

# Part I: Tabular Solution Methods

- look at all core RL methods in simple ways
- + state / action spaces are small enough to to approximate value functions in arrays / tables
- first chapters: single state RL problems
- + known as bandit problem
- then formalize with Markov definition (finite)
- solving methods next: dynamic programming, Monte Carlo methods, temporal diff. learning
- + DP: well developed mathematically, but require complete & accurate model of env.
- + Monte Carlo doesn't need model & are simple, but don't do well for incremental computation
- + temporal-difference: no model, incremental solving, but complex to analyze
- \* diff speed + convergence issues too
- finally, how do you combine each to get the best of all worlds?