

~~Large-scale planning for resource management problems~~

Why I love POMDPs and so should you

Anthony Corso


acorso@stanford.edu

October 21, 2022



Models for decision making


Single-shot




(States, Actions, Rewards)




Markov Decision Process (MDP)



(States, Actions, Rewards, [Transition](#))



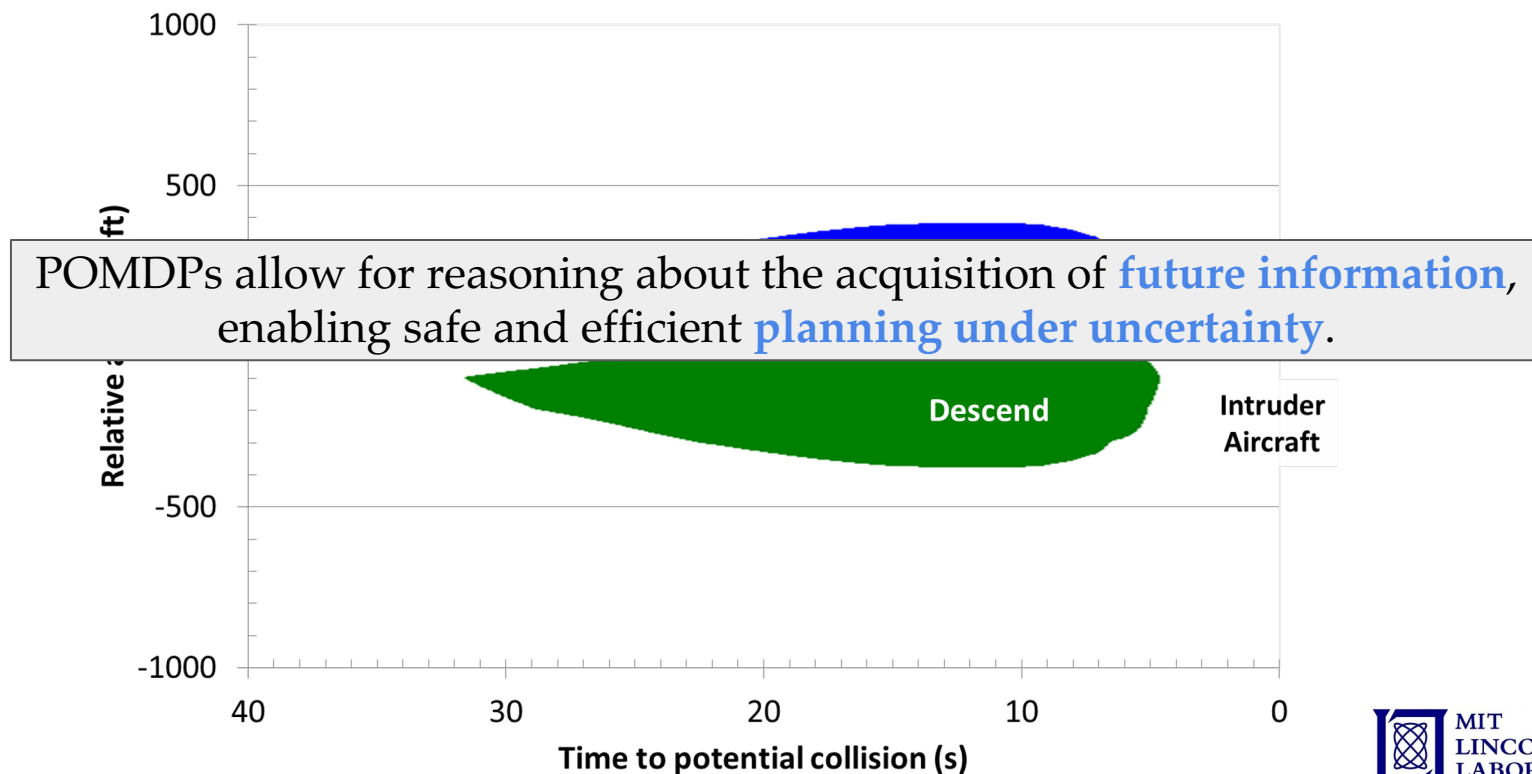
Partially Observable (POMDP)



(States, Actions, Rewards, Transition, [Observations](#))




Aircraft Collision Avoidance




Geological Carbon Storage


State Uncertainty



Transitions




Actions



Reward


Trapped Volume
Exited Volume
Max Pressure

Observations



Modern Solution Techniques

Deep Surrogate Modeling



Forward Tree Search




Image credit: DeepMind, Alpha Go, 2016.

Method	Trapped volume (Gt)	Exited volume (Gt)	Total reward
Random policy	844.855	5022.96	-1.17864e6
Expert solution	1915.2	8.83861	1331.67
POMDP solver	2805.4	0.0	5403.65

Lets collaborate!

