Adapting Representation to the Problem

Typically, **representations given and/or chosen manually**

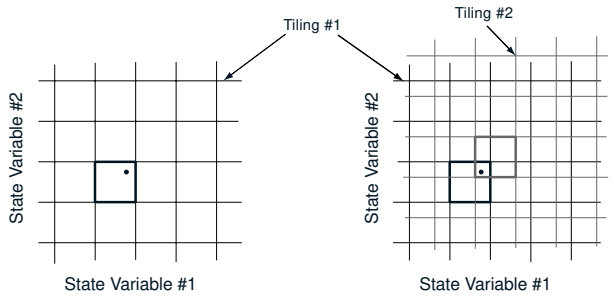

Adapting Representation to the Problem

Typically, **representations given and/or chosen manually**

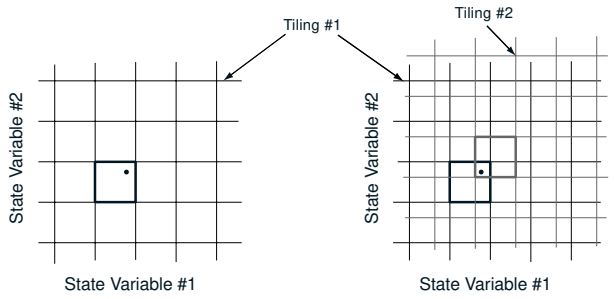

The crucial factor for a successful approximate algorithm is the choice of the parametric approximation architecture. . . ." [Lagoudakis & Parr,'03]

Adapting Representation to the Problem

- How do we *represent* our solution?
- Example: using *neural networks*

- Too simple: suboptimal performance
- Divergence and catastrophic performance
  [Baird 1995] [Boyan & Moore 1995]

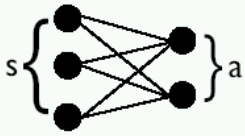

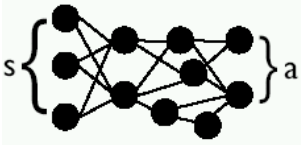

Too complex: infeasibly slow learning

Adapting Representation to the Problem

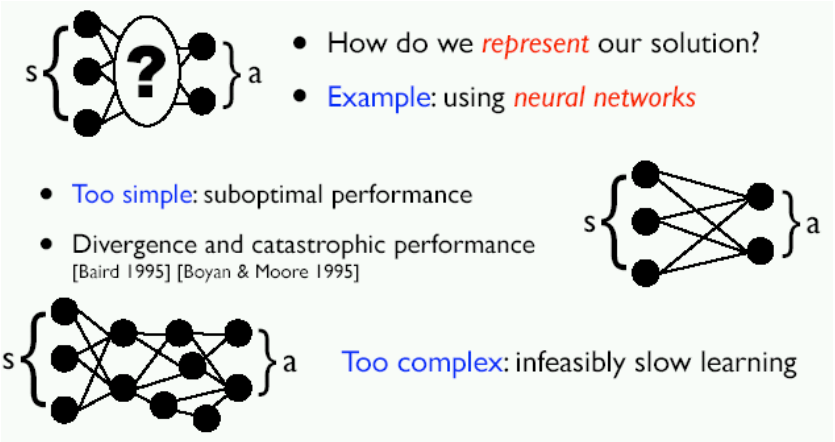

Can RL agents automatically **learn** effective **representations?**

NEAT+Q [Whiteson & Stone, JMLR 2006]

Evolve agents that are better able to **learn**

- Evolution chooses representation and initial weights
  - NEAT learns NN topologies [Stanley & Miikkulainen, '02]
- Q-learning learns weights that approximate value function

NEAT+Q [Whiteson & Stone, JMLR 2006]

Evolve agents that are better able to **learn**

- Evolution chooses representation and initial weights
  - NEAT learns NN topologies [Stanley & Miikkulainen, '02]
- Q-learning learns weights that approximate value function

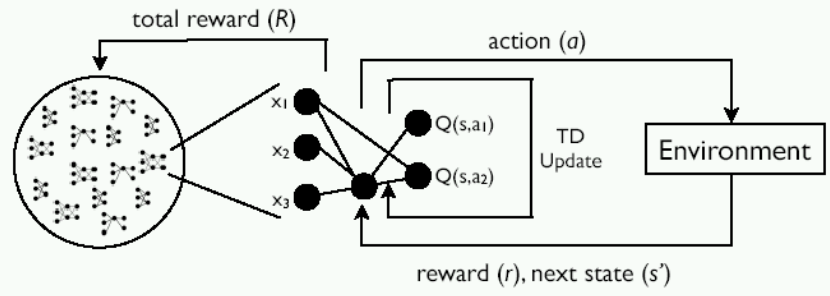

NEAT+Q Results

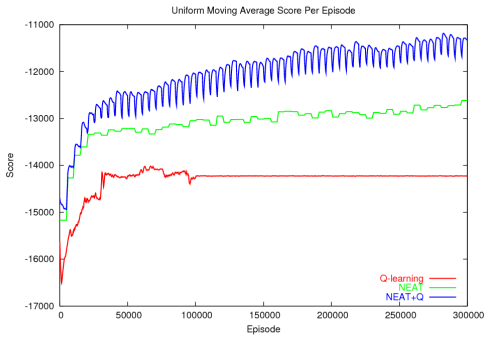

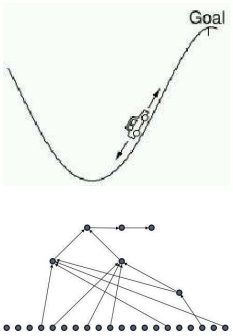

- Neural net function approx. works on mountain car!
- Tested Q-learning with 24 manual configurations