# Inverse Reinforcement Learning

Andrew Ng and Stuart Russell

ICML 2000

Reinforcement Learning: Given reward structure and state mechanics, derive the optimal policy.

Inverse Reinforcement Learning: Given state mechanics and optimal policy derive reward structure.

When is IRL useful?

Details:

- Ambiguity in reward functions

- Solution method: linear programming

- Settings with tested algorithms

  - Simple discrete state space with known mechanics and known policy

  - Simple continuous state space with known mechanics and known policy

  - Simple continuous state space with known mechanics and simulatable policy