CSE 473: Artificial Intelligence

Course Wrapup

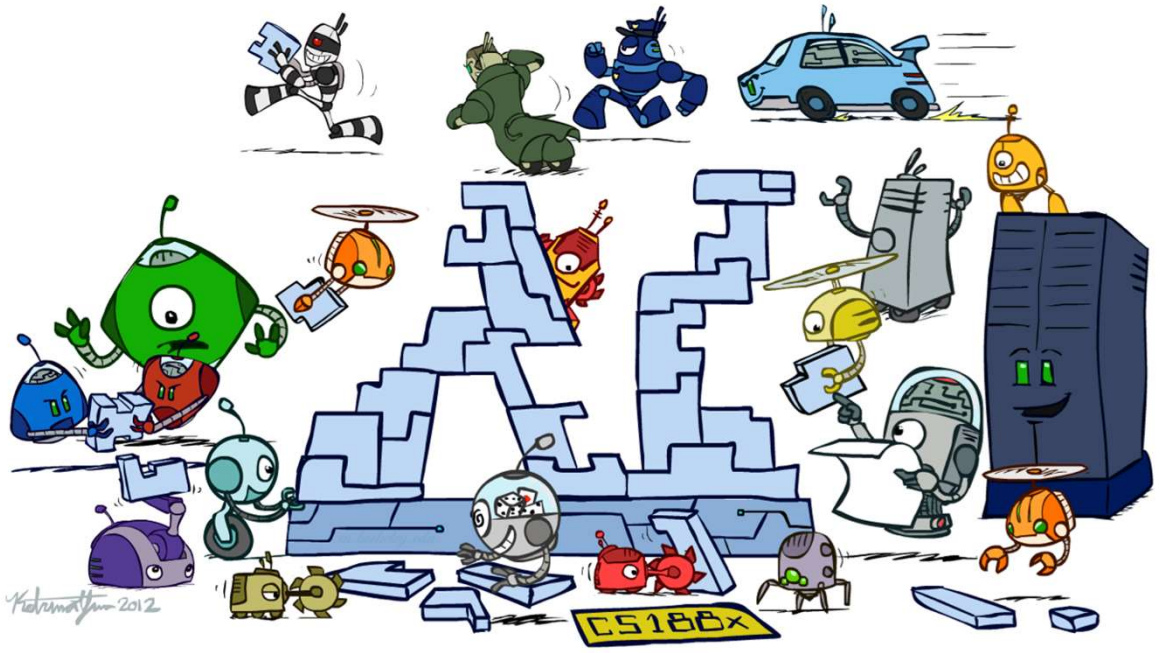

Steve Tanimoto --- University of Washington

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

Exam Topics

- Search
  - Problem spaces
  - BFS, DFS, UCS, A* (tree and graph), local search
  - Completeness and Optimality
  - Heuristics: admissibility and consistency; pattern DBs
- CSPs
  - Constraint graphs, backtracking search
  - Forward checking, AC3 constraint propagation, ordering heuristics
- Games
  - Minimax, Alpha-beta pruning,
  - Expectimax, Evaluation Functions
- MDPs
  - Bellman equations
  - Value iteration, policy iteration
- Reinforcement Learning
  - Exploration vs Exploitation
  - Model-based vs. model-free, Q-learning
- Markov Models
  - Diagrams: Bayes-net, state transition, trellis
  - Stationary probability distribution
- Hidden Markov Models
  - DBNs
  - Forward algorithm
  - Particle Filters
- Bayesian Networks
  - Joint distributions, probabilistic inference, var. elim.
  - Basic definition, independence, D-separation
  - Conditional independence, Bayes' rule
- Learning
  - Perceptrons, training algorithm, linear separability
- Natural Language Processing
  - Document comparison using cosine similarity
- Miscellaneous
  - Types of agents and environments
  - Turing test
  - Asimov's three laws of robotics

What is intelligence?

- (bounded) Rationality
  - Agent has a performance measure to optimize
  - Given its state of knowledge
  - Choose optimal action
  - With limited computational resources
- Human-like intelligence/behavior

Search in Discrete State Spaces

- Every discrete problem can be cast as a search problem.
  - states, actions, transitions, cost, goal-test
- Types
  - uninformed systematic: often slow
    - DFS, BFS, uniform-cost, iterative deepening
  - Heuristic-guided: better
    - Greedy best first, A*
    - relaxation leads to heuristics
  - Local: fast, fewer guarantees; often local optimal
    - Hill climbing and variations
    - Simulated Annealing: global optimal

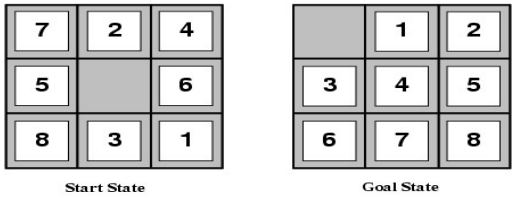

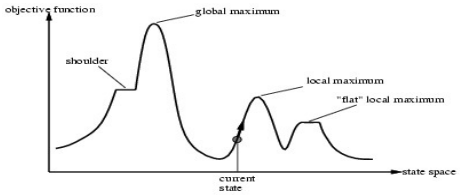

Which Algorithm?

- A*, Manhattan Heuristic

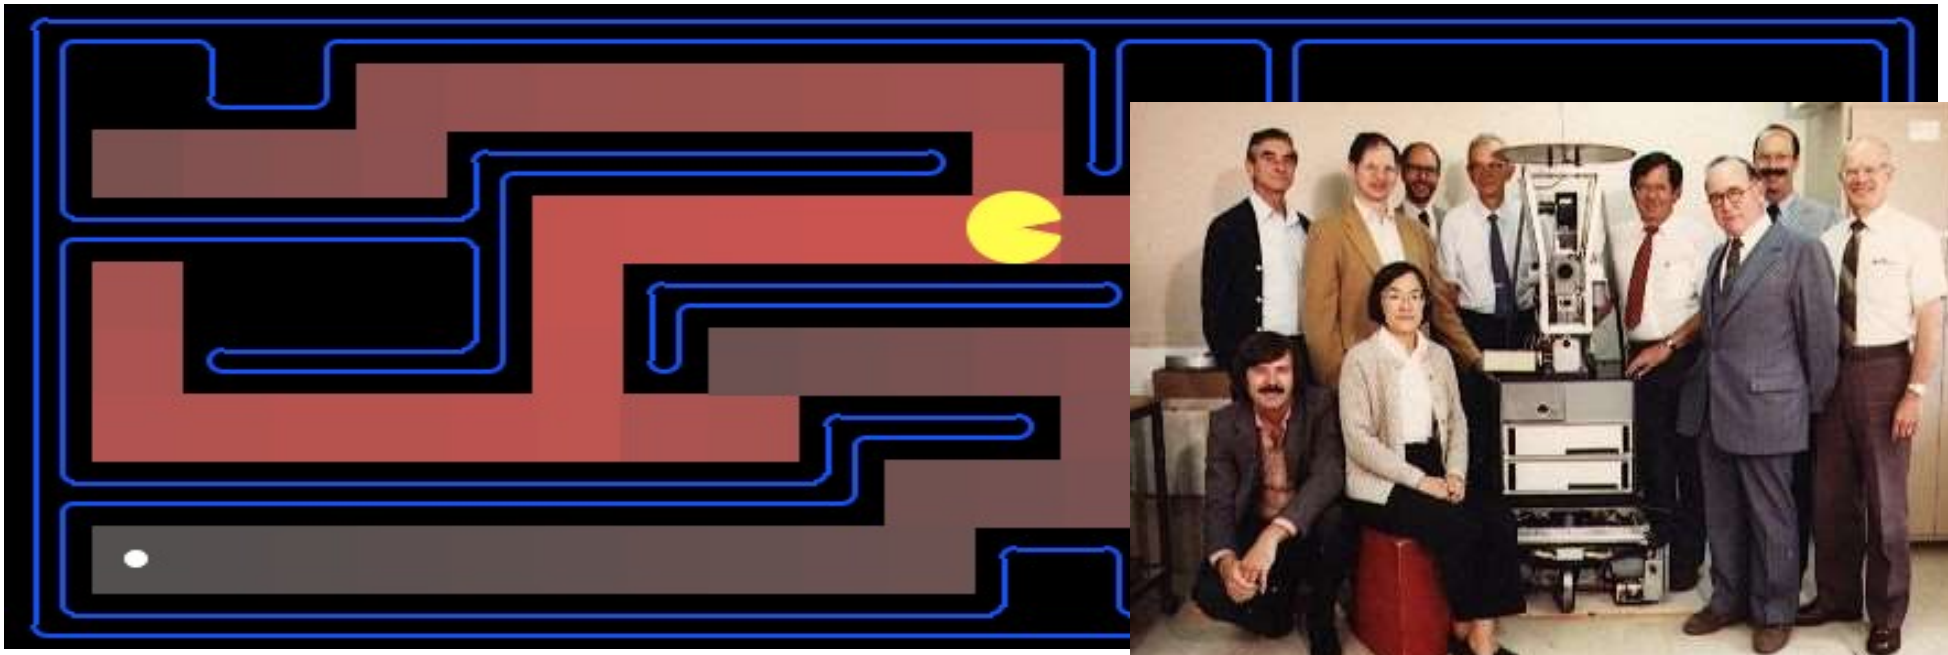

Adversarial Search

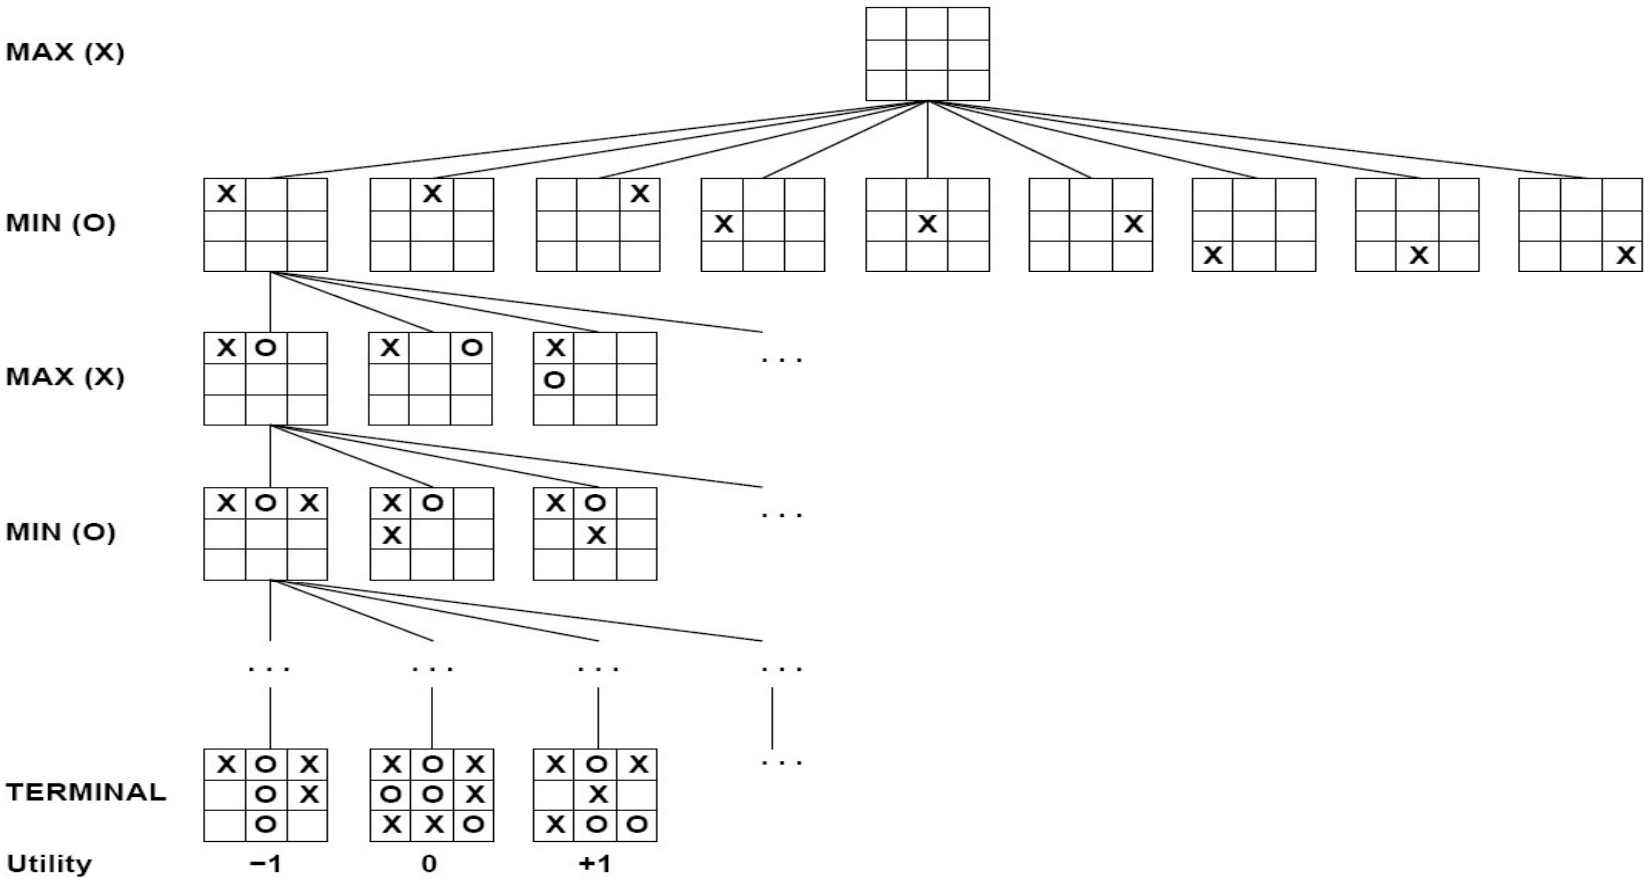

Adversarial Search

- AND/OR search space (max, min)
- minimax objective function
- minimax algorithm (~dfs)
  - alpha-beta pruning
- Utility function for partial search
  - Learning utility functions by playing with itself
- Openings/Endgame databases

Knowledge Representation and Reasoning

- Representing: what agent knows

  Propositional logic
  Constraint networks
  HMMs
  Bayesian networks
  …

- Reasoning: what agent can infer

  Search
  Dynamic programming
  Preprocessing to simplify

Search+KR&R Example: CSP

- Representation
  - Variables, Domains, Constraints
- Reasoning:
  - Arc Consistency (k-Consistency)
  - Solving
    - Backtracking search: partial var assignments
      - Heuristics: min remaining values, min conflicts
    - Local search: complete var assignments

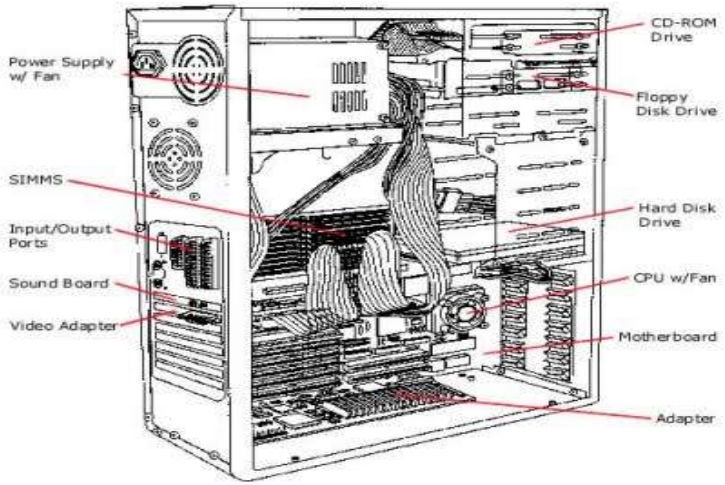

KR&R: Markov Decision Process

- Representation
  - states, actions, probabilistic outcomes, rewards

$$V^*(s) = \max_a Q^*(s,a)$$

$$Q^*(s,a) = \sum_{s'} T(s,a,s') \left[R(s,a,s') + \gamma V^*(s') \right]$$

- Reasoning: V*(s)
  - Expectimax
  - Value Iteration: dynamic programming
- Reinforcement Learning:
  - Exploration / exploitation
  - Learn model or learn Q-function?

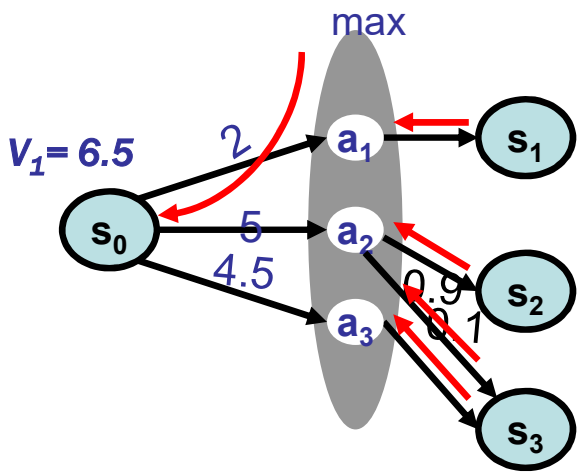

Pac-Man Beyond the Game!

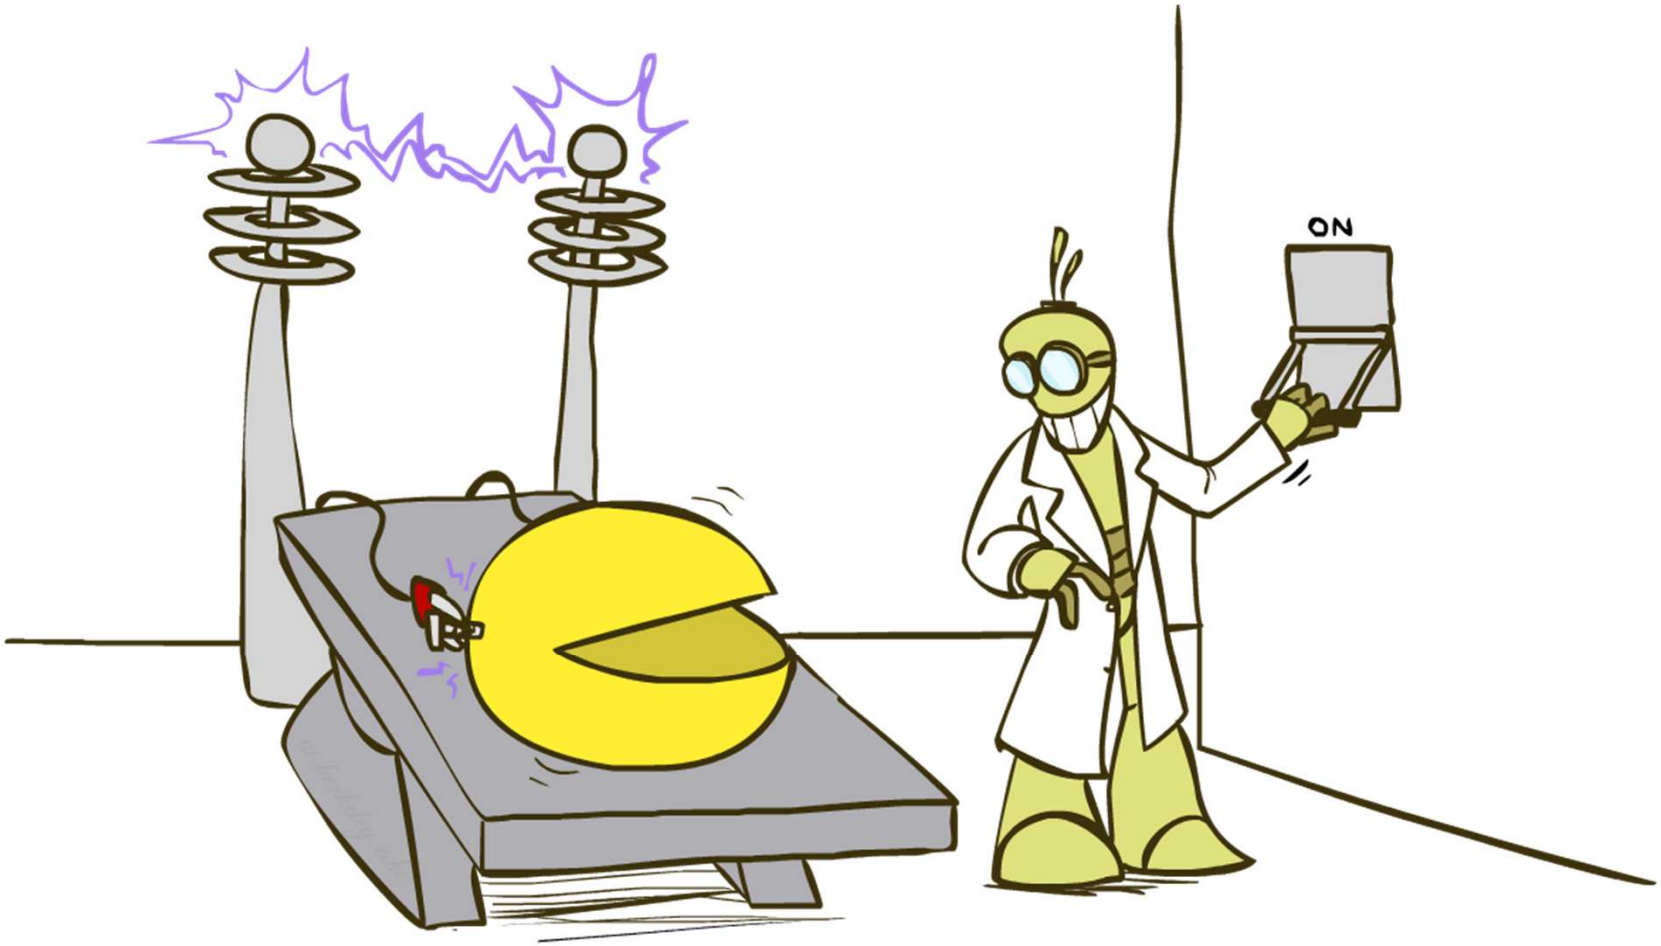

Pacman: Beyond Simulation?

ROOMBA PACMAN

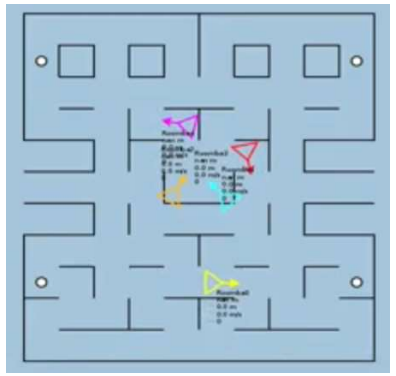

Students at Colorado University: http://pacman.elstonj.com

[VIDEO: Roomba Pacman.mp4]

Pacman: Beyond Simulation!

KR&R: Probability

- Representation: Bayesian Networks
  - encode probability distributions compactly
    - by exploiting conditional independences

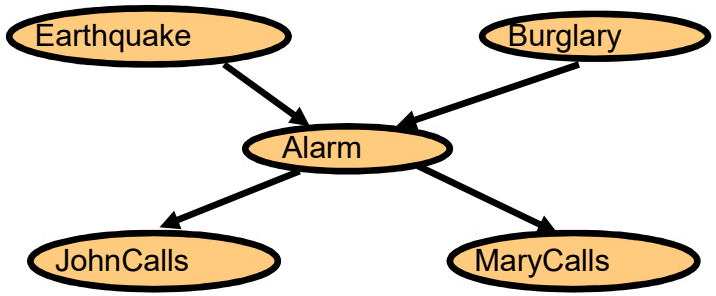

- Reasoning
  - Exact inference: var elimination
  - Approx inference: sampling based methods
    - rejection sampling, likelihood weighting, MCMC/Gibbs

KR&R: Hidden Markov Models

- Representation
  - Sequence model
  - One hidden state, one observation

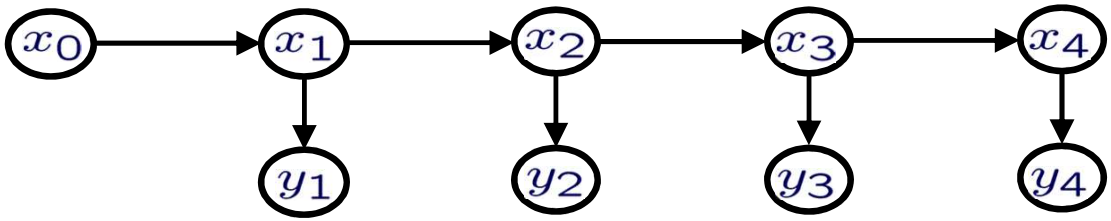

- Reasoning/Search
  - most likely state sequence: Viterbi algorithm
  - marginal prob of one state: forward-backward

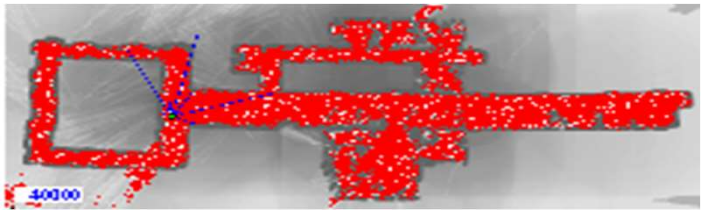

Personal Robotics

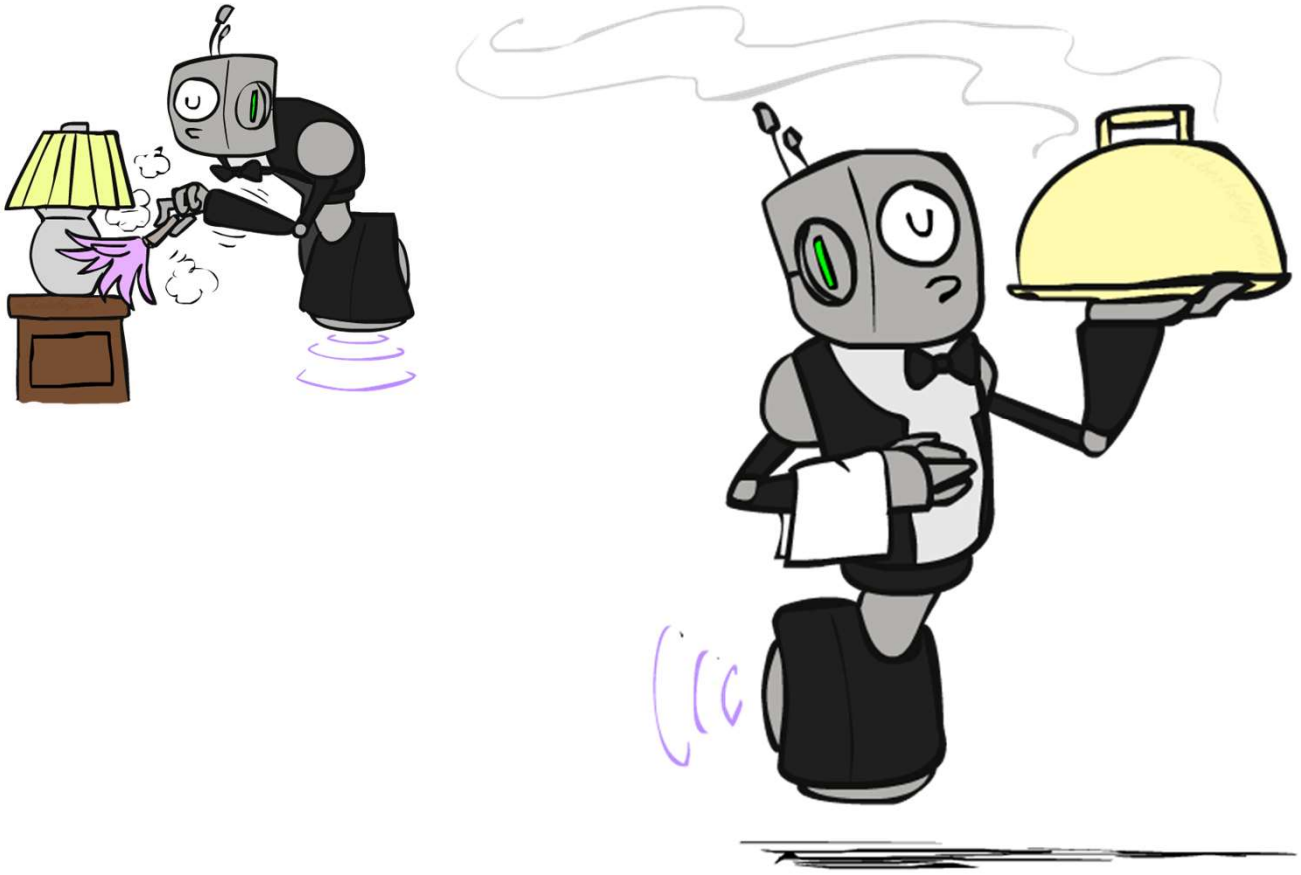

Where to Go Next?

That's It!

- Please help out with some course evaluations.
- Thanks to TAs Ben, Emilia, Kenny, Vardhman, Nicholas.
- Thanks to you all for your interest in AI and your participation in the course.
- Best wishes for the summer and after, and always maximize your expected utilities!

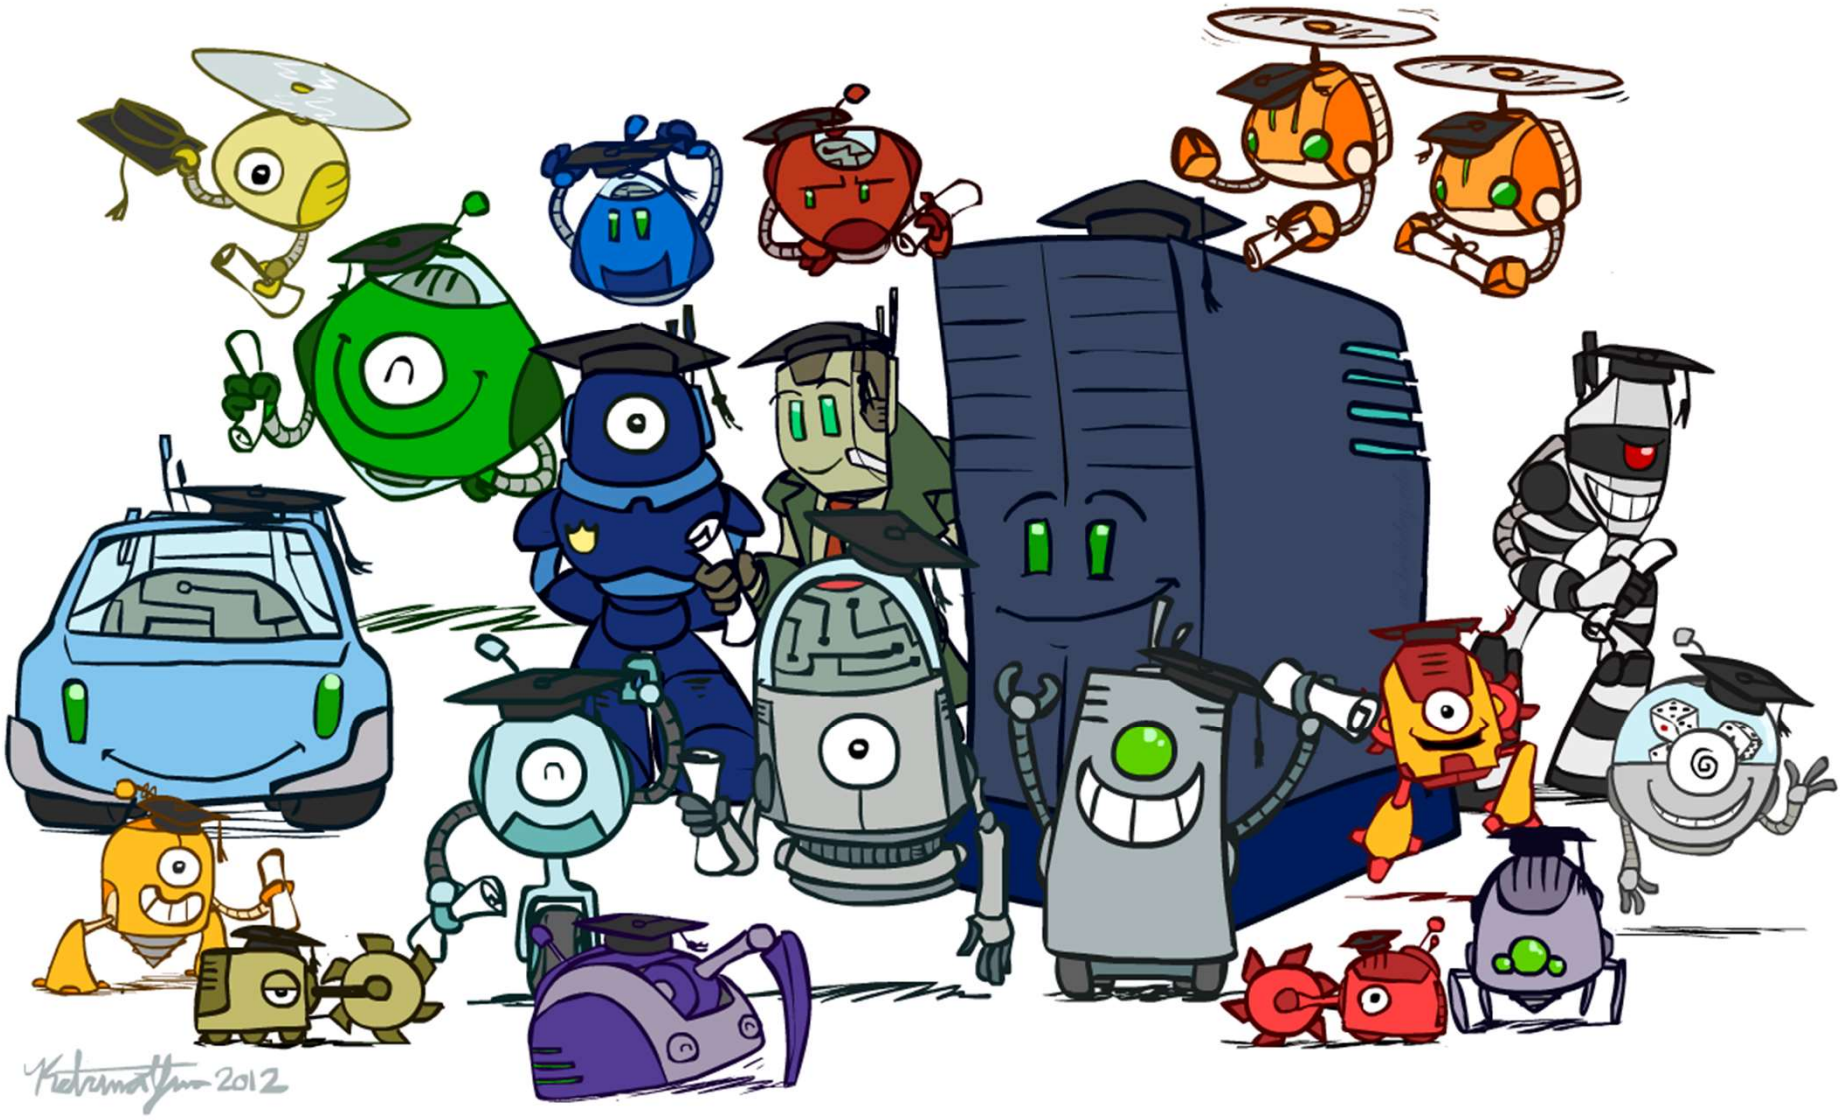
2012