

# Learning in Search-based Planning for Robotics

*Maxim Likhachev*

*Associate Professor*

*Carnegie Mellon University and Waymo*

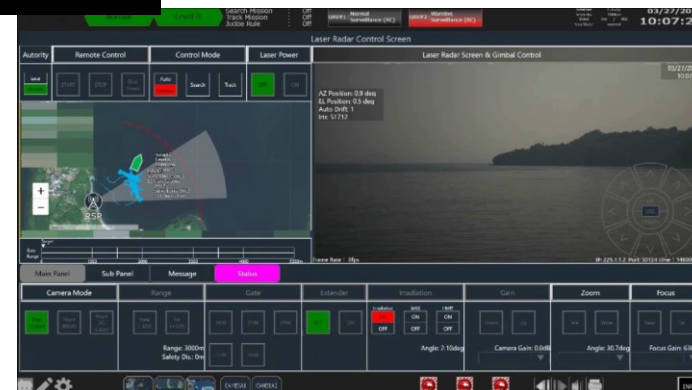
*Search-based Planning Lab (SBPL)*

*Joint work with*

*F. Islam, O. Salzman, A. Vemula (and others)*


- Planning, Decision-making and Learning in robotic systems
- General algorithmic methods with rigorous theoretical guarantees
- Applications to real-world robotic problems/systems






- Generate a systematic graph representation of the planning problem
- Search the graph for a solution with a heuristic search (e.g., A\* search)
- Can interleave the construction of the representation with the search (i.e., construct only what is necessary)

2D grid-based graph representation for 2D  $(x,y)$  search-based planning:




Lattice-based graph representation for 3D  $(x,y,\theta)$  planning:




- Generate a systematic graph representation of the planning problem

*Pros:*

- rigorous guarantees on solution quality
- good cost minimization
- consistency in solution  
(similar queries = similar solutions)




Lattice-based graph representation for 3D  $(x, y, \theta)$  planning:




- Generate a systematic graph representation of the planning problem

## Challenges:

- high-dimensionality/graph size
- expensive edge cost evaluation
- edge construction for dynamic systems
- reliance on the accuracy of the model



Lattice-based graph representation for 3D  $(x, y, \theta)$  planning:



- Challenges in Search-based Planning
- Constant-time Motion Planning (CTMP) – offline learning for online planning
- CMAX/CMAX++ for handling inaccurate models
- Summary and thoughts on research directions


- Challenges in Search-based Planning
- Constant-time Motion Planning (CTMP) – offline learning for online planning
- CMAX/CMAX++ for handling inaccurate models
- Summary and thoughts on research directions




- Generate a systematic graph representation of the planning problem

## Challenges:

- *high-dimensionality/graph size*
- *expensive edge cost evaluation*
- *edge construction for dynamic systems*
- *reliance on the accuracy of the model*




Lattice-based graph representation for 3D  $(x, y, \theta)$  planning:



## 5-D trajectory planning


$(x, y, z, \theta, v)$



## 12-D full-body planning

(3D base pose, 1D torso height,

6DOF object pose, 2 redundant DOFs in arms)



## 5-D trajectory planning

$(x, y, z, \theta, v)$




**over 500M states!**

(for 10km by 10km area discretized into 25m cubes, 32 yaw angles, 5 velocities)

## 12-D full-body planning

(3D base pose, 1D torso height,

6DOF object pose, 2 redundant DOFs in arms)



**over  $10^{19}$  states!**

(for small indoor space, joint angle resolution = 10 degrees)

5-D trajectory

path planning

height,

arms)

### *Methodologies to address it:*

- *implicit graphs*
- *compact graph representations including adaptive dimensionality [Gochev et al. 11]*
- *sub-optimal/anytime search [Pearl 84, Likhachev et al. 04, Hansen & Zhou 07, Thayer & Ruml 08,...]*
- *incremental planning, especially within sub-optimal/anytime search [Koenig & Likhachev 04, Likhachev et al. 08, ...]*
- *deriving multiple heuristics, each corresponding to a low-dimensional version of the problem, and using these via Multi-Heuristic A\* [Aine et al. 15]*

***over 500M states!***

*(for 10km by 10km area discretized into 25m cubes, 32 yaw angles, 5 velocities)*


***over  $10^{19}$  states!***

*(for small indoor space, joint angle resolution = 10 degrees)*


- Generate a systematic graph representation of the planning problem

## Challenges:

- high-dimensionality/graph size
- *expensive edge cost evaluation*
- edge construction for dynamic systems
- reliance on the accuracy of the model




Lattice-based graph representation for 3D  $(x, y, \theta)$  planning:






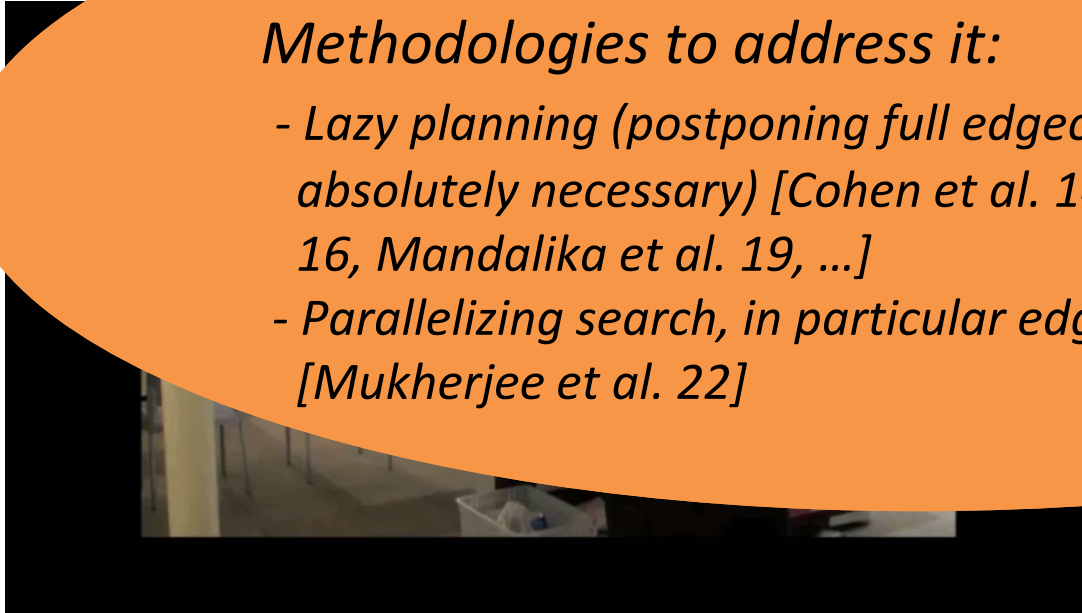
*3-D  $(x,y,\theta)$  planning with **full-body** collision checking*



*Work done in collaboration with Willow Garage  
[Hornung et al. 12]*

## *Methodologies to address it:*


- *Lazy planning (postponing full edgecost evaluation until absolutely necessary) [Cohen et al. 14, Dellin & Srinivasa 16, Mandalika et al. 19, ...]*
- *Parallelizing search, in particular edge evaluations [Mukherjee et al. 22]*




- Generate a systematic graph representation of the planning problem

## Challenges:



- high-dimensionality/graph size
- expensive edge cost evaluation
- *edge construction for dynamic systems*
- *reliance on the accuracy of the model*



Lattice-based graph representation for 3D  $(x, y, \theta)$  planning:








*Planning for a highly dynamic driving (collaboration between Thyssenkrupp and RobotWits, now part of Waymo)*



*Hard/impossible to construct transitions whose endpoints land at the centers of high-d cells*



## *Methodologies to address it:*

- *Soft duplicate detection for search without state discretization but de-prioritizing states that are similar to previously expanded states (i.e., are soft duplicates) [Du et al. 19, Maray et al. 22]*


*Hard/impossible to construct transitions whose endpoints land at the centers of high-d cells*




- Generate a systematic graph representation of the planning problem

## Challenges:

- *high-dimensionality/graph size*
- *expensive edge cost evaluation*
- *edge construction for dynamic systems*
- *reliance on the accuracy of the model*



Lattice-based graph representation for 3D  $(x, y, \theta)$  planning:



*Planning for tasks requiring heavy interaction with the world*



*[Saleem & Likhachev 20]*



### *Methodologies to address it:*

- *Use of a physics-based simulator to compute/evaluate transitions while minimizing the number of calls to the simulator [Saleem & Likhachev 20, Saxena et al. 21, ...]*
- *Updating a model during execution [Sutton 91]. Requires lots of samples and an updateable model.*

*[Saleem & Likhachev 20]*

- Generate a systematic graph for a given planning problem

## Challenges:


- *high-dimensionality/graph size*
- *expensive edge cost evaluation*
- *edge construction for dynamic systems*
- *reliance on the accuracy of the model*

discretize:  Construct  $S_1, S_2, S_3, \dots$  search the graph for solution:

## Lots of work on learning to address these challenges:

- *learning heuristics [Bhardwaj et al. 17, ...]*
- *learning collision detection estimator [Das & Yip 19, Huh & Lee 16, ...]*
- *learning soft duplicate measure [Maray et al. 22]*
- *learning residual models [Nagabandi et al. 19, ...]*
- ....

*motion primitives*





- Generate a systematic graph for a given planning problem

## Challenges:

- high-dimensionality/graph size
- expensive edge cost evaluation
- edge construction for dynamic systems
- reliance on the accuracy of the model



CTMP algorithms  
[Islam et al. 21]

CMAX  
[Vemula 22]

## Lots of work on learning to address these challenges:

- learning heuristics [Bhardwaj et al. 17, ...]
- learning collision detection estimator [Das & Yip 19, Huh & Lee 16, ...]
- learning soft duplicate measure [Maray et al. 22]
- learning residual models [Nagabandi et al. 19, ...]
- ....

motion primitives





- Challenges in Search-based Planning
- Constant-time Motion Planning (CTMP) – offline learning for online planning
- CMAX/CMAX++ for handling inaccurate models
- Summary and thoughts on research directions





[Cowley et al., 2013], joint work with CJ Taylor



<https://youtu.be/mTFBuSuYuZI>


Autonomous truck unloading  
(joint work with Honeywell & NREC - Herman, Pires, etc.)

- Planning often needs to be fast and “constant-time”
- while tasks are often repetitive




Can we provide a **provably** bounded planning time by pre-processing?

[Cowley et al., 2013], joint work with CJ Taylor



<https://youtu.be/mTFBuSuYuZI>



Autonomous truck unloading  
(joint work with Honeywell & NREC - Herman, Pires, etc.)


- Planning often needs to be fast and “constant-time”
- while tasks are often repetitive

## ***Constant-time Motion Planning (CTMP) class of algorithms***

*[Islam et al., ICAPS'19], [Islam et al., RSS'20], [Islam et al., ICRA'21]*

***Algorithms that learn offline data structures which enable online planners to guarantee to find a solution (if one exists) within a (small) user-defined time***

- Given (known) environment
- Given set of potential start configurations
- Given set of corresponding goal regions (e.g., we need to handle goal perturbations)






- Given (a mostly known) environment
- Given set of potential start configurations
- Given set of corresponding goal regions (e.g., we need to handle goal perturbations)




Goal Region




Goal Region




Goal Region




- Given a start state and a goal region

 $R_{goal}$  $s_{start}$ 

- Given a start state and a goal region
- Discretize the goal region

 $R_{goal}$  $s_{start}$

- Given a start state and a goal region
- Discretize the goal region
- ~~We could pre compute all paths but too long, too much memory... and not interesting 😊.~~






- Given a start state and a goal region
- Discretize the goal region

**Key idea:**


***Given a potential function and any “attractor” state  $S^{att}$ , there is typically a large region of states that can reach the attractor state following the potential function***



- Given a start state and a goal region
- Discretize the goal region

**Key idea:**


***Given a potential function and any “attractor” state  $S^{att}$ , there is typically a large region of states that can reach the attractor state following the potential function***



- Given a start state and a goal region
- Discretize the goal region

**Key idea:**

*Given a potential function and any “attractor” state  $S^{att}$ , there is typically a large region of states that can reach the attractor state following the potential function*




**Pre-processing algorithm:**

*Decompose  $R_{goal}$  into a set of subregions  $R_{1...K}$ , each defined by  $\{S_i^{att}, r_i\}$ ,  
s. t.  $\bigcup_i^K R_i$  completely covers  $R_{goal}$   
where  $r_i$  – radius of subregion  $R_i$*

- Given a start state and a goal region
- Discretize the goal region

**Key idea:**

*Given a potential function and any “attractor” state  $S^{att}$ , there is typically a large region of states that can reach the attractor state following the potential function*



**Pre-processing algorithm:**


*Decompose  $R_{goal}$  into a set of subregions  $R_{1...K}$ , each defined by  $\{S_i^{att}, r_i\}$ ,  
s. t.  $\bigcup_i^K R_i$  completely covers  $R_{goal}$   
where  $r_i$  – radius of subregion  $R_i$*

- Given a start state and a goal region
- Discretize the goal region

**Online Planning:**

Given an  $s_{goal}$  in  $R_{goal}$ :

- a) find which  $R_i$  contains  $s_{goal}$
- b) follow the potential function towards  $S_i^{att}$ ,
- c) once reached, follow the stored path from  $S_i^{att}$  to  $s_{start}$



**Pre-processing algorithm:**


Decompose  $R_{goal}$  into a set of subregions  $R_{1...K}$ , each defined by  $\{S_i^{att}, r_i\}$ ,  
s. t.  $\bigcup_i^K R_i$  completely covers  $R_{goal}$   
where  $r_i$  – radius of subregion  $R_i$

- Given a start state and a goal region
- Discretize the goal region

## Online Planning:

### Details on pre-processing:

- each subregion  $R_i$  can be computed via single backward search from  $S_i^{att}$
- $S_{i+1}^{att}$  can be any state that lies in the portion of  $R_{goal}$  that isn't yet covered
  - more details including analysis in [Islam et al., ICAPS'19]




Provides provable guarantee on finding a plan to the goal (if one exists) within a given user-defined time

where  $r_i$  – radius of subregion  $R_i$




- On bin picking 7DOF arm planning

|                    | PRM (4T)    | MQ-RRT (4T) | E-graph        | RRT-Connect | Our method       |
|--------------------|-------------|-------------|----------------|-------------|------------------|
| Planning time [ms] | 21.7 (59.6) | 21.2 (35.5) | 497.8 (9678.5) | 1960 (9652) | <b>1.0 (1.6)</b> |
| Success rate [%]   | 86          | 69.75       | 76.5           | 83.8        | <b>100</b>       |
| Memory usage [Mb]  | 1,828       | 225.75      | <b>2.0</b>     | -           | 7.8              |



(a)



(b)

## **Constant-time Motion Planning (CTMP) class of algorithms**

*[Islam et al., ICAPS'19], [Islam et al., RSS'20], [Islam et al., ICRA'21]*

**Algorithms that learn offline data structures which enable online planners to guarantee to find a solution (if one exists) within a (small) user-defined time**


### **Examples of observations exploited by these algorithms:**

- *Goal region can be decomposed into subsets within which one can get to its center by following a potential function [Islam et al., ICAPS'19]*
- *Paths can be precomputed so as to guide Experience-based planner to find a solution within  $X$  expansions [Islam et al., RSS'20]*
- *Disjoint paths guarantee that at least one is feasible given a potential for an obstacle blocking one [Islam et al., ICRA'21]*



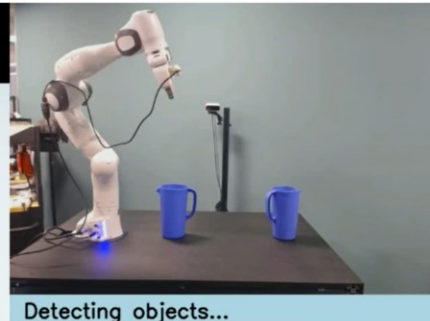
CTMP for picking dynamic objects off a conveyor with imperfect perception [Islam et al., RSS'20]

Simulation Experiments



100% task success rate in simulation

Real World Experiments



Detecting objects...

• Planning happens between "Detecting objects" and "Following trajectory" (notice the blink)

**Task:** The robot must pick up the bowl while avoiding collisions with the pitchers and the table

CTMP for picking up objects in partially-known environments [Islam, Paxton, Eppner, Peele, Likhachev, Fox et al., ICRA'21] (collaboration with Nvidia)




CTMP for Shield-based Protection project (work in progress)

- Challenges in Search-based Planning
- Constant-time Motion Planning (CTMP) – offline learning for online planning
- CMAX/CMAX++ for handling inaccurate models
- Summary and thoughts on research directions




- Planning models for real world tasks are often complex (e.g., physics-based simulators, analytically computed motion primitives, etc.), yet often imperfect
- Learning a model on-the-fly requires too many samples for goal-oriented execution



How to interleave planning and execution to  
**guarantee task achievement**  
despite the inaccuracies in the model?

Related to  
“cost poisoning” [Zucker et al. 2011]

- Main points behind CMAX [Vemula, Oza, Bagnell & Likhachev, RSS'20]
  - instead of updating dynamics, inflates the cost of transitions discovered to be incorrect
  - does not require updates to the dynamics of the model
  - uses limited expansion search as planner to bound computation
  - uses function approximation to scale to large state spaces
- Guarantees task achievement under certain conditions



- Objective:

Provably reach the goal online, despite having an inaccurate dynamical model, ***without any resets***

<sup>1</sup>*Resets allow the robot to “reset” to a state, usually a previously visited state*

- The problem is formulated as a shortest path problem  $M=(\mathcal{S}, \mathcal{A}, \mathcal{G}, f, c)$   
 $\mathcal{S}$  : State space,  $\mathcal{A}$ : Discrete action space,  $\mathcal{G}$ : Goal space  
 Cost function:  $c: \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$   
 Unknown Deterministic True Dynamics:  $f: \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$   
 Access to Approximate Dynamics:  $f': \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$   
 State is fully observable



- Incorrect transitions:

Transitions where true and approximate dynamics differ  
for example,  $f(s,a) \neq f'(s,a)$  or  $\|f(s,a) - f'(s,a)\| > \xi$

$\xi$  - smallest discrepancy handled by low-level feedback controller

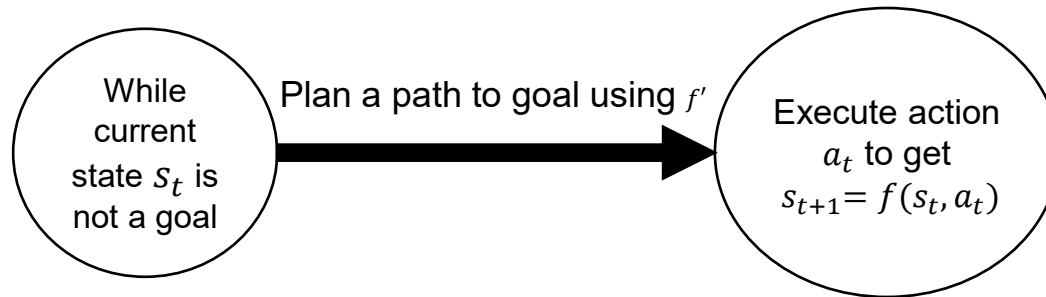
$\mathcal{X} \subseteq \mathcal{S} \times \mathcal{A}$  = set of “incorrect” transitions

$\mathcal{X}$  is not known beforehand, and is only discovered through online executions

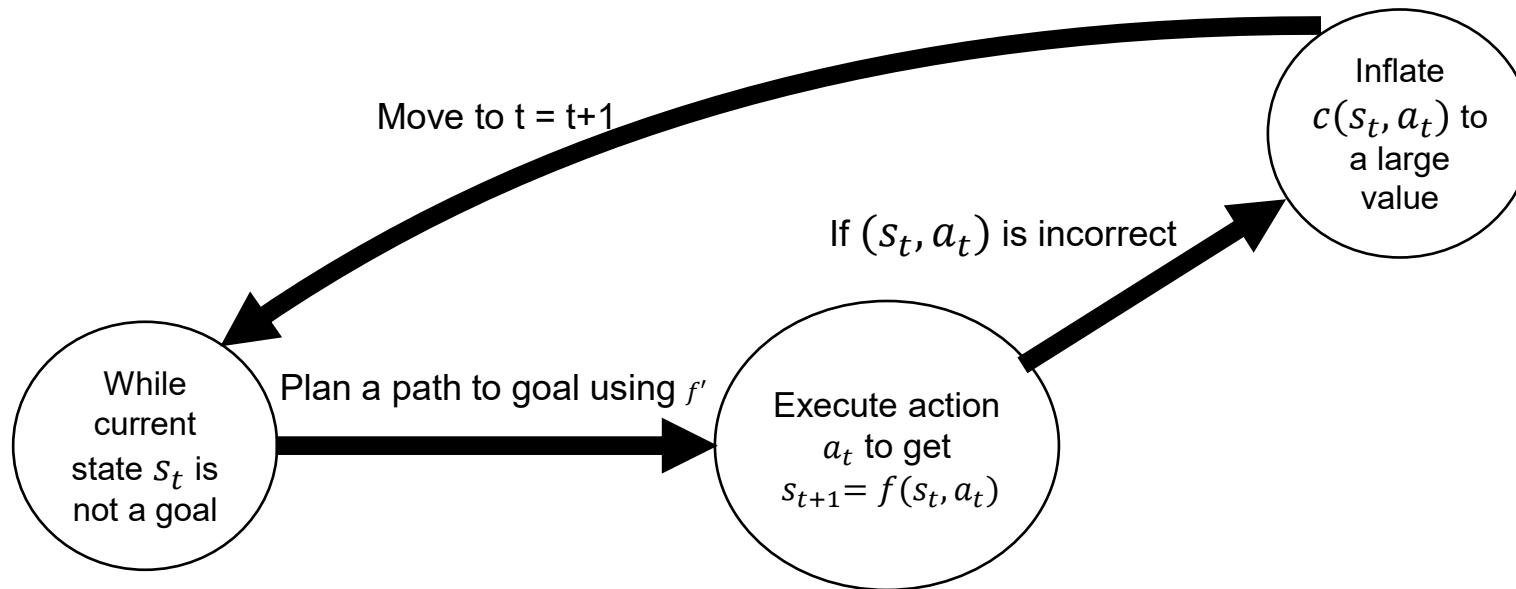
## Key Idea:

Instead of learning the true dynamics, CMAX maintains  
a running estimate of the set of incorrect transitions  $\mathcal{X}$  and  
biases the planner to avoid using transitions known to be incorrect

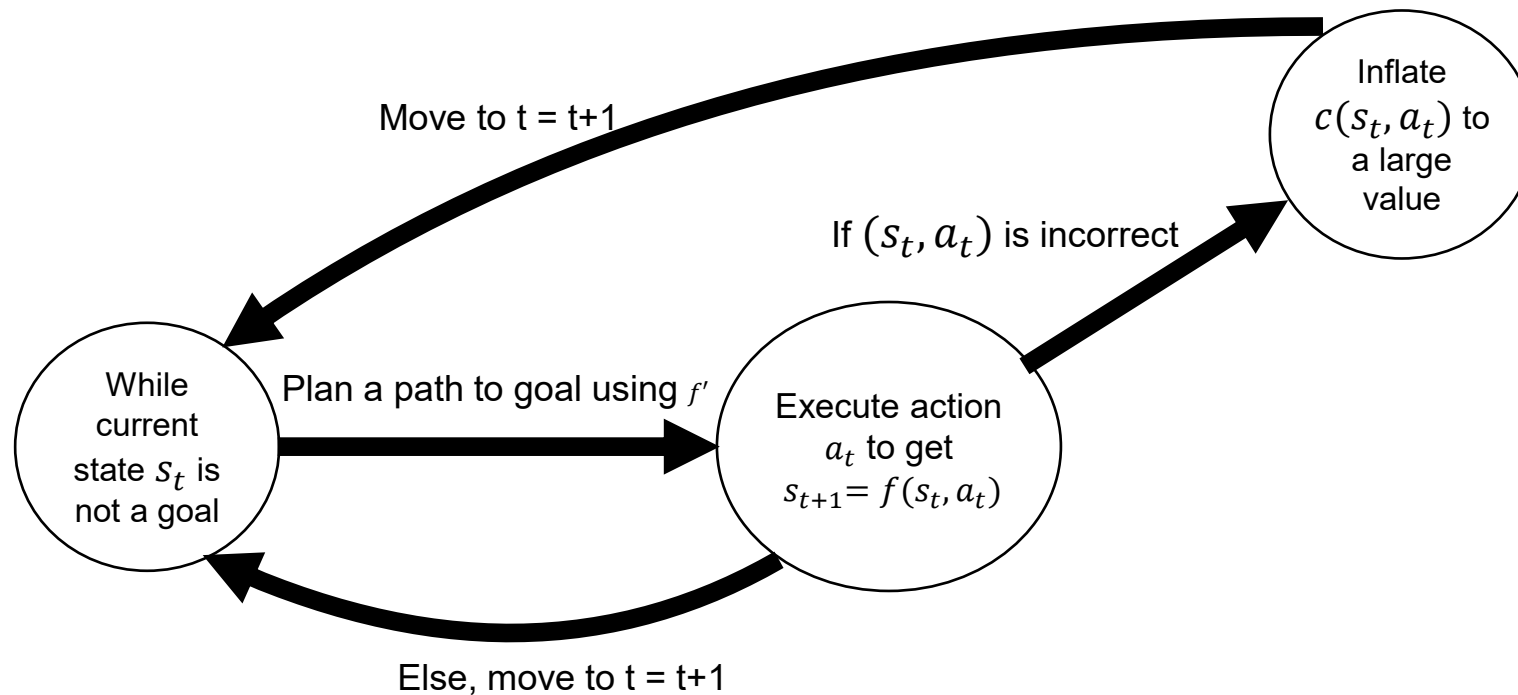
- Key idea:



- Key idea:



- Key idea:



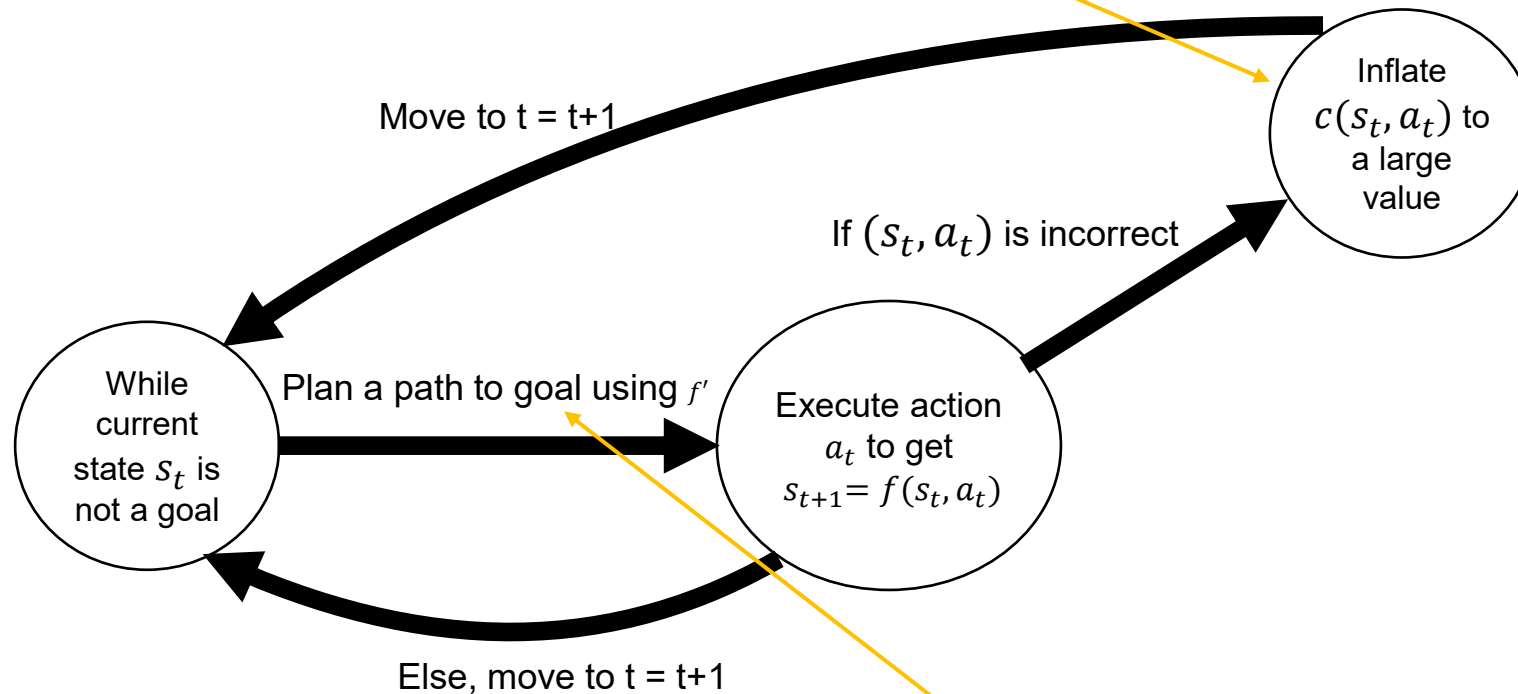
**Does not update approximate dynamics  $f'$  !**

**Theorem.** *The robot is guaranteed to reach a goal (accomplish its task), i.e. CMAX is task-complete, under the assumption that there always exists a path from  $s_t$  to a goal that does not contain any transitions  $(s, a)$  known to be incorrect, i.e.  $(s, a) \in \mathcal{X}_t$*

- Key idea:

**To handle large state-spaces:**

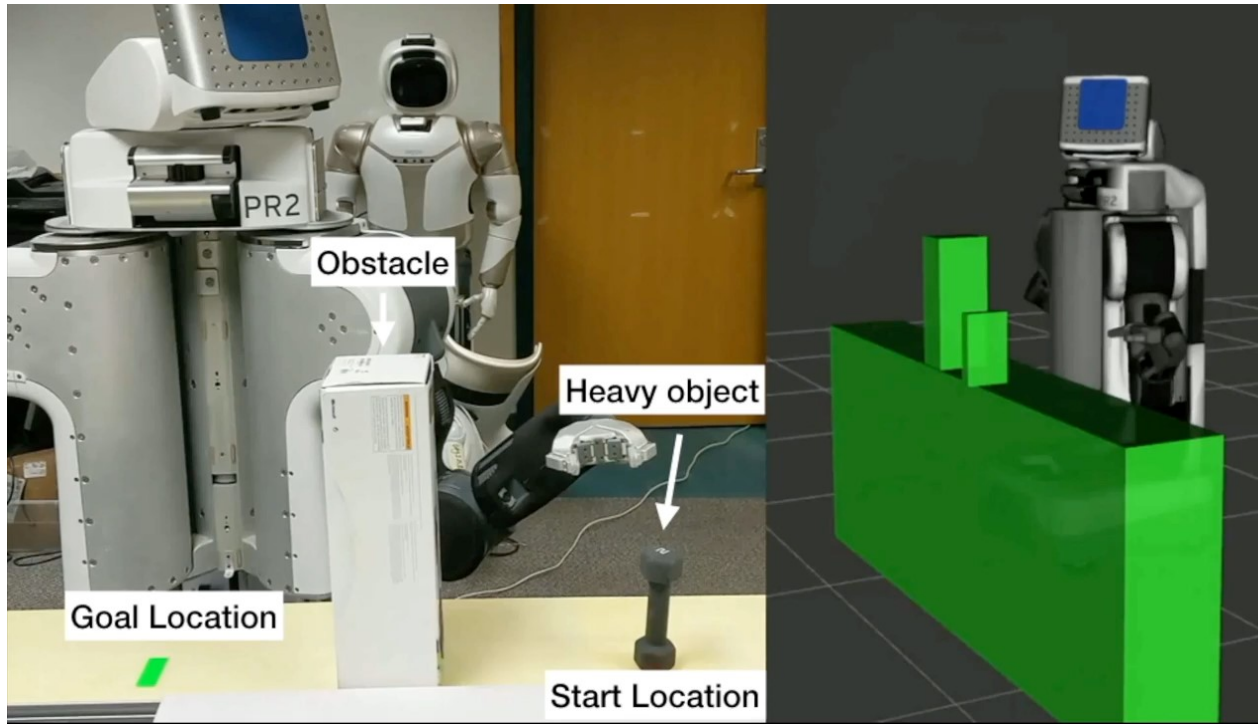
uses function approximators to learn where discrepancies are




**To handle large state-spaces:**

extended to support limited-horizon search (e.g., LRTA\* [Korf 90])


- CMAX in action:




- Task achievement vs. optimality for repeated tasks



*CMAX*  
after 1<sup>st</sup> iteration



*CMAX*  
after n<sup>th</sup> iteration



*potentially highly sub-optimal path*




- Task achievement vs. optimality for repeated tasks

## **CMAX++ [Vemula, Bagnell & Likhachev, AAAI'21]**

- *combines CMAX with model-free Q-learning*
- *learns optimal Q-values of incorrect transitions over time and slowly switches to using those during planning*
- ***Guarantees task achievement under certain conditions AND convergence to optimal execution of repeated tasks***

- Task achievement vs. optimality for repeated tasks



## **CMAX++ [Vemula, Bagnell & Likhachev, AAAI'21]**

- combines CMAX with model-free Q-learning
- learns optimal Q-values of incorrect transitions over time and slowly switches to using those during planning
- **Guarantees task achievement under certain conditions AND convergence to optimal execution of repeated tasks**

- Challenges in Search-based Planning
- Constant-time Motion Planning (CTMP) – offline learning for online planning
- CMAX/CMAX++ for handling inaccurate models
- Summary and thoughts on research directions

*Challenges in Search-based Planning:*

- *high-dimensionality/graph size*
- *expensive edge cost evaluation*
- *edge construction for dynamic systems*
- *reliance on the accuracy of the model*

*CTMP algorithms*  
*[Islam et al. 21]*

*CMAX*  
*[Vemula 22]*

## *Challenges in Search-based Planning:*


*Challenges in learning to reduce reliance on the accuracy of the model:*

- 1) How to do it while providing goal-directed behavior and task achievement guarantees?*
- 2) How to combine learning from demos with model-based planning while guaranteeing safety?*
- 3) How to get good confidence estimates in learned policies and how to incorporate those into planning?*

*Two main goals:*

- 1) Learning with the aim of speeding up planning*
- 2) Learning with the aim of reducing dependency on the accuracy of the model*

- All the students, postdocs and collaborators!



- Funding sources for this research:
  - ONR
  - ARL
  - NSF
  - DARPA
  - Mitsubishi
  - Honeywell