Co-financed by the Connecting Europe Facility of the European Union

Advanced Topics in AI

Local Search

Instructor: Prof. Dr. techn. Wolfgang Nejdl

Leibniz University Hannover

[These slides were created by Dan Klein and Pieter Abbeel for CS188 Intro to AI at UC Berkeley. All materials are available at http://ai.berkeley.edu.]

Local Search

- Tree search keeps unexplored alternatives on the fringe (ensures completeness)

- Local search: improve a single option until you can't make it better (no fringe!)

- New successor function: local changes

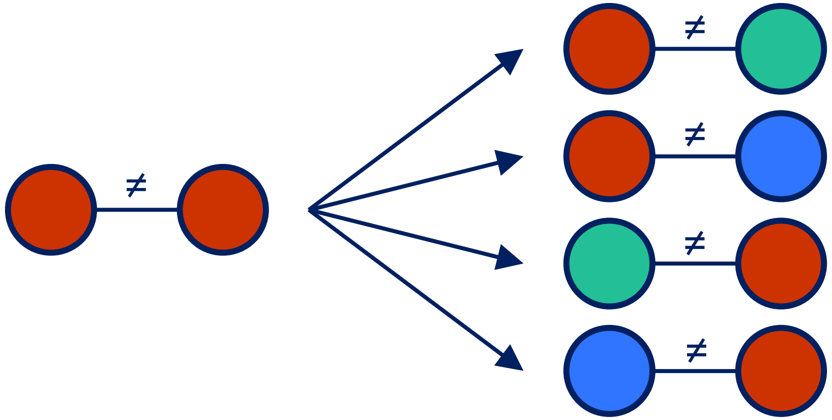

- Generally much faster and more memory efficient (but incomplete and suboptimal)

Co-financed by the Connecting Europe Facility of the European Union

Hill Climbing

- Simple, general idea:
  - Start wherever
  - Repeat: move to the best neighboring state
  - If no neighbors better than current, quit

- What's bad about this approach?

- What's good about it?

Co-financed by the Connecting Europe Facility of the European Union

Hill Climbing Diagram

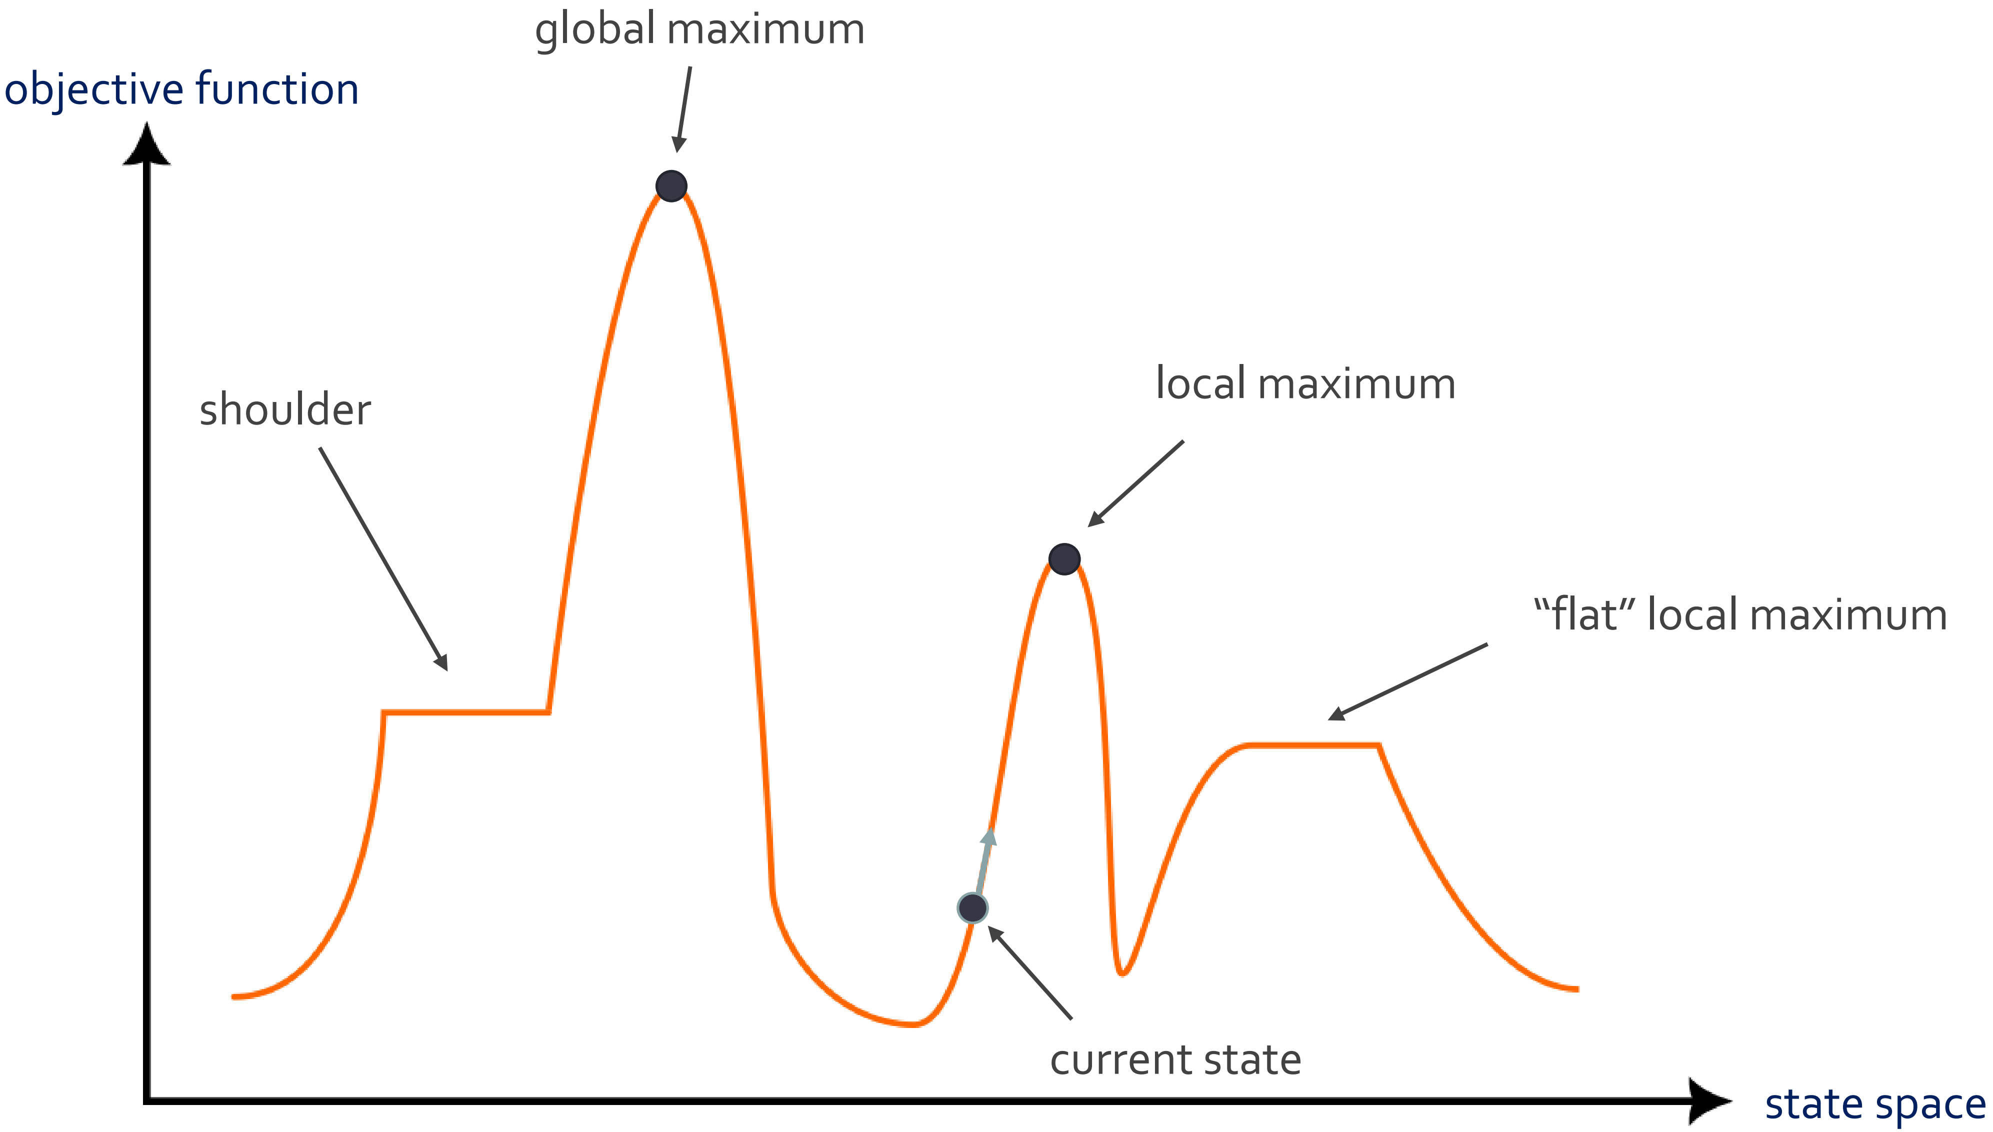

Co-financed by the Connecting Europe Facility of the European Union

Hill Climbing Quiz

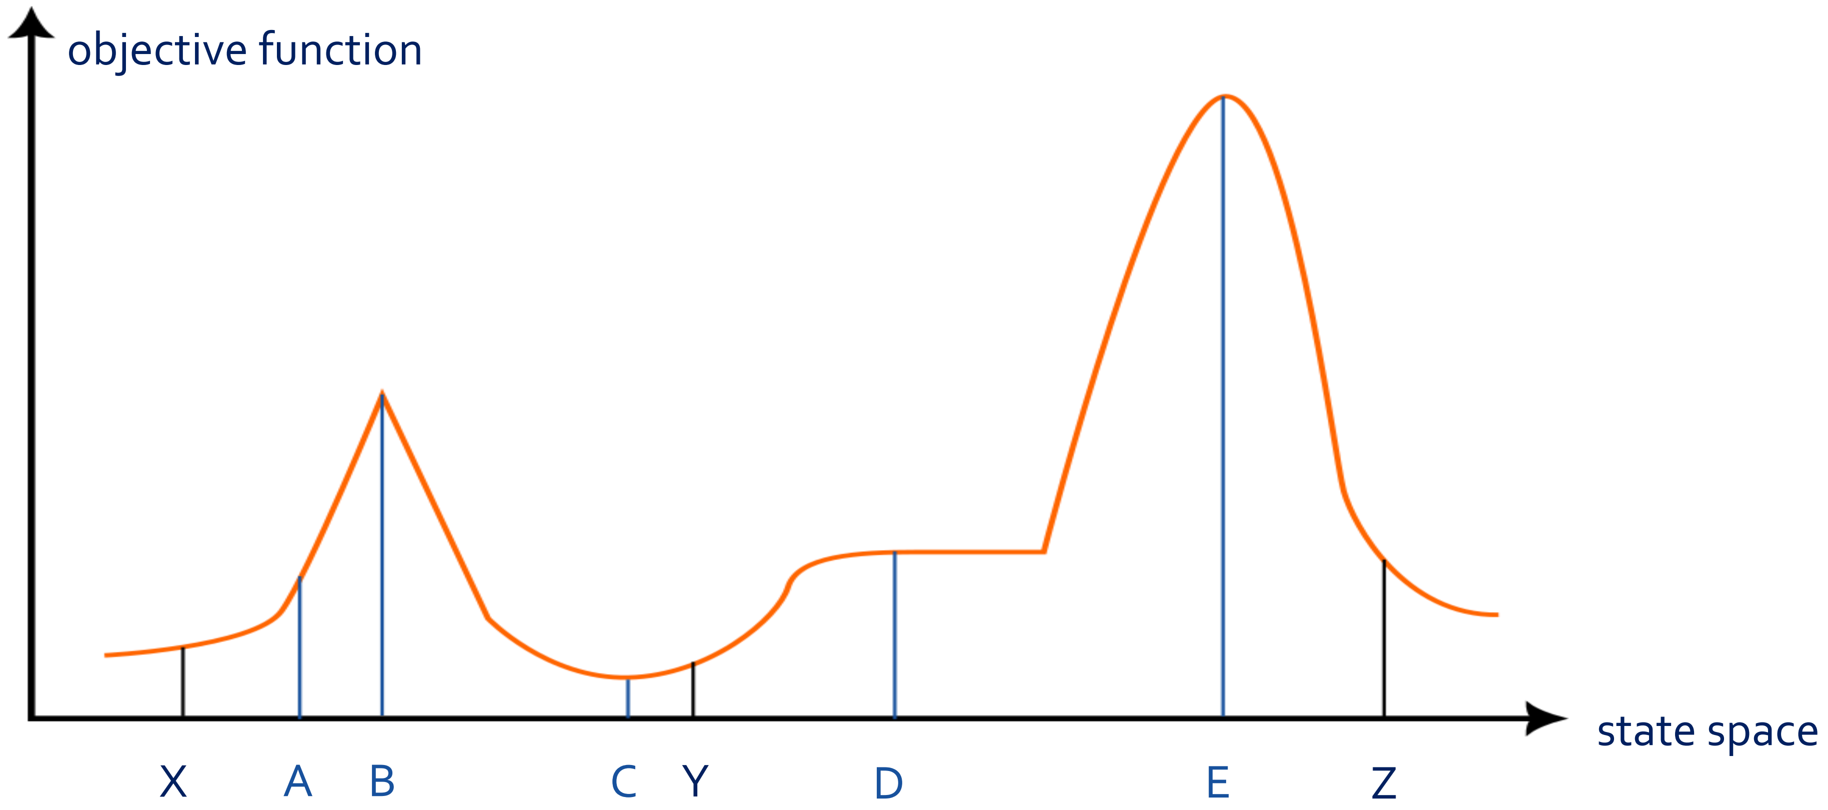

- Starting from X, where do you end up?
- Starting from Y, where do you end up?
- Starting from Z, where do you end up?

Co-financed by the Connecting Europe Facility of the European Union

Simulated Annealing

- Idea: Escape local maxima by allowing downhill moves
  - But make them rarer as time goes on

```
function SIMULATED-ANNEALING(problem, schedule) returns a solution state
    inputs:     problem, a problem
                schedule, a mapping from time to "temperature"
    local variables:    current, a node
                        next, a node
                        T, a "temperature" controlling probability of downward steps

    current ← MAKE-NODE(INITIAL-STATE[problem])
    for t ← 1 to ∞ do
        T ← schedule[t]
        if T = 0 then return current
        next ← a randomly selected successor of current
        ΔE ← VALUE[next] – VALUE[current]
        if ΔE > 0 then current ← next
        else current ← next only with probability e^(ΔE/T)
```

Co-financed by the Connecting Europe Facility of the European Union

Simulated Annealing

- Theoretical guarantee:
  - Stationary distribution: $p(x) \propto e^{\frac{E(x)}{kT}}$
  - If T decreased slowly enough,
    will converge to optimal state!

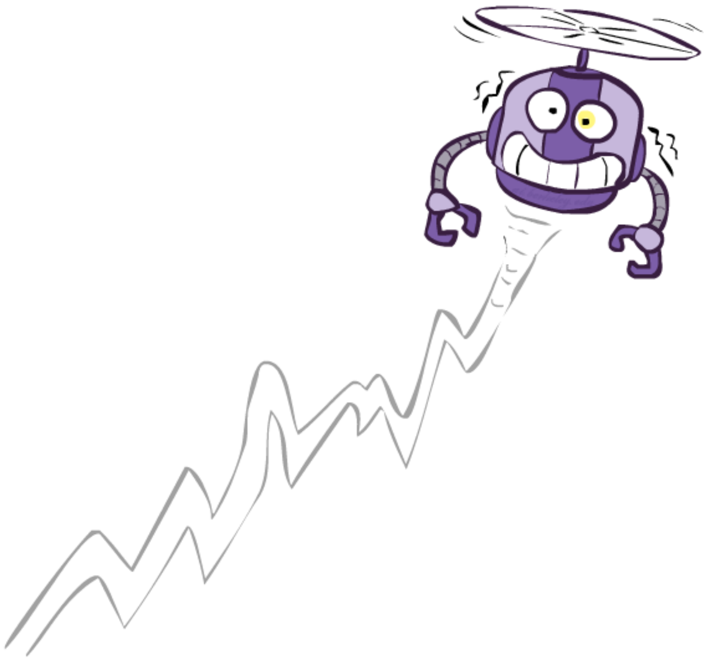

- Is this an interesting guarantee?

- Sounds like magic, but reality is reality:
  - The more downhill steps you need to escape a local optimum, the less likely you are to ever make them all in a row
  - People think hard about *ridge operators* which let you jump around the space in better ways

Co-financed by the Connecting Europe Facility of the European Union

Genetic Algorithms

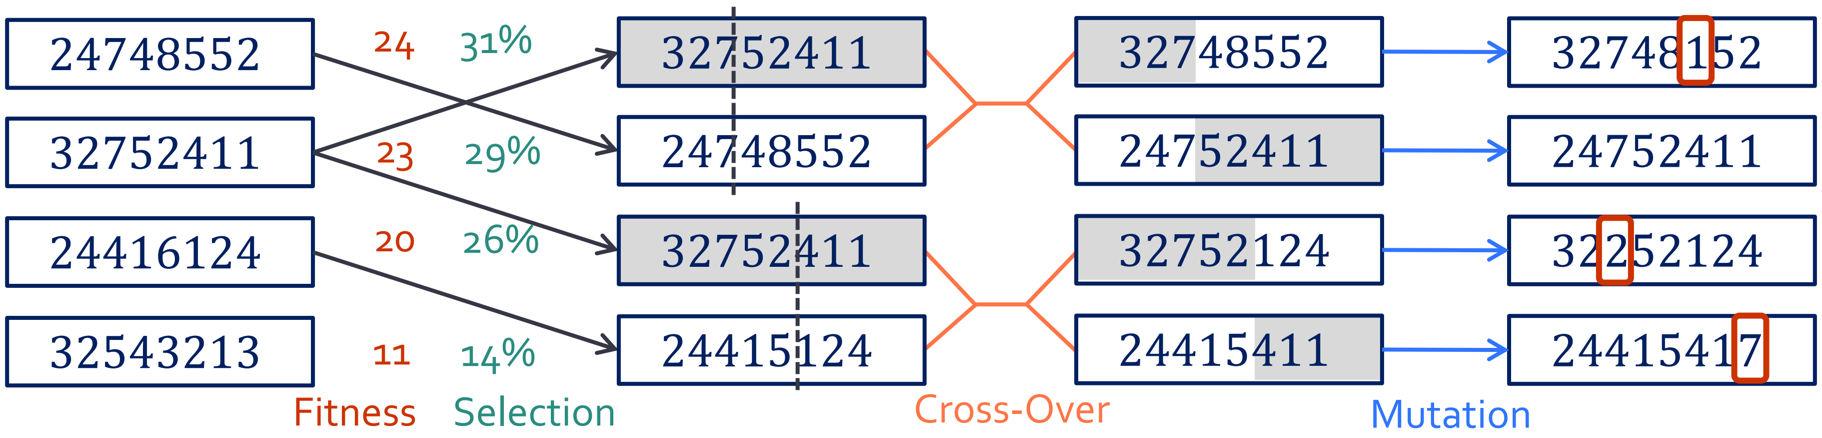

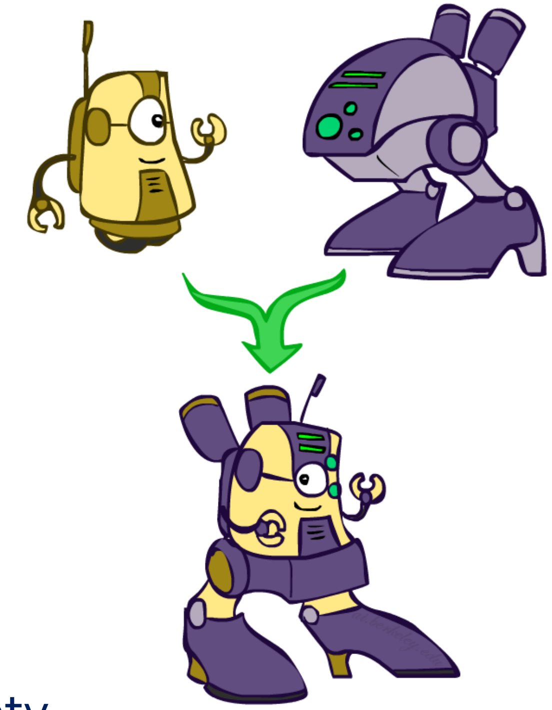

- Genetic algorithms use a natural selection metaphor
  - Keep best N hypotheses at each step (selection) based on a fitness function
  - Also have pairwise crossover operators, with optional mutation to give variety

- Possibly the most misunderstood, misapplied (and even maligned) technique around

Co-financed by the Connecting Europe Facility of the European Union

Example: N-Queens

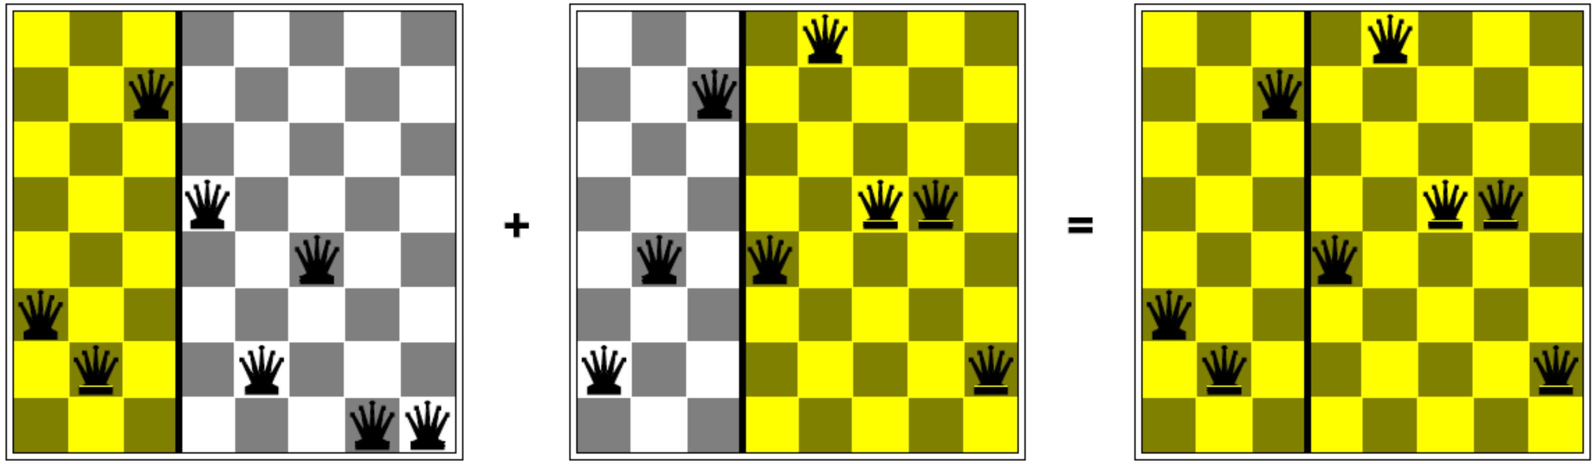

- Why does crossover make sense here?
- When wouldn't it make sense?
- What would mutation be?
- What would a good fitness function be?

Co-financed by the Connecting Europe Facility of the European Union

Advanced Topics in AI

Next topic: Search with Other Agents

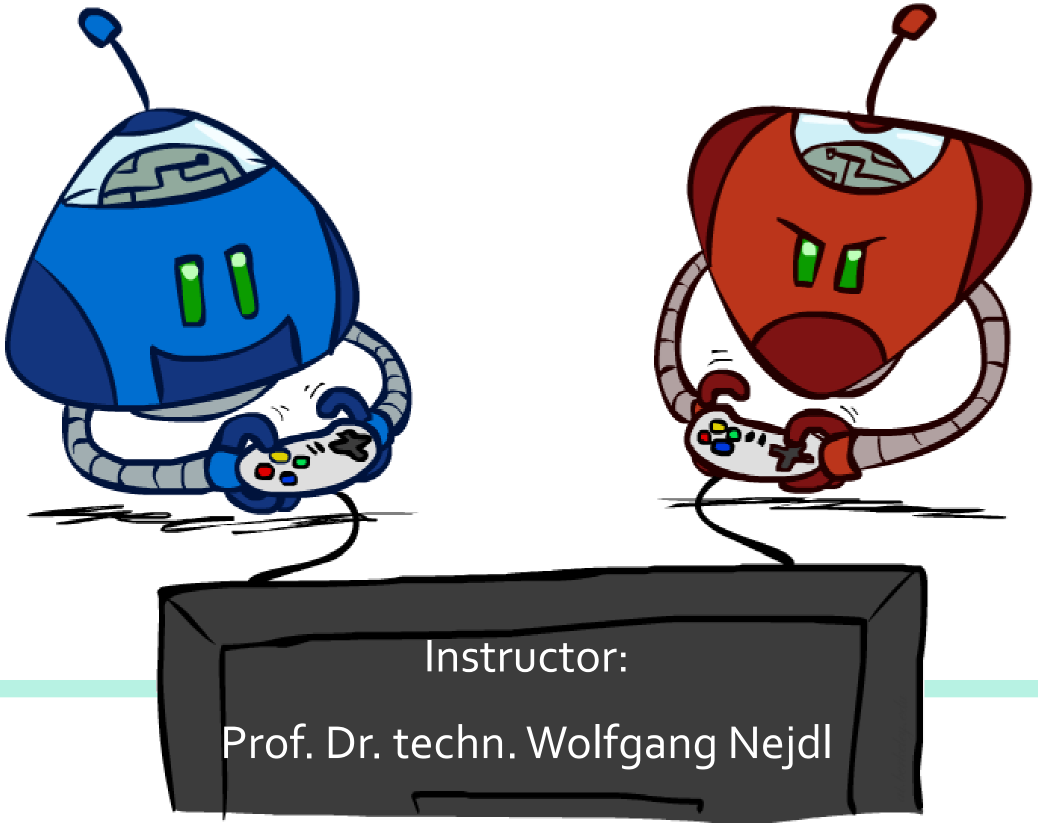

Instructor:

Prof. Dr. techn. Wolfgang Nejdl

Leibniz University Hannover

Co-financed by the Connecting Europe Facility of the European Union

[These slides were created by Dan Klein and Pieter Abbeel for CS188 Intro to AI at UC Berkeley. All materials are available at http://ai.berkeley.edu.]