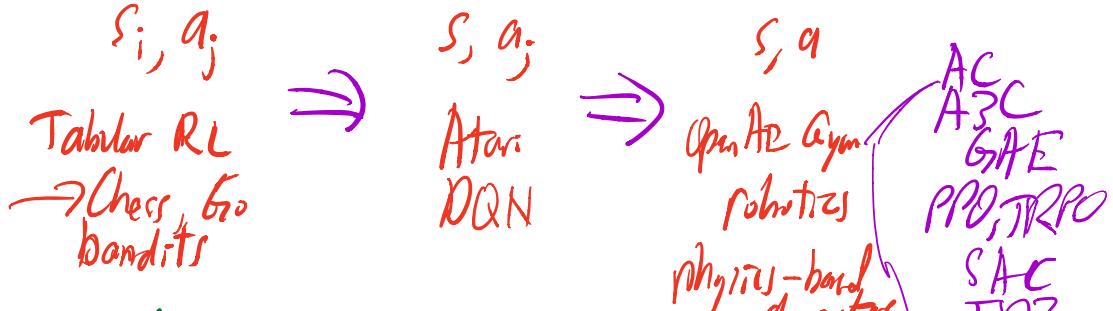


The RL Landscape

~~X~~ MDP: states, actions, rewards, dynamics, initial states



- single agent
- single task
- known rewards, initial state
- stationary
- Δt
- tabula rasa

Efficient / Assumptions / Generalization
Prior knowledge / Rep. Learning

Hyperparameters matter
Adversarial env
Visualize
RL is fragile

RL Systems!
Sutton's "bitter lesson"
Scalability
"cake"

Mult. Agent RL
AAMAS / AAAI
Coop / competitive

Hierarchical RL
Options
Modular RL
multiple time scales

