

Learning to Plan with Logical Automata


Brandon Araki^{1*}, Kiran Vodrahalli^{2*}, Thomas Leech^{1,3},
Mark Donahue³, Cristian-Ioan Vasile¹, Daniela Rus¹


¹Massachusetts Institute of Technology

²Columbia University

³MIT Lincoln Laboratory

*Equal contributors





Goals

Learn to plan in an environment with rules

1. Learn the rules in a way that they can be easily **interpreted** by humans
2. Incorporate the rules into planning so that **modifying** the rules results in predictable changes in behavior

Packing a Lunchbox

Pack a burger or a sandwich; then pack a banana



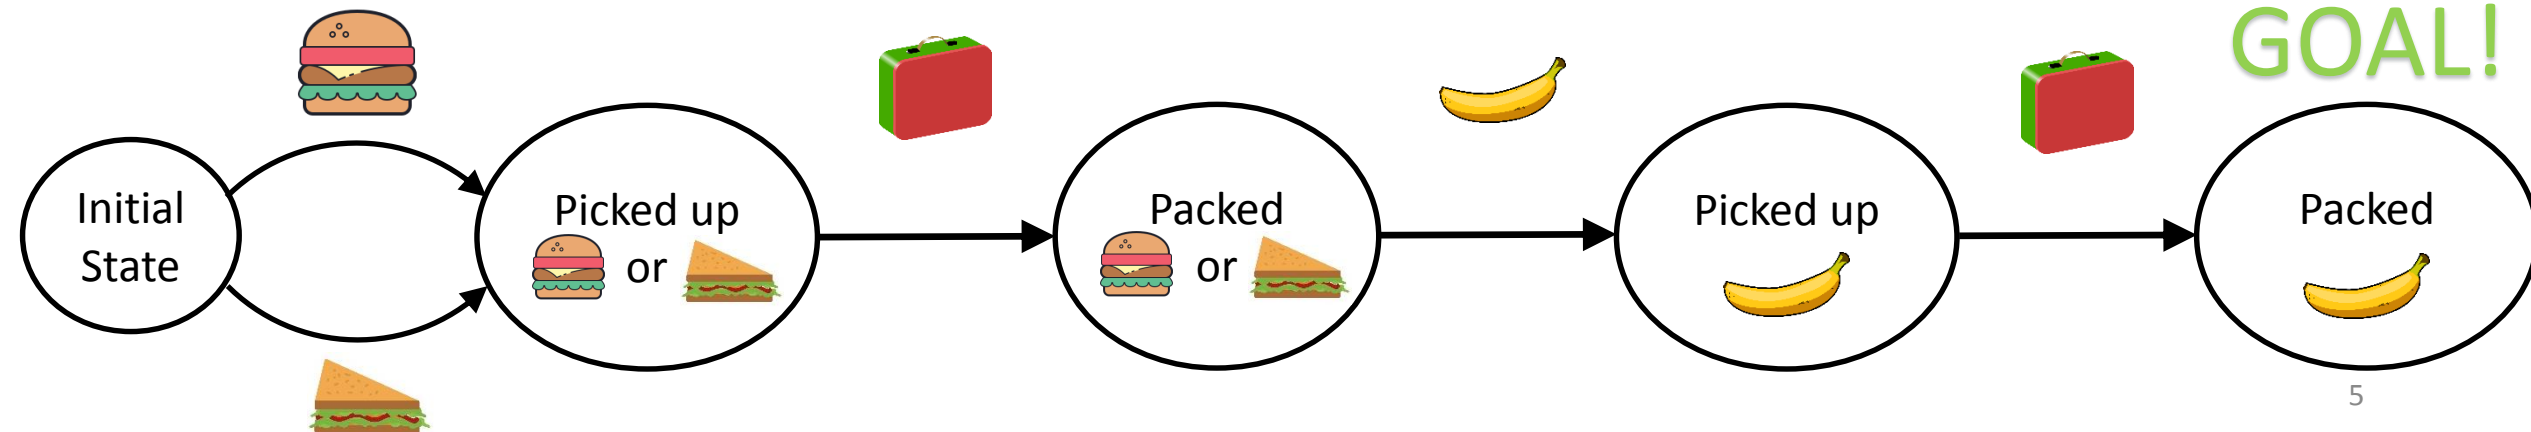
Goal 1 – Interpretability

Rules

Pack a burger or a sandwich;
then pack a banana



Finite State Automaton




Factoring the Environment

High-level MDP


Pack sandwich or burger;
Then pack banana
Avoid obstacles

Low-level MDP




Representing the Environment


Pack sandwich or burger;
Then pack banana
Avoid obstacles



Finite state automaton




Discrete 2D gridworld



Goal 2 – Manipulability


Incorporate FSA into planning







S0

					o	∅
S0						
S1						
S2						
S3						
G						
T						

Differentiable Recursive Planning







Experiments - Interpretability

		Propositions					
S0						o	∅
FSA States	S0						
	S1						
	S2						
	S3						
	G						
	T						

Experiments - Interpretability

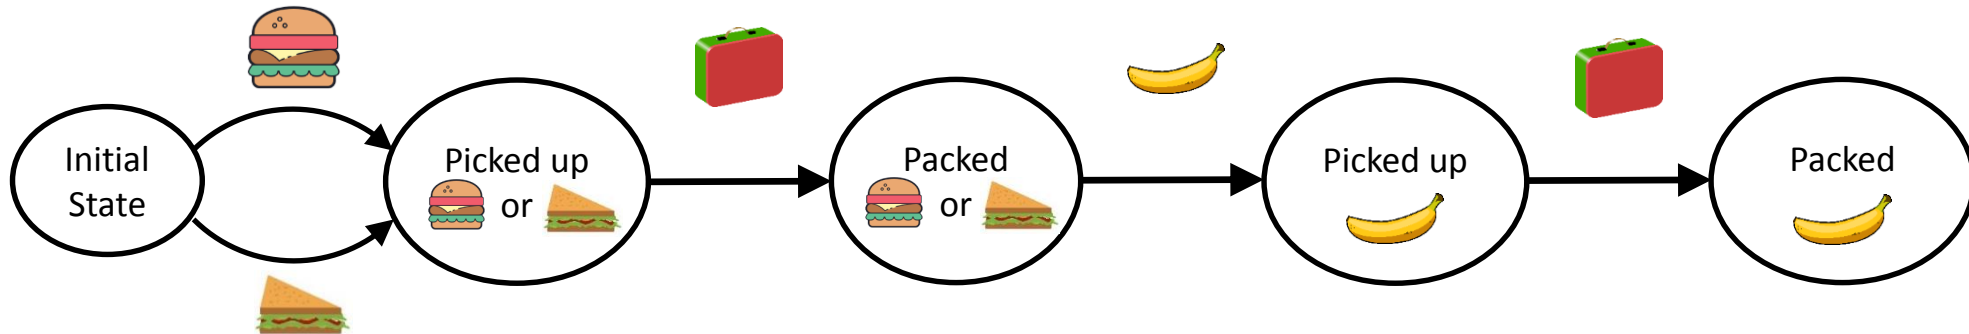
Picking up the **sandwich** or the **hamburger** causes a transition to the next state

S0

					o	∅
S0						
S1						
S2						
S3						
G						
T						


Experiments – Manipulability

We can modify the FSA so that it will only pick up the burger and not the sandwich.



Experiments – Manipulability





We can modify the FSA so that it will only pick up the burger and not the sandwich.



Experiments – Manipulability

We can modify the FSA so that it will only pick up the burger and not the sandwich.





S0

					o	∅
S0						
S1						
S2						
S3						
G						
T						

Experiments – Manipulability

We can modify the FSA so that it will only pick up the burger and not the sandwich.





S0

					o	∅
S0						
S1						
S2						
S3						
G						
T						

Experiments – Manipulability

We can modify the FSA so that it will only pick up the burger and not the sandwich.

S0

					o	∅
S0						
S1						
S2						
S3						
G						
T						

Learning to Plan with Logical Automata

Brandon Araki^{1*}, Kiran Vodrahalli^{2*}, Thomas Leech^{1,3},
Mark Donahue³, Cristian-Ioan Vasile¹, Daniela Rus¹

¹Massachusetts Institute of Technology

²Columbia University

³MIT Lincoln Laboratory

*Equal contributors

