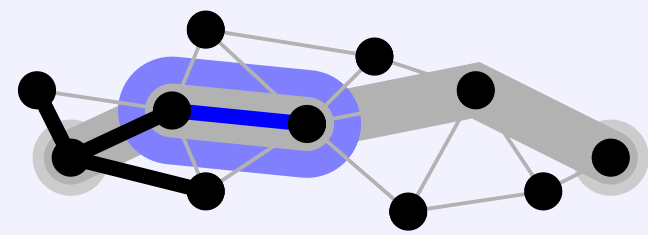

Reverse
(last unevaluated edge)

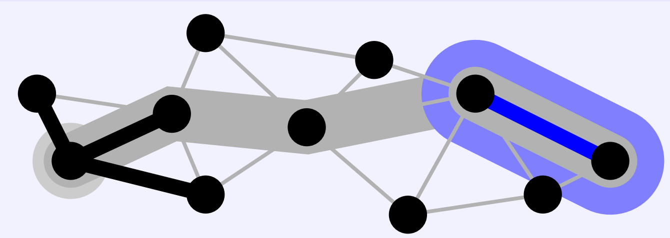

Alternate
(alternate Forward and Reverse)

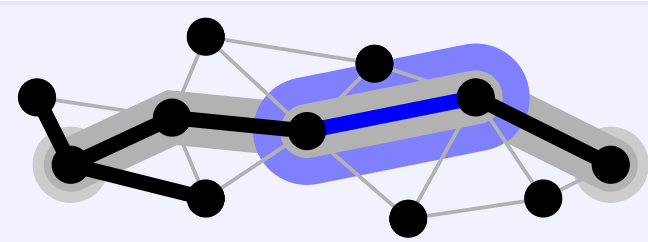

Bisect
(furthest from an unevaluated edge)

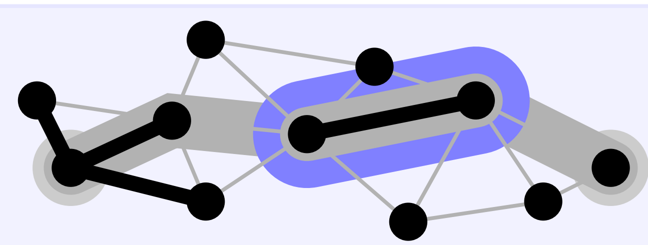

The Oracle is a LazySP Selector!

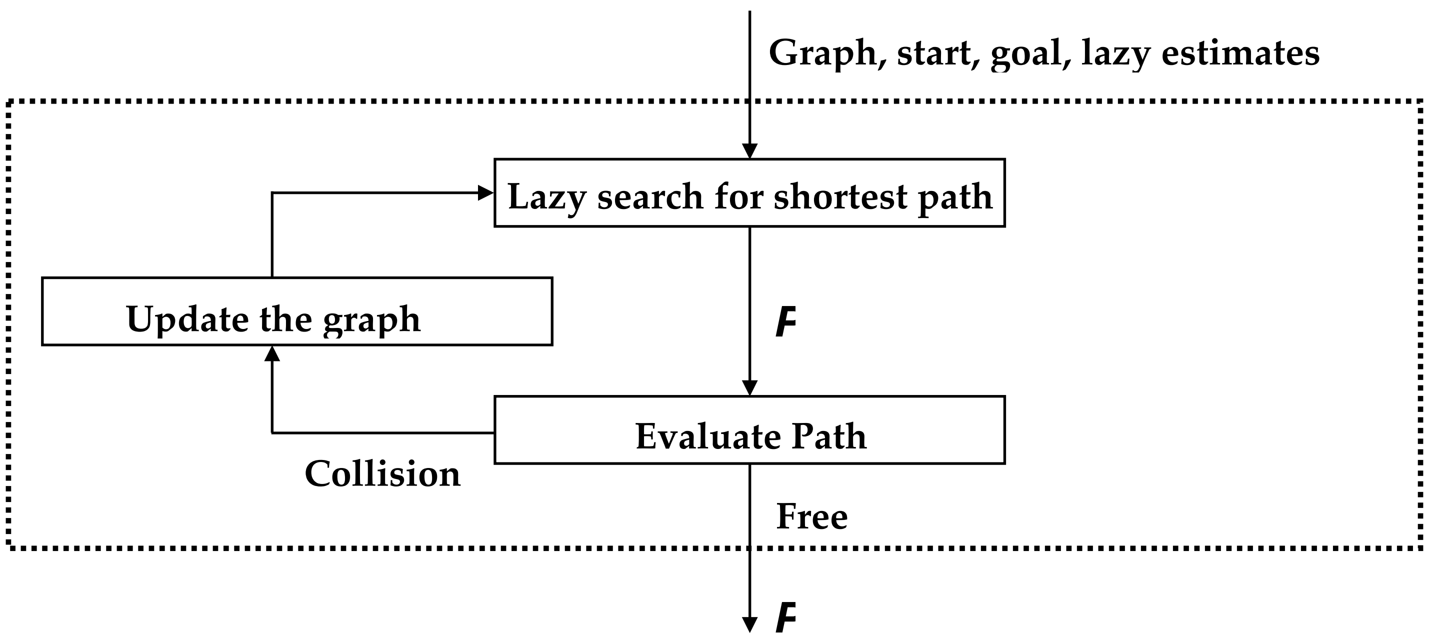

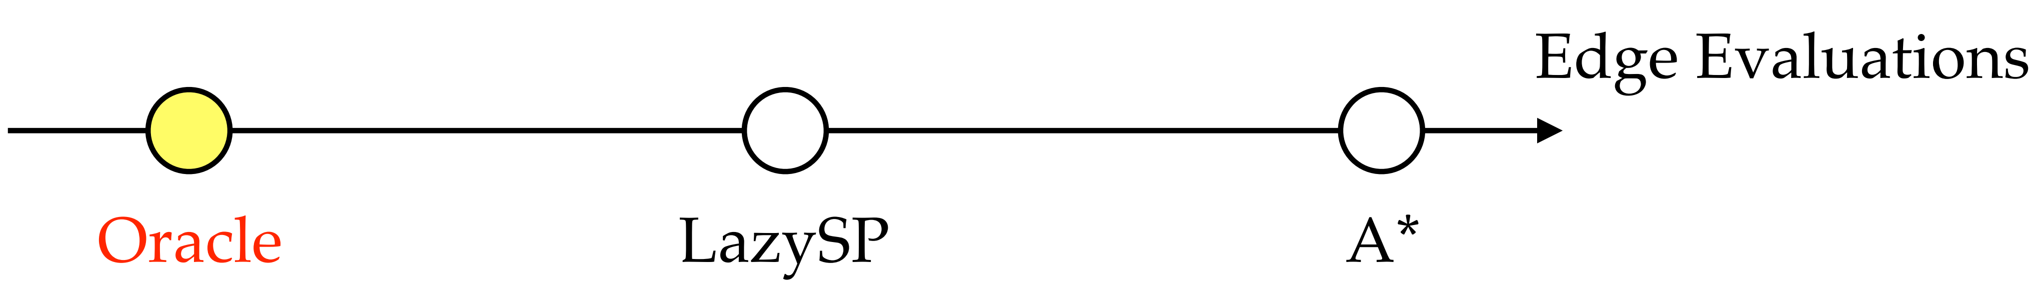

The Story Continues …

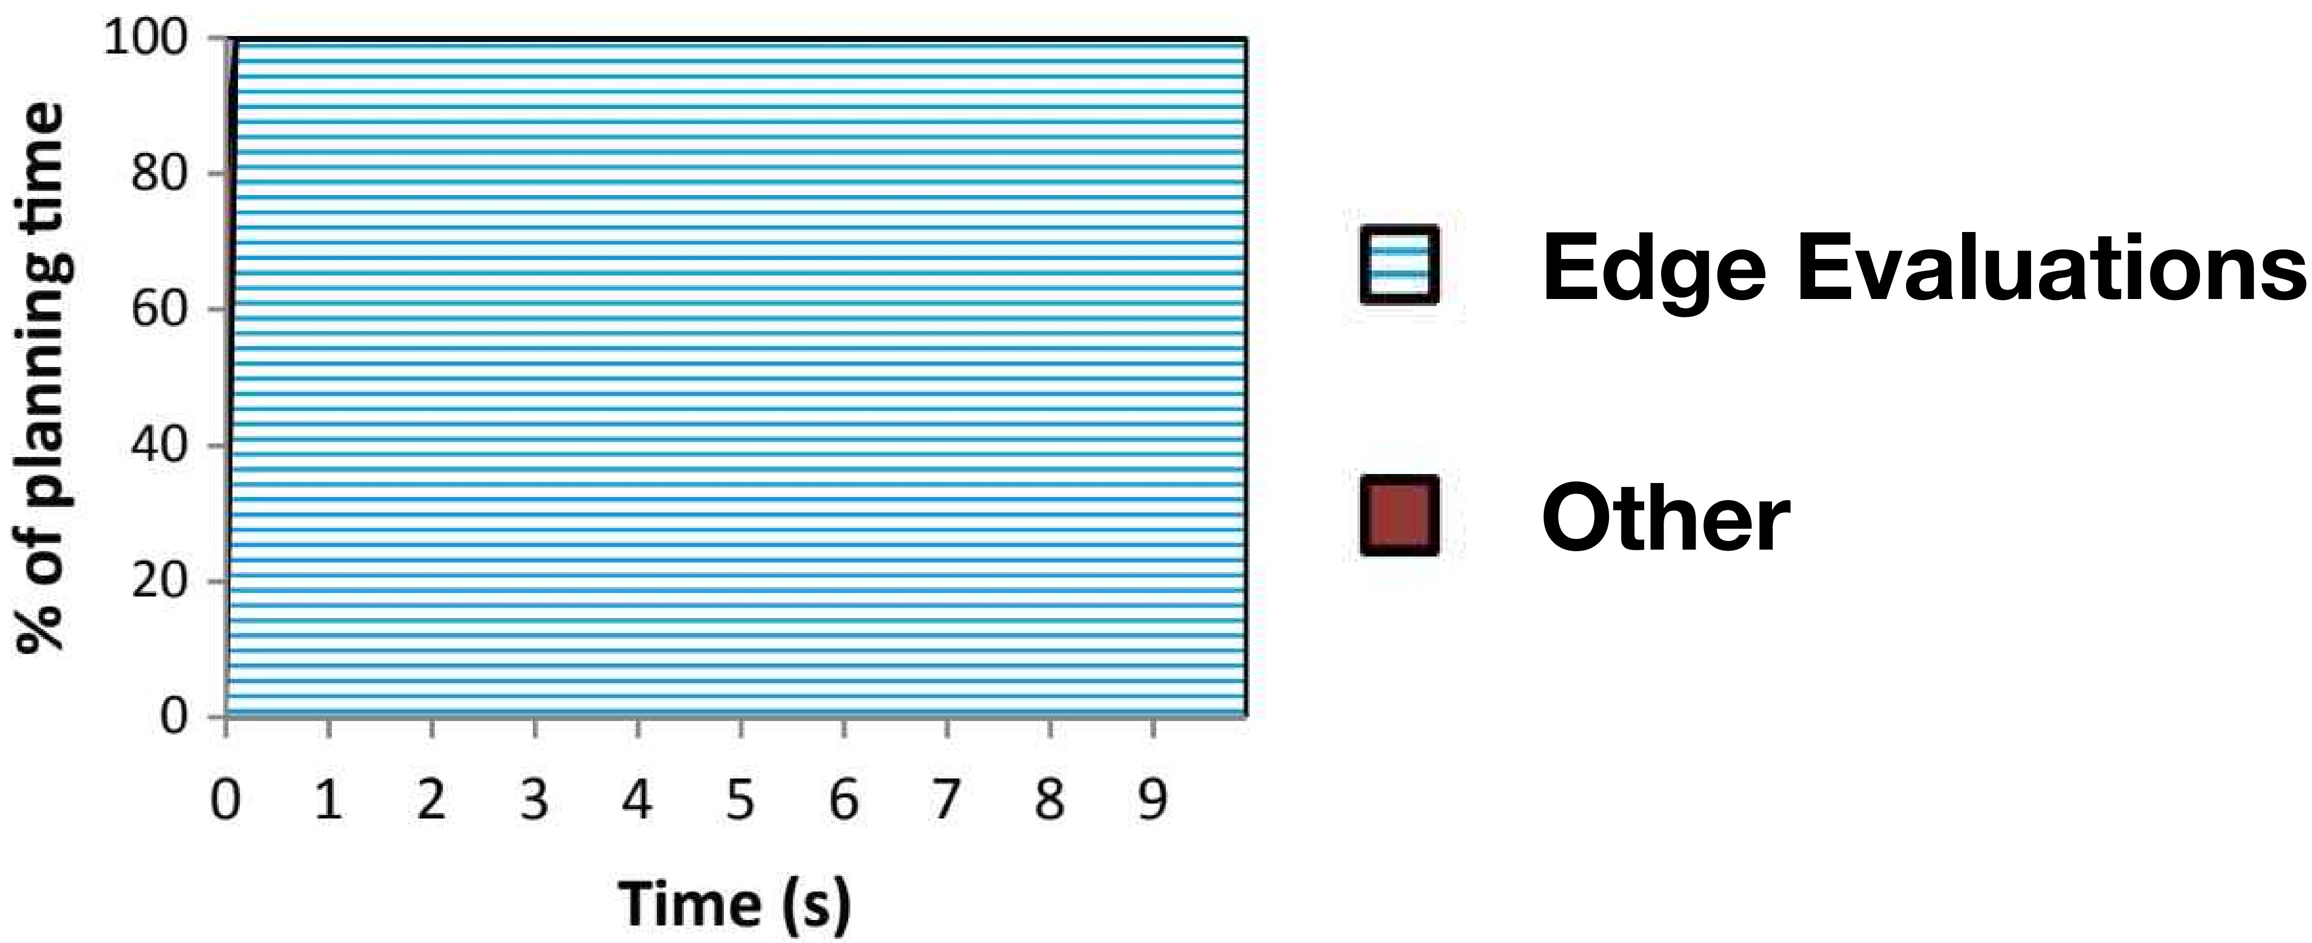

Hauser, Kris., Lazy collision checking in asymptotically-optimal motion planning. *ICRA* 2015

The Story Continues …

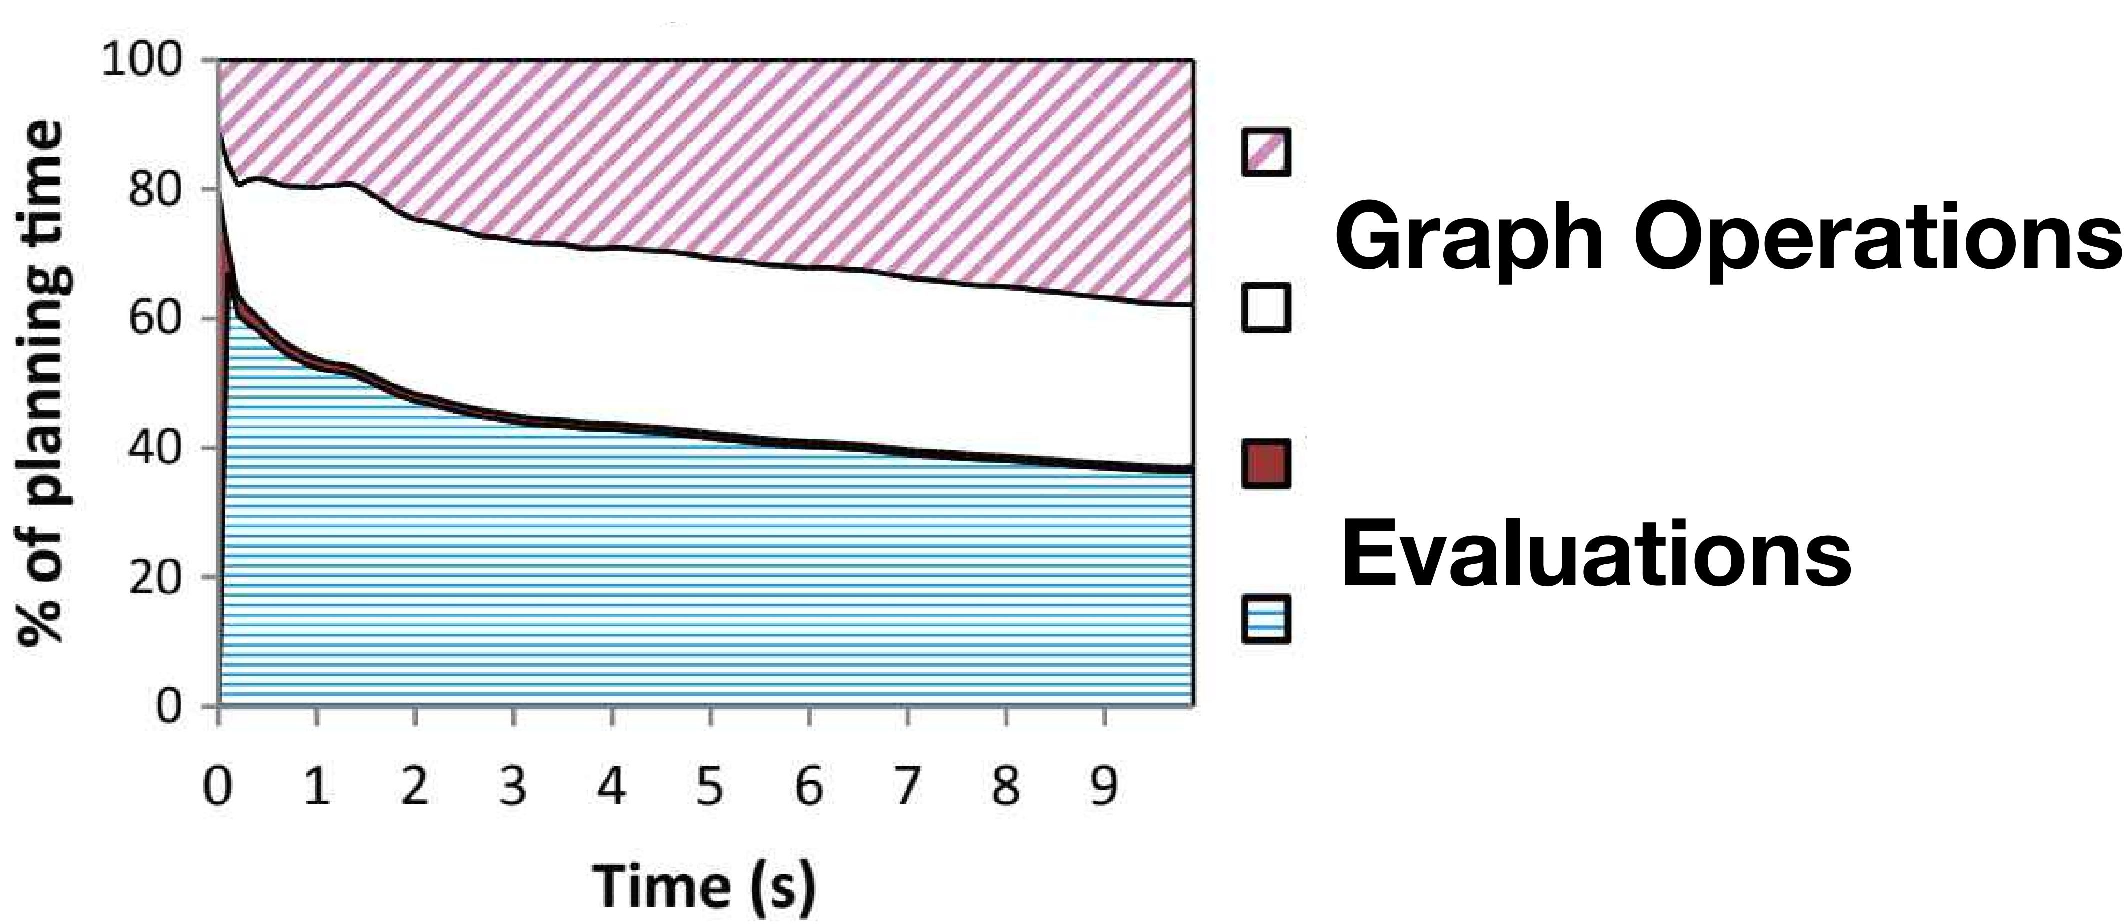

Is there a Search Algorithm
that Minimizes
the Number of Edge Evaluations?

LazySP

ICAPS 2018 [Best Conference Paper Award Winner]

First Provably Edge-Optimal A*-like Search Algorithm

The Piano Movers' Problem

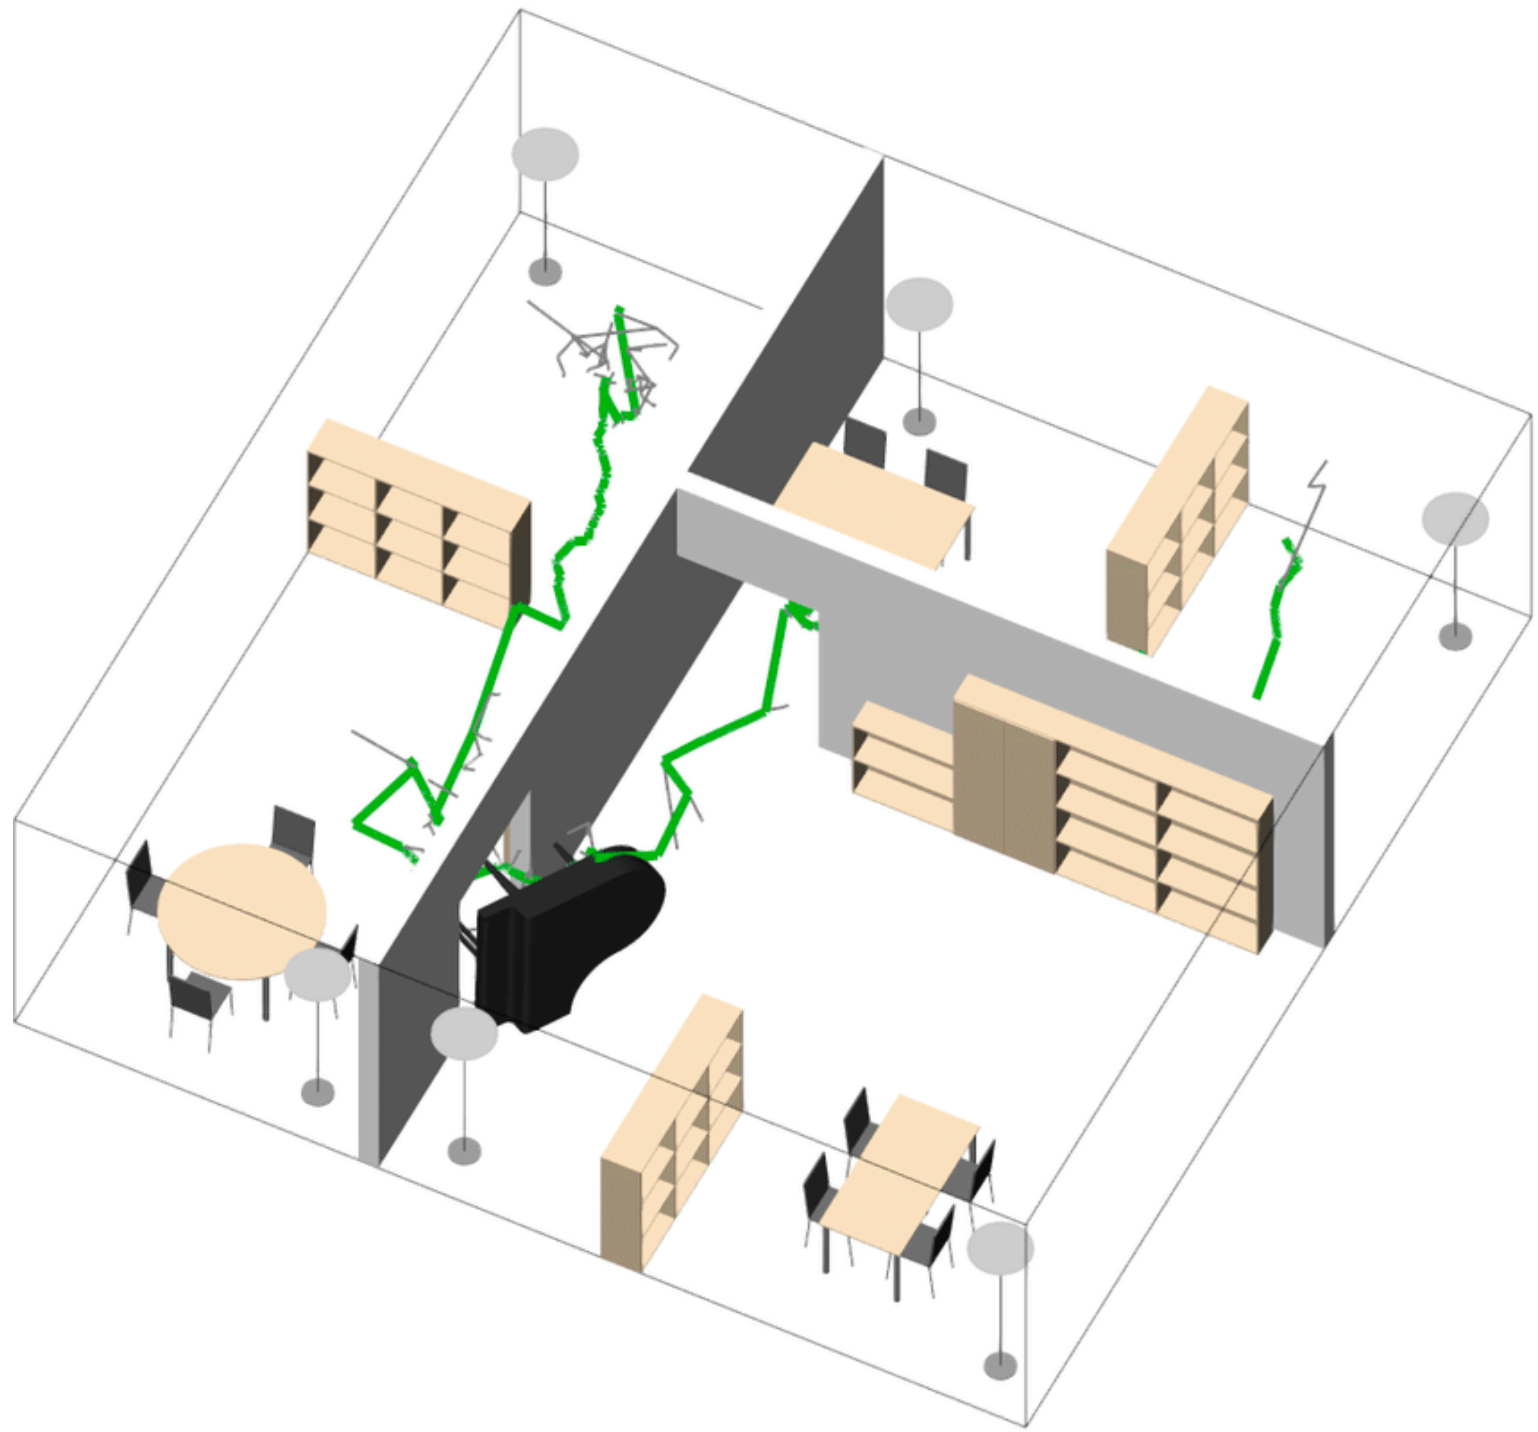

The Experienced Piano Movers' Problem

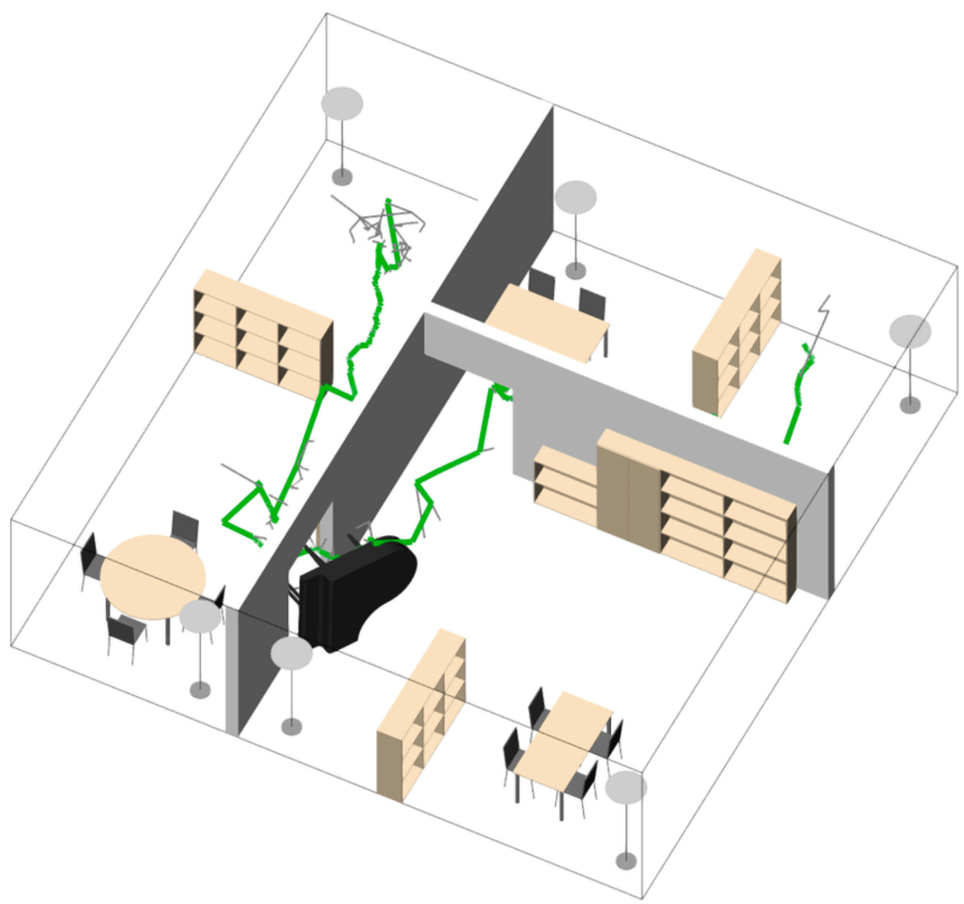

New Piano.
New House.
Same Mover.

A Bayesian Approach to Edge Evaluation

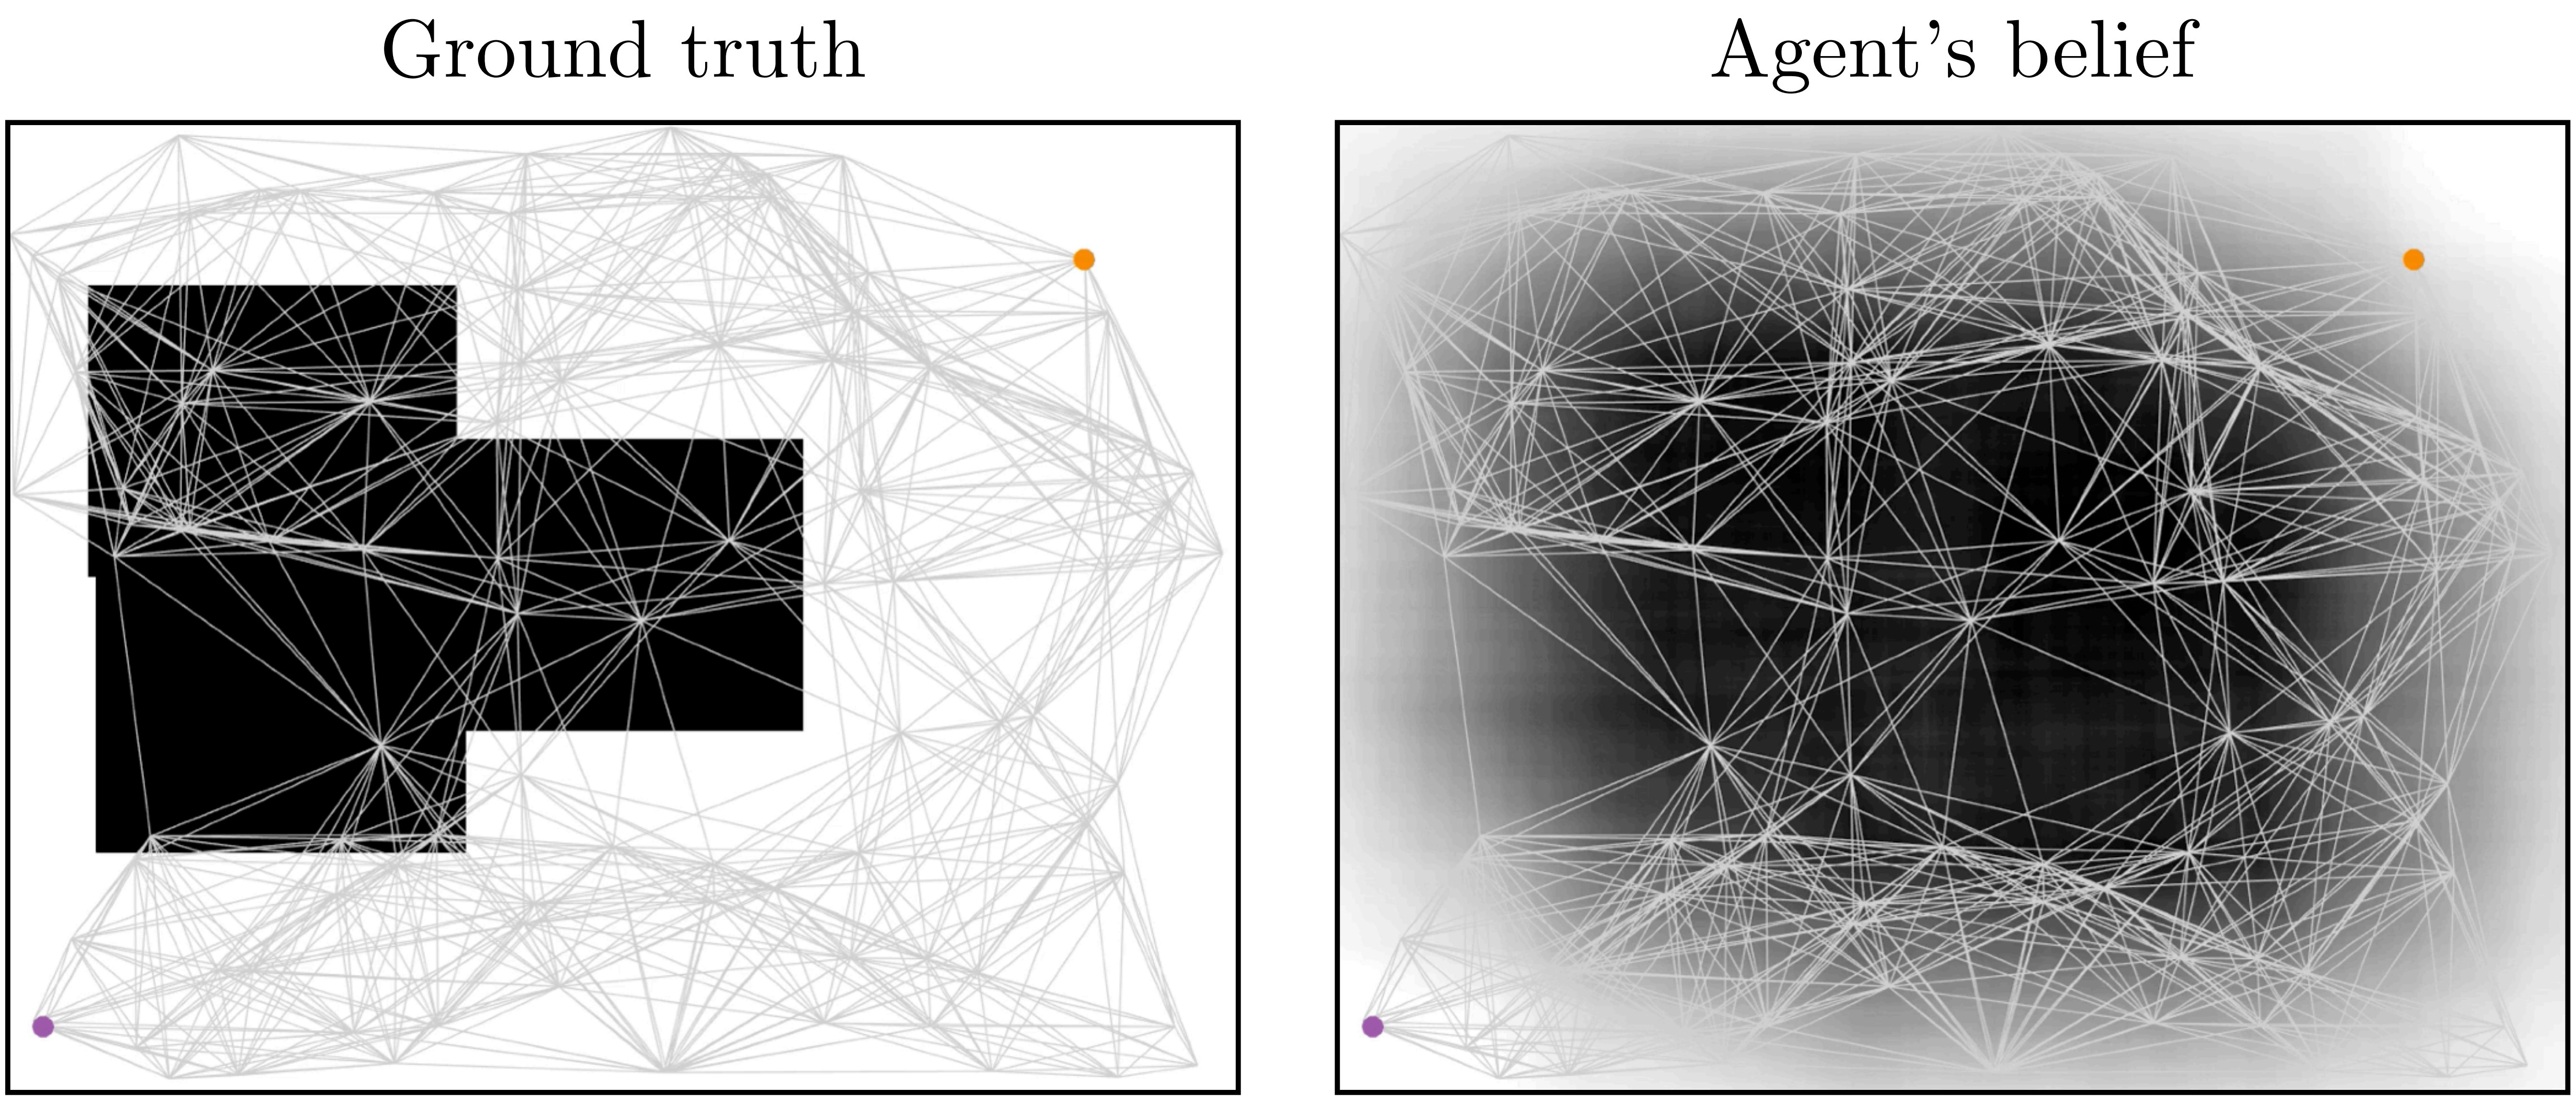

Approach: Offline decision tree via DIRECT

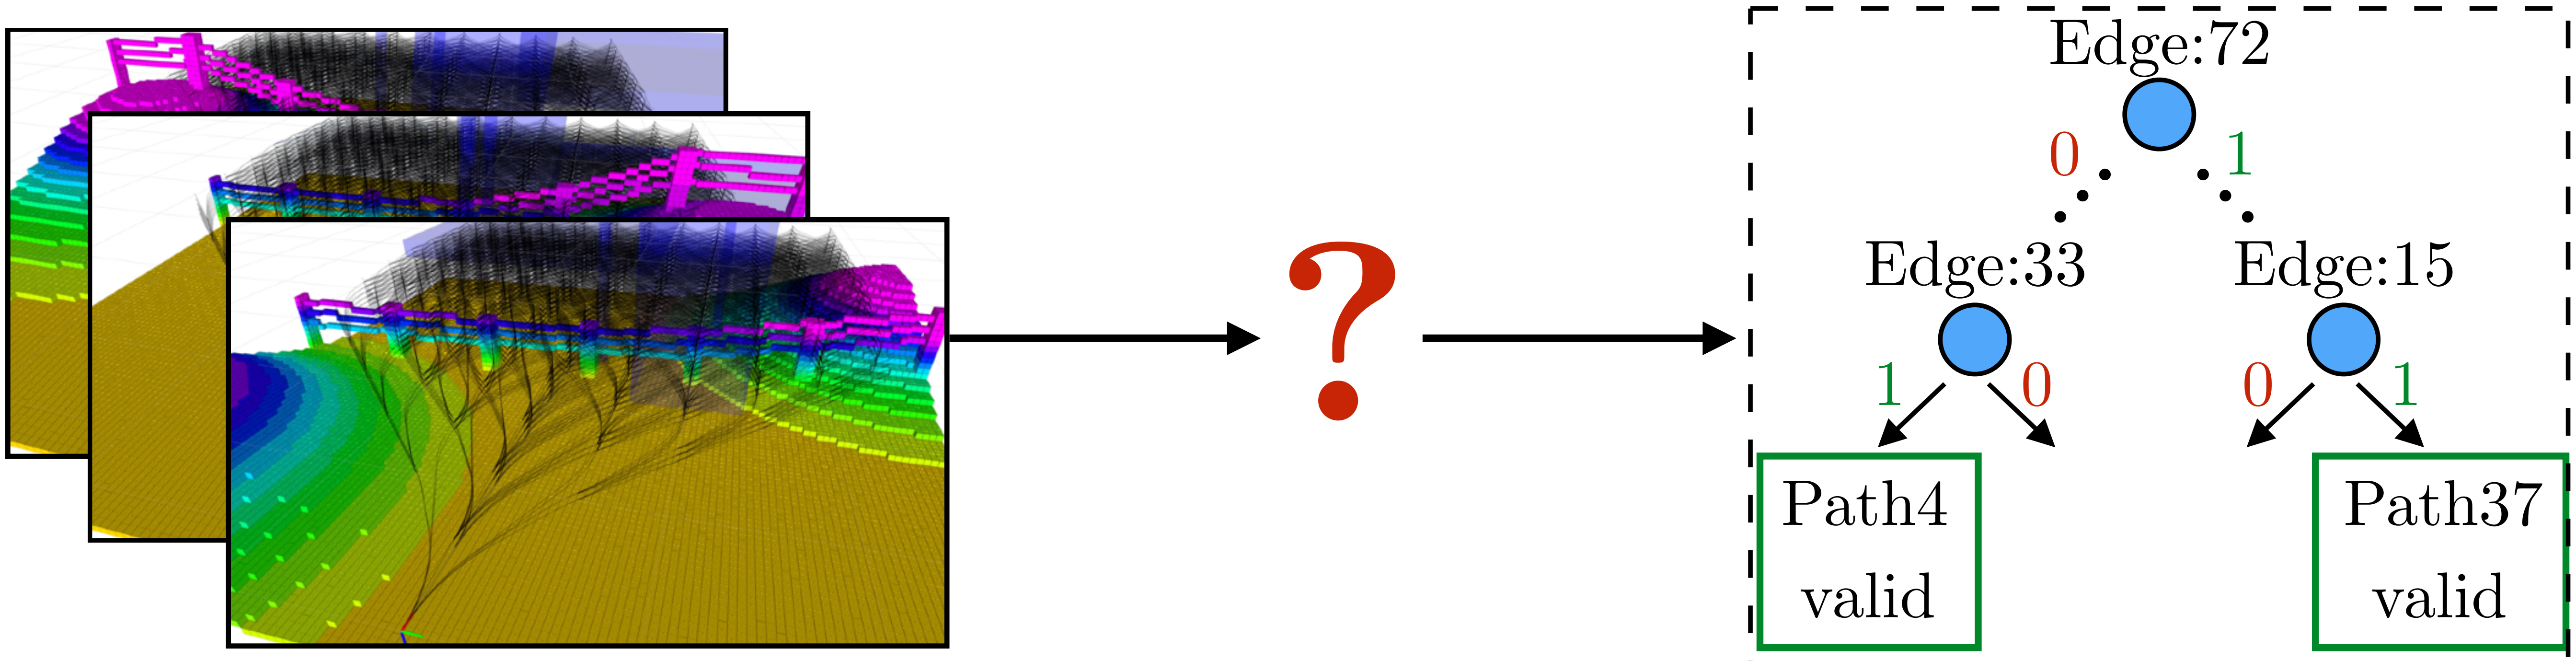

Dataset

Problem is NP-Hard (Javdani et al. '14)

DIRECT (Chen et al '15) frames it as a graph cut problem

Greedy policy ($\mathcal{O}(|E|\ |\Xi|\ |\Phi|)$) is near-optimal

Executing DIRECT online is expensive, cached as a decision tree

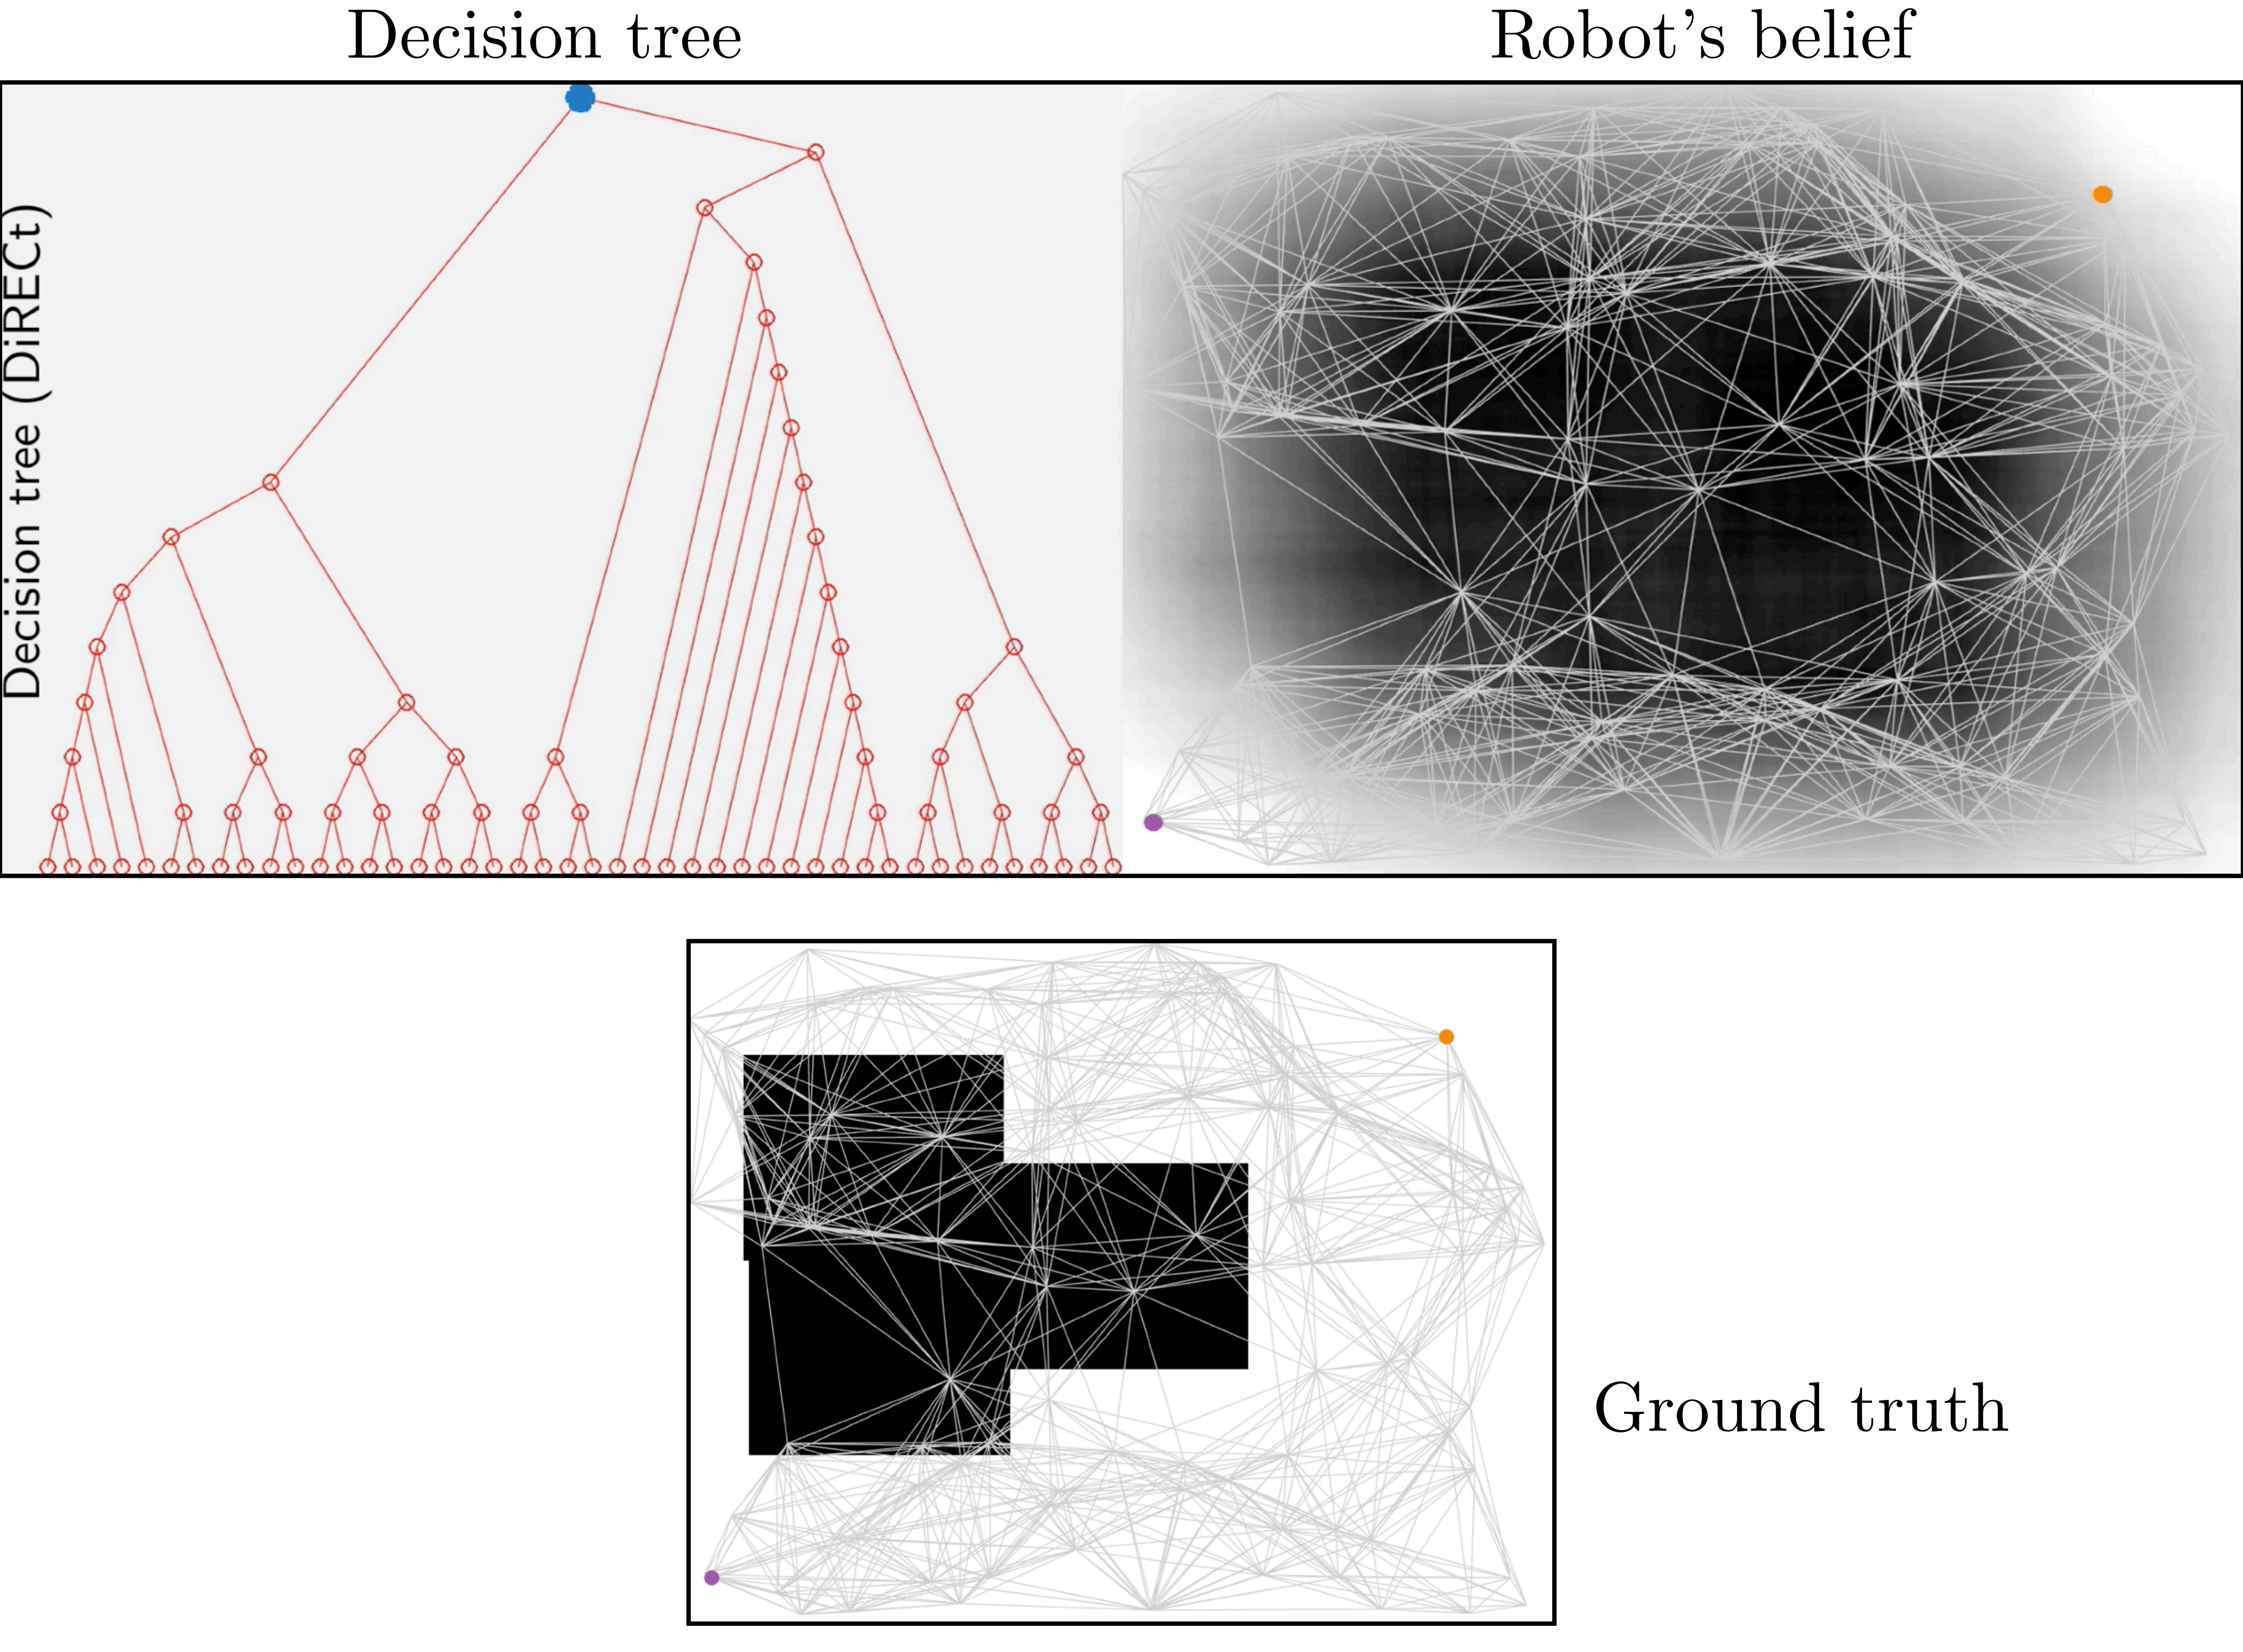

Decision tree
Robot's belief
Decision tree (DiRECt)
Ground truth

Problem: What if test world not in dataset?

This can happen because we cannot enumerate all $\mathcal{O}(2^{|E|})$ worlds.

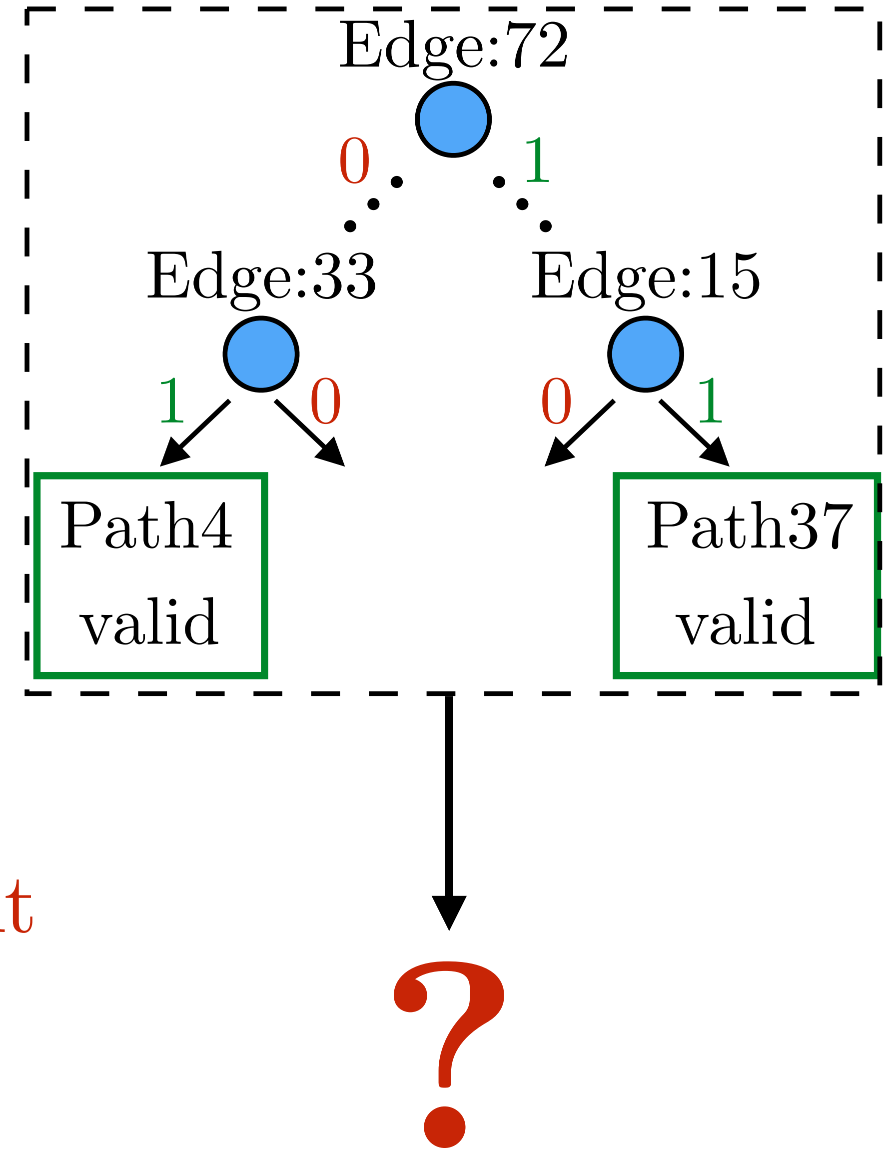

Solution: We developed BISECT [NIPS'17] that can reason about all worlds implicitly.

1. BISECT assumes edges are independent Bernoulli r.v.

2. BISECT has linear complexity: $\mathcal{O}(|E||\Xi|)$

3. We execute BISECT from the leaf of the tree

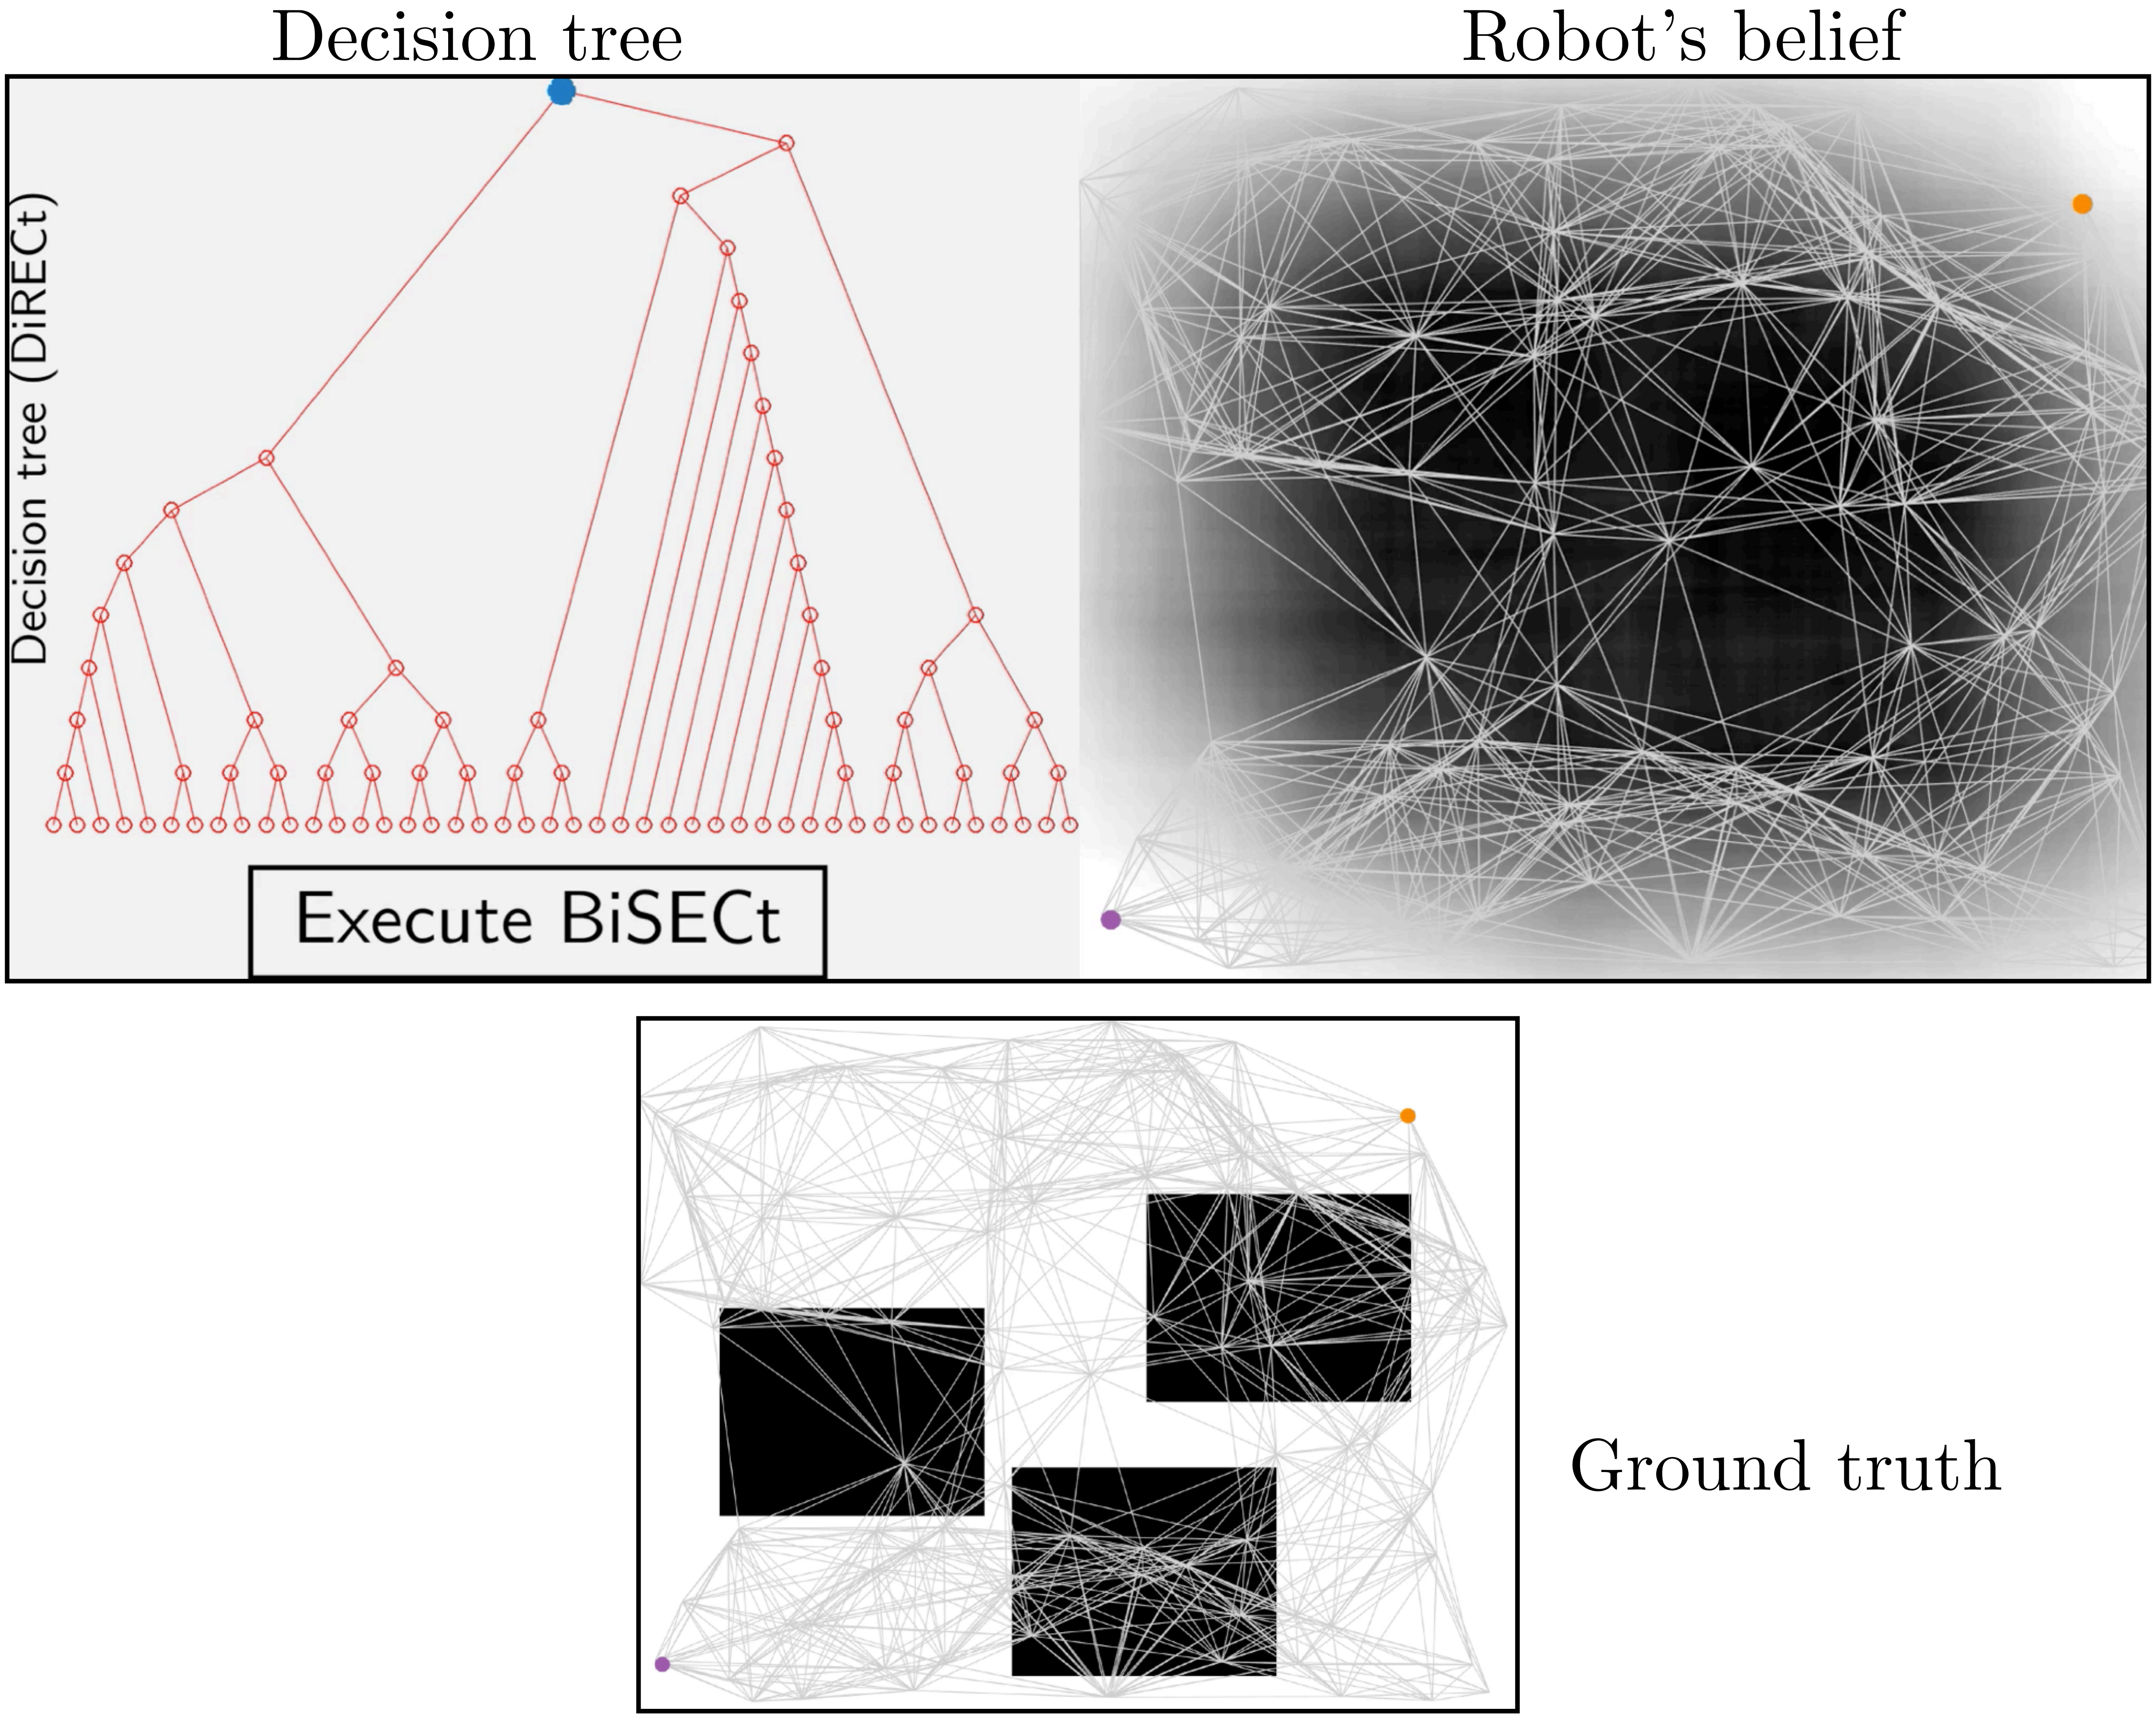

Decision tree
Robot's belief
Decision tree (DiRECt)
Execute BiSECt
Ground truth

DIRECT + BISECT exploits structure

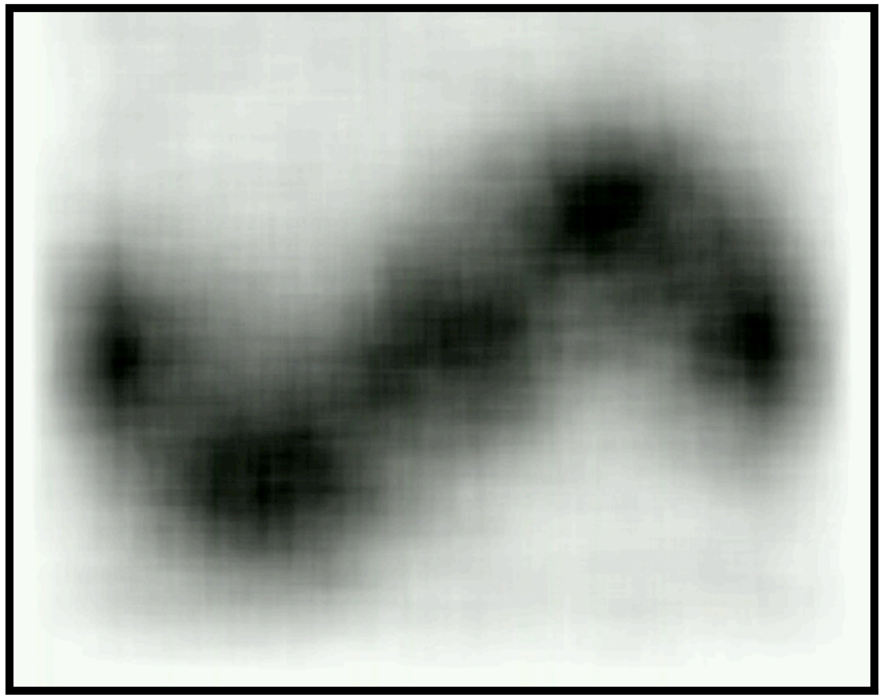

DIRECT + BISECT (14)

DIRECT + BISECT (24)

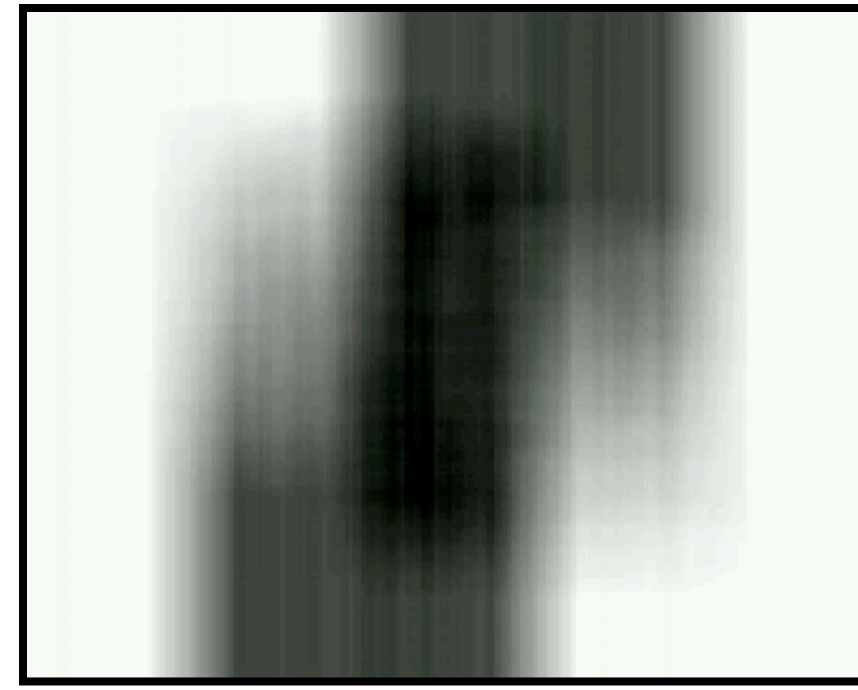

DIRECT + BISECT (23)

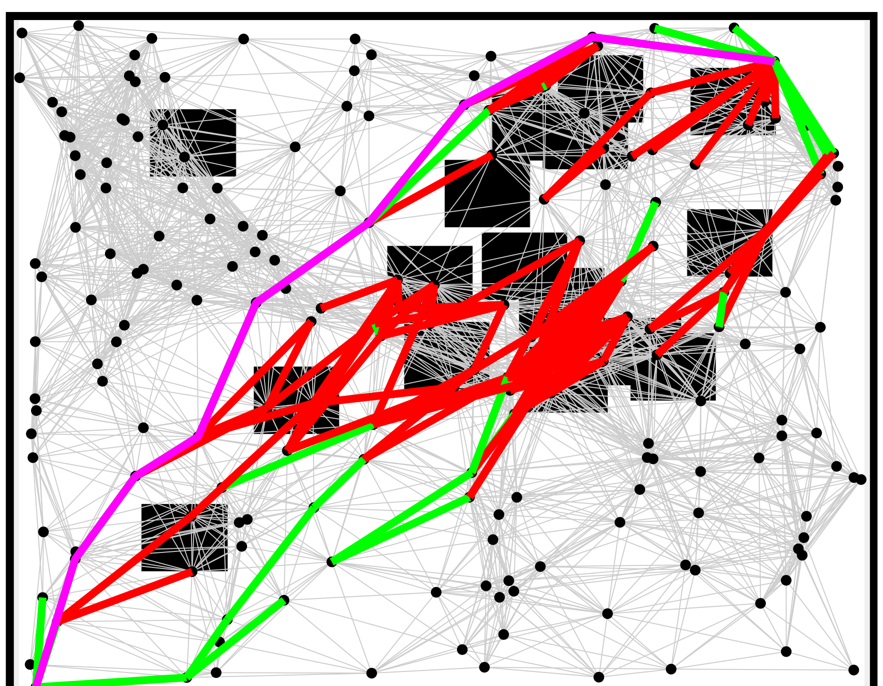

LazySP (155)

LazySP (838)

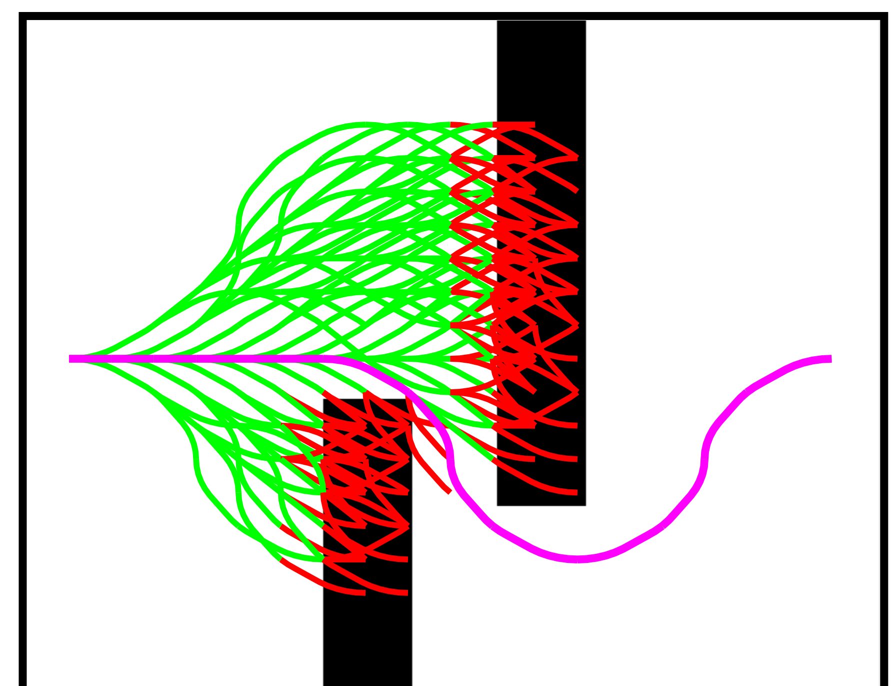

LazySP (251)

The Piano Movers' Problem

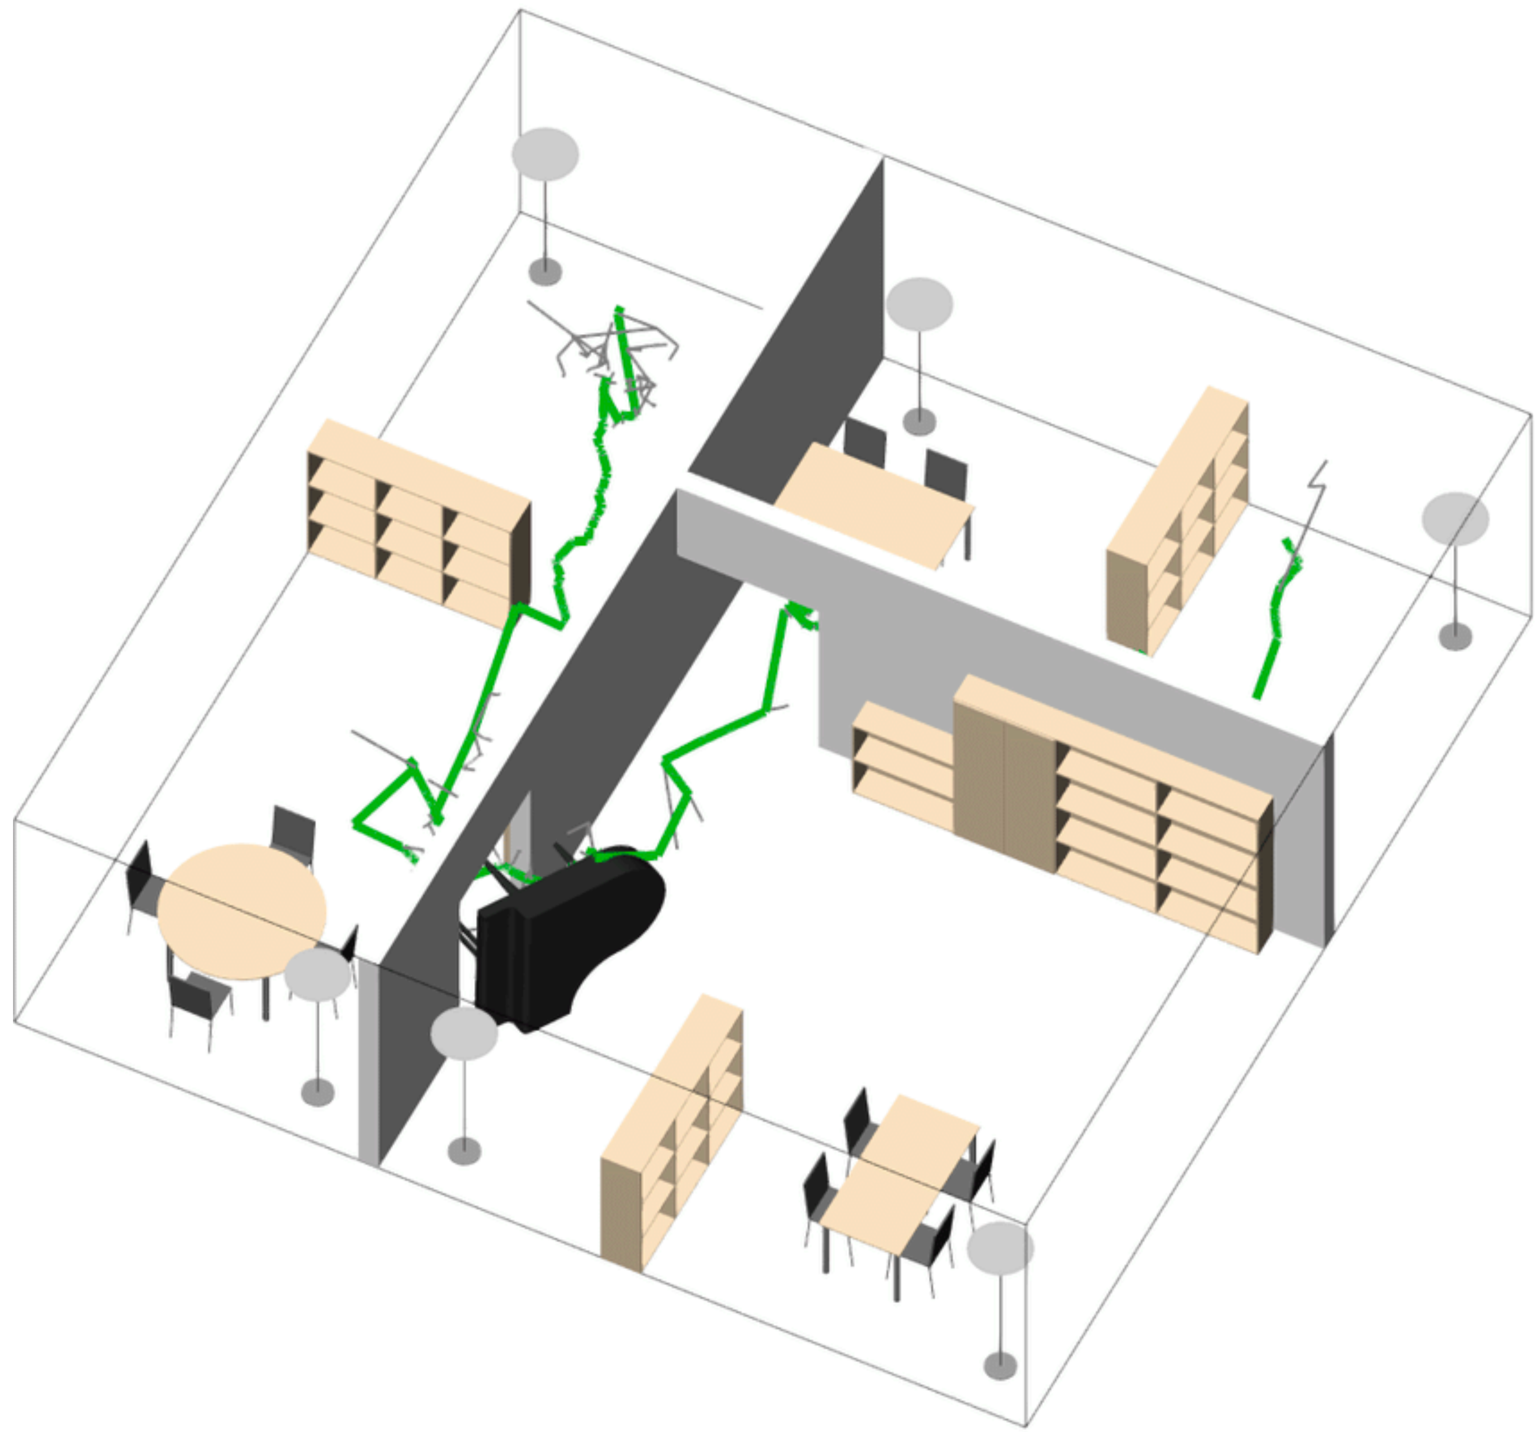

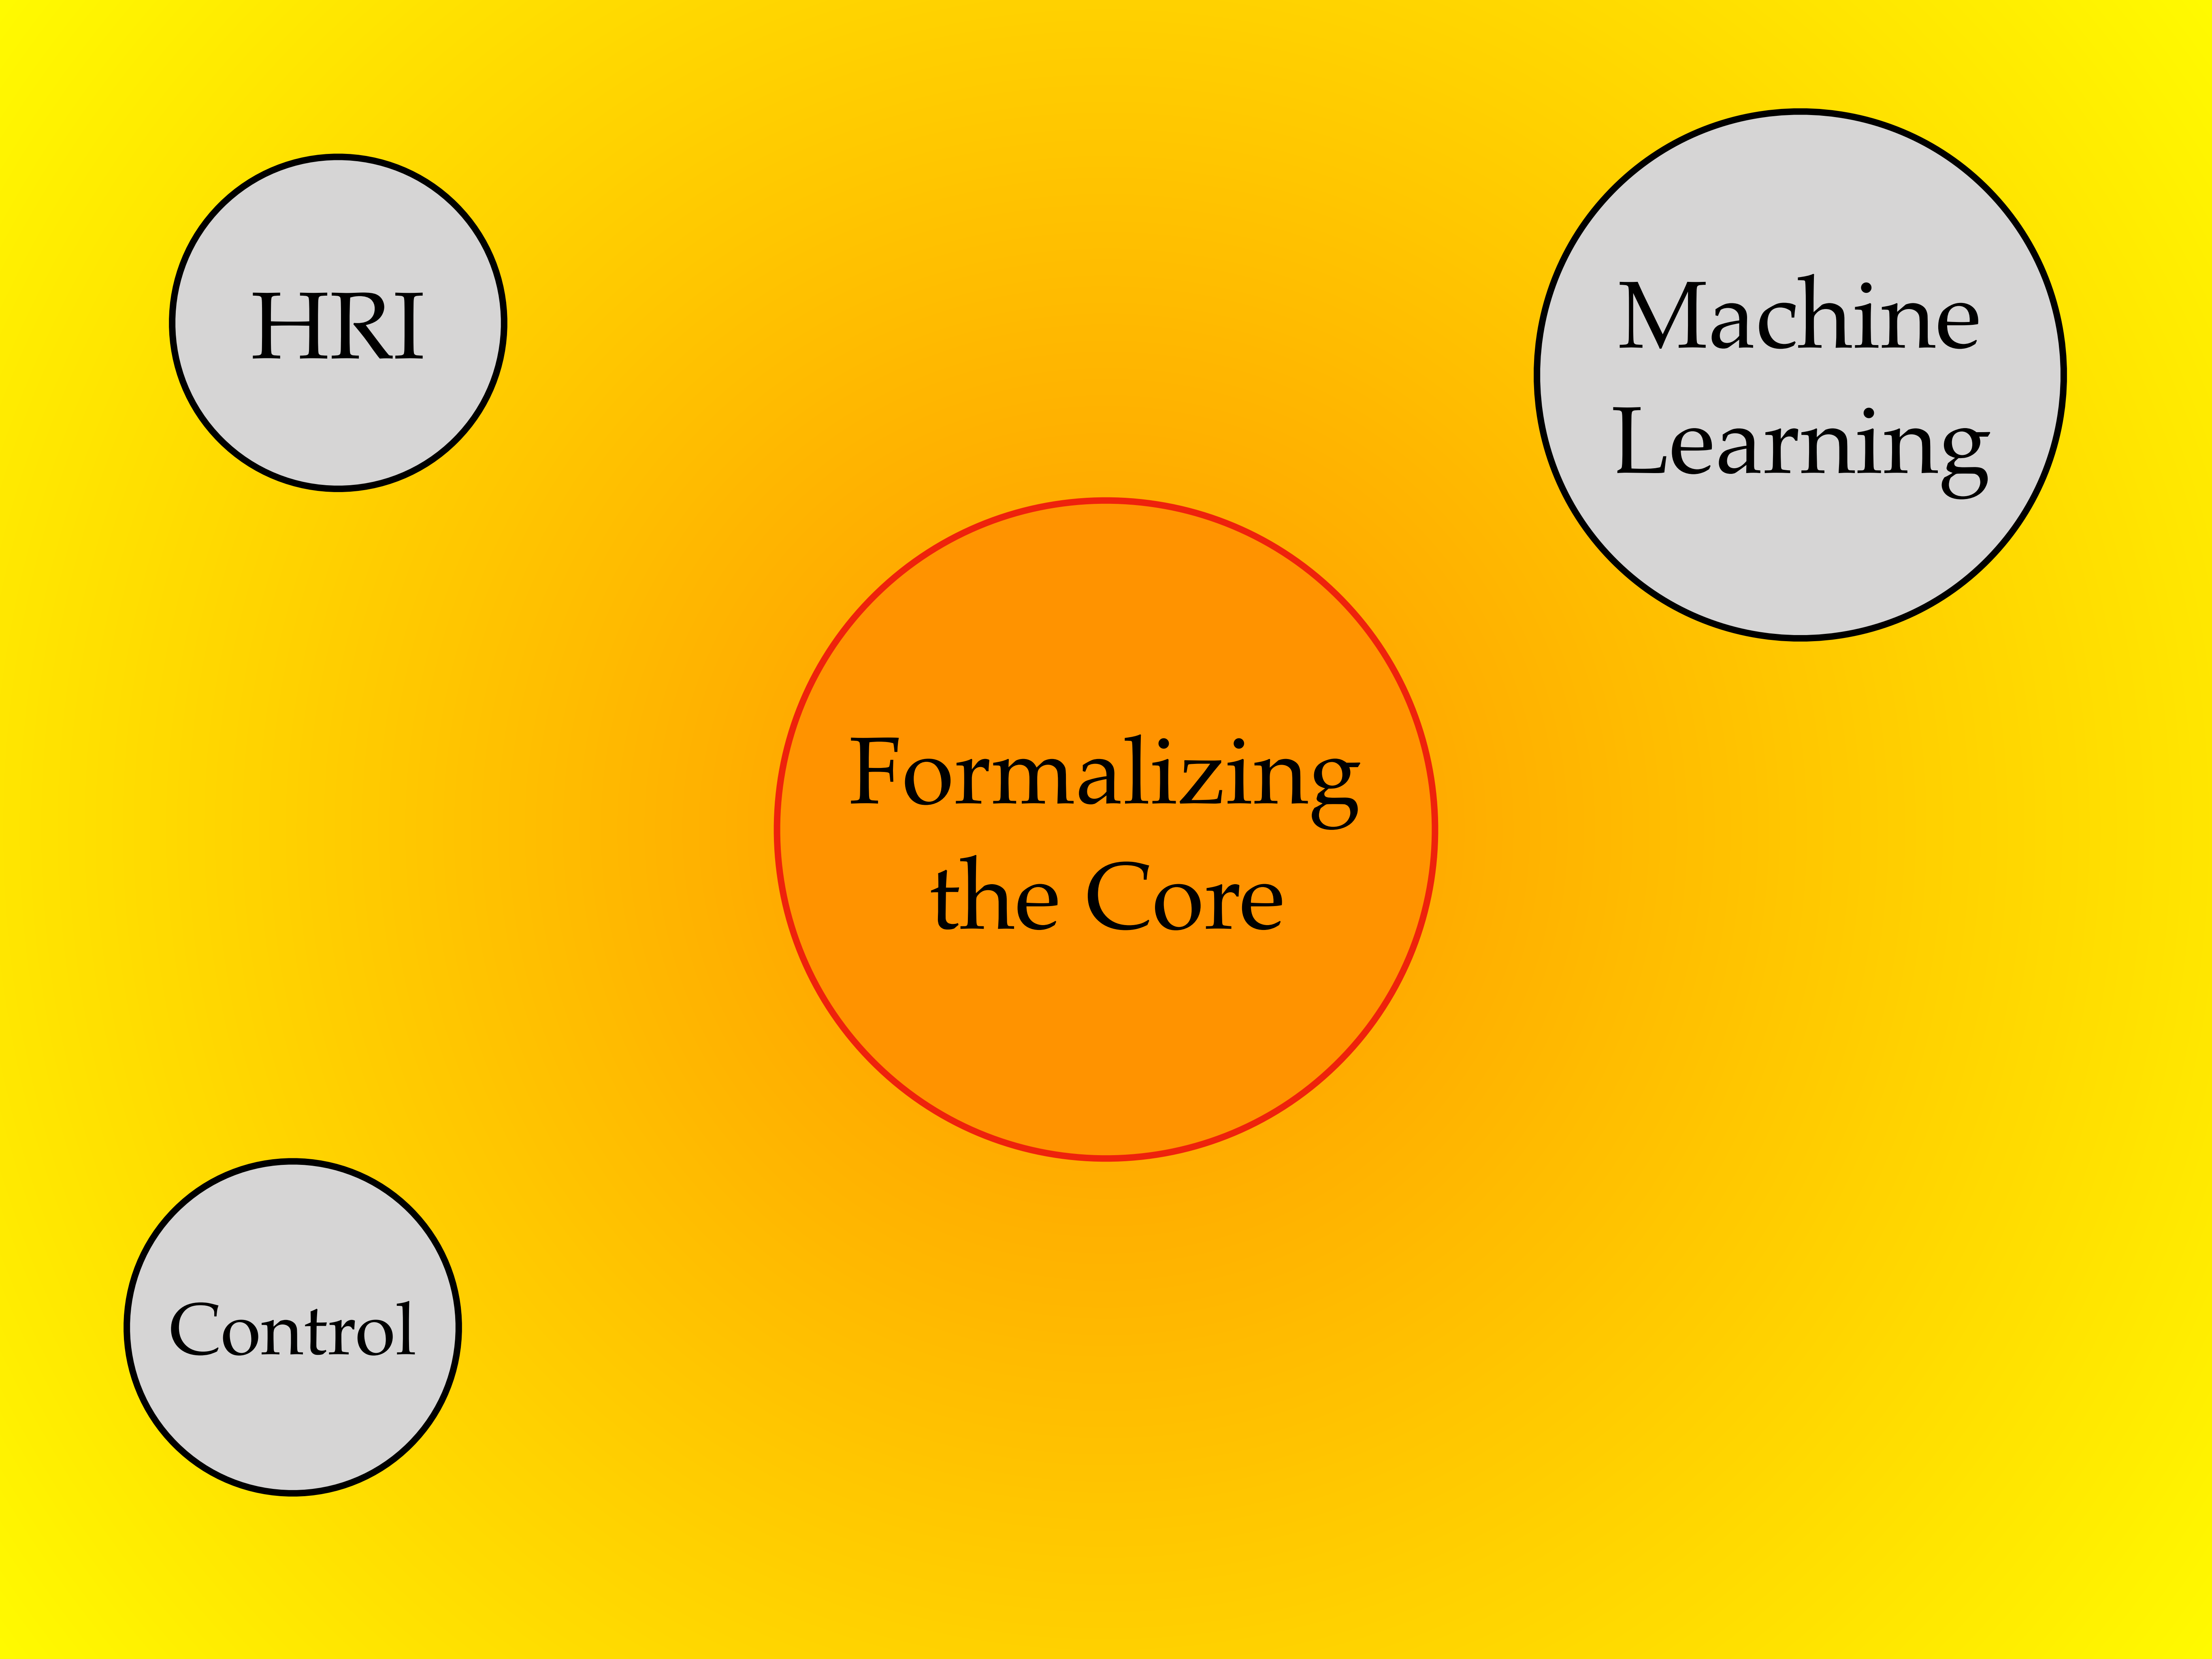
HRI
Machine Learning
Formalizing the Core
Control

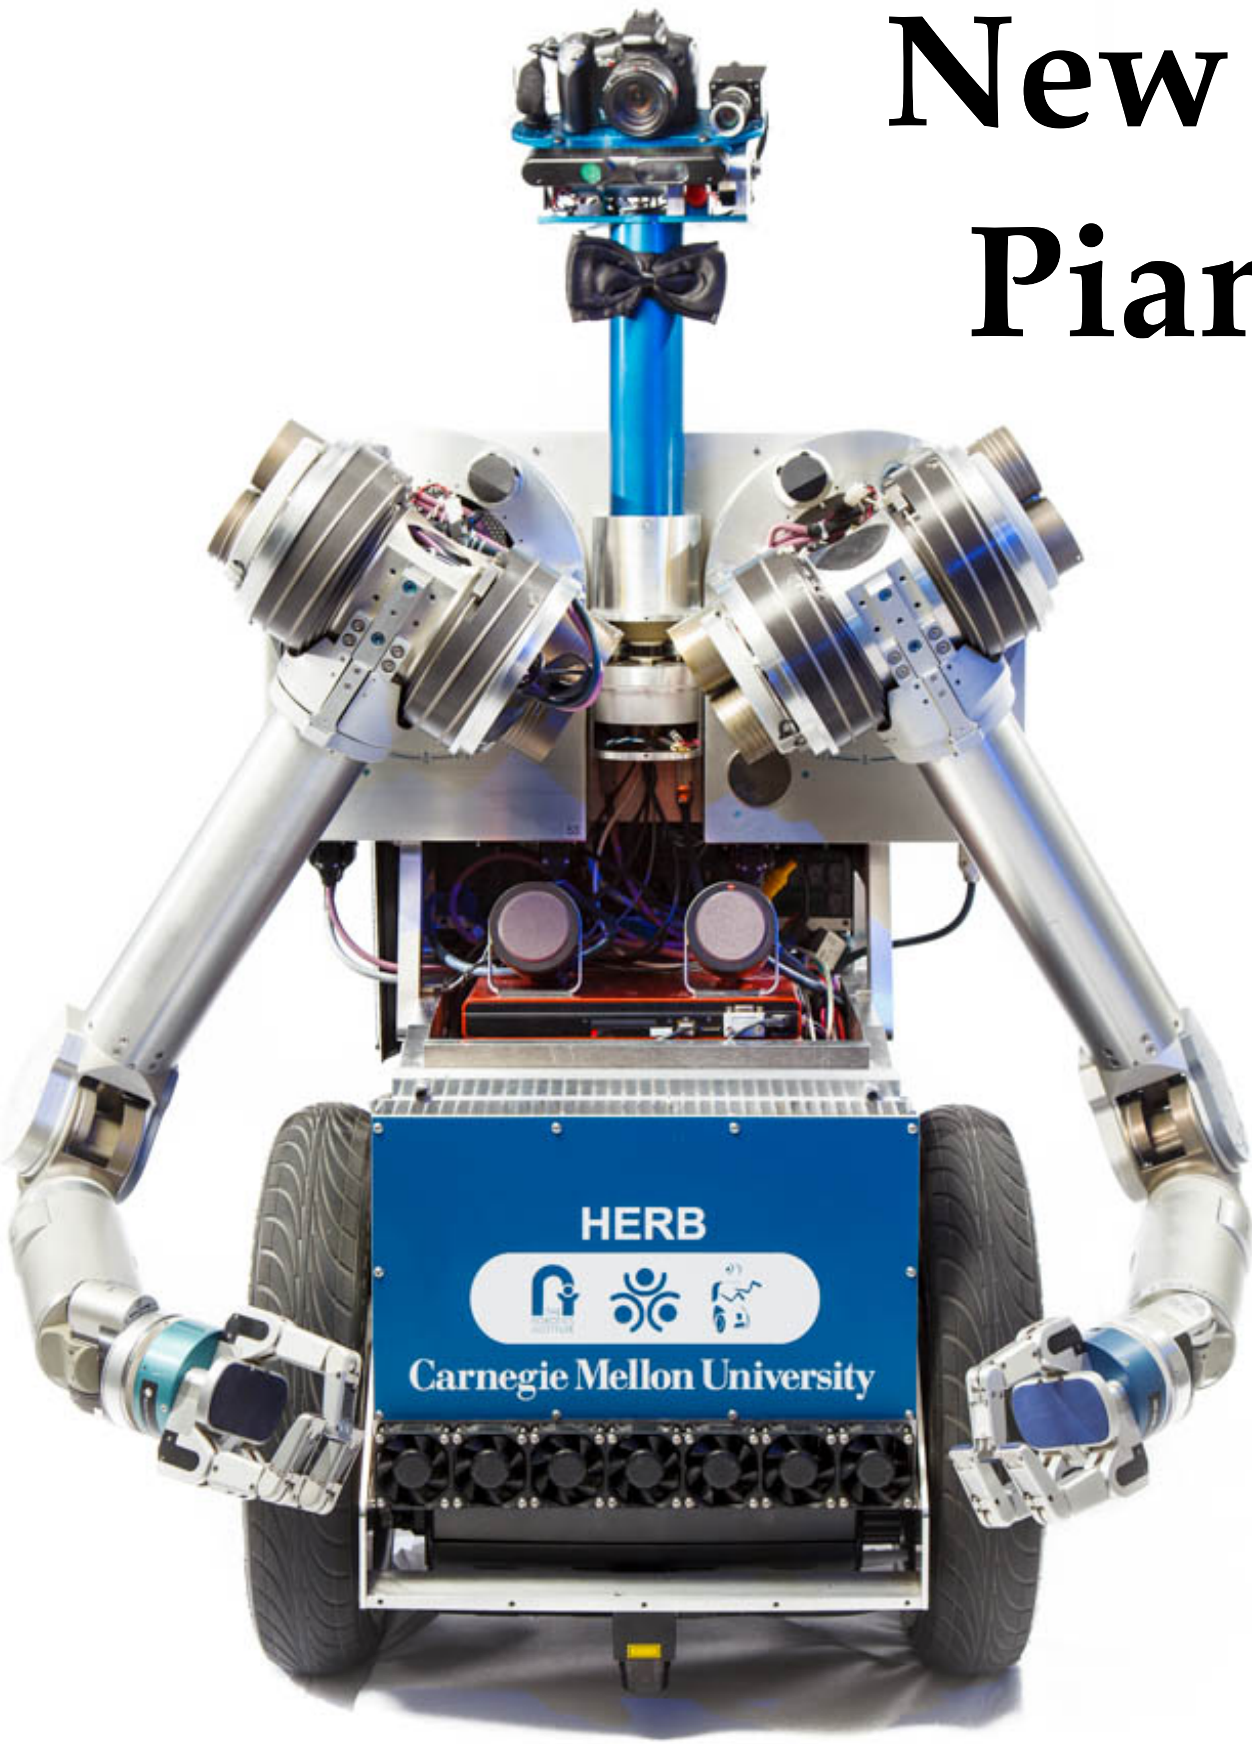

New Perspectives on the Piano Movers' Problem

Siddhartha Srinivasa
Boeing Endowed Professor
Personal Robotics Lab
University of Washington