CS440/ECE448 Lecture 35: Exam 3 Review

Mark Hasegawa-Johnson

Outline

- Mechanics
- Topics
- Course eval

Mechanics

1. Go to prairietest.org, and sign up for a final exam. There are three available: Two at 8am, one at 1:30pm on May 13.
2. Tuesday, May 13: Come to the exam room. Bring:
   – Device (computer or tablet)
   – ID
   – Pencils
   – Show up early, so we can check you in.

Bring your own device!

- Bring a laptop or tablet on which you will take the exam
- It should be a device with an internet connection, on which you can take exams at https://prairietest.org
- Cell phones are not guaranteed to work

Bring your UIUC ID!!!

- You will need your ID to check in before taking the exam
- You will need your ID to check out, and then check back in again, if you need a bathroom break
- You will need your ID to check out if you finish the exam early
- If you have a non-UIUC ID, you will have to wait in the slow check-in line

Bring your pencils!!!!

- We will provide scratch paper
- We will not provide pencils
- Bring your pencil

Bring nothing else!!!!

- Your pockets should be empty, except ID and pencils
  - Cell phone is in your backpack
  - Earphones are in your backpack
  - Baseball cap is backward on your head, or in your backpack
- Backpack is at the side of the room, on the floor
  - NOT under your own desk

Outline

- Mechanics
- Topics
- Course eval

Topics covered on the exam

- 8 questions: topics from exam 1
- 8 questions: topics from exam 2
- 2 questions: zero-sum and non-zero-sum games

Topics from Exam 1

- Probability
- Decision theory and fairness
- Naïve bayes
- Learning and linear regression
- Logistic regression and word2vec
- Nonlinear regression and relevance backprop
- Softmax and transformer
- HMMs and Bayesian networks

Topics from Exam 2

- Configuration space
- Search: BFS, DFS, A*
- Theorem proving
- Markov Decision Processes: Value & Policy Iteration
- Model-based RL, Exploration vs Exploitation
- Q-learning: TD, SARSA
- Policy learning (Imitation, Actor-Critic)
- Computer vision: Pinhole camera, vanishing point
- Convolution & Max-pooling

Minimax and Alpha-Beta

- Minimax search
  - Max node: $v(s) = \max_a v(\text{child}(s, a))$
  - Min node: $v(s) = \min_a v(\text{child}(s, a))$
- Limited-horizon computation and heuristic evaluation functions
  - It's impossible to search all the way to the end of the game!
  - Instead, search a fixed number of steps, then estimate $v(s)$ using the best approximation you can think of
- Alpha-beta search
  - Alpha is the highest score that Max knows how to force Min to accept
  - Beta is the lowest score that Min knows how to force Max to accept
  - If Beta ever falls below Alpha, prune the rest of the children
- Computational complexity of minimax and alpha-beta
  - Minimax is $O\{b^d\}$. With optimal move ordering, alpha-beta is $O\{b^{d/2}\}$.

Monte Carlo Search & Expectiminimax

- Monte Carlo search:

$$v(s) = \frac{(\text{\# times Max wins})}{N}$$

- Expectiminimax:

$$v(s) = \max_a \sum_{s'} P(S_{t+1} = s' | S_t = s, a) v(s')$$

$$v(s') = \min_{a'} \sum_{s''} P(S_{t+2} = s'' | S_{t+1} = s', a') v(s'')$$

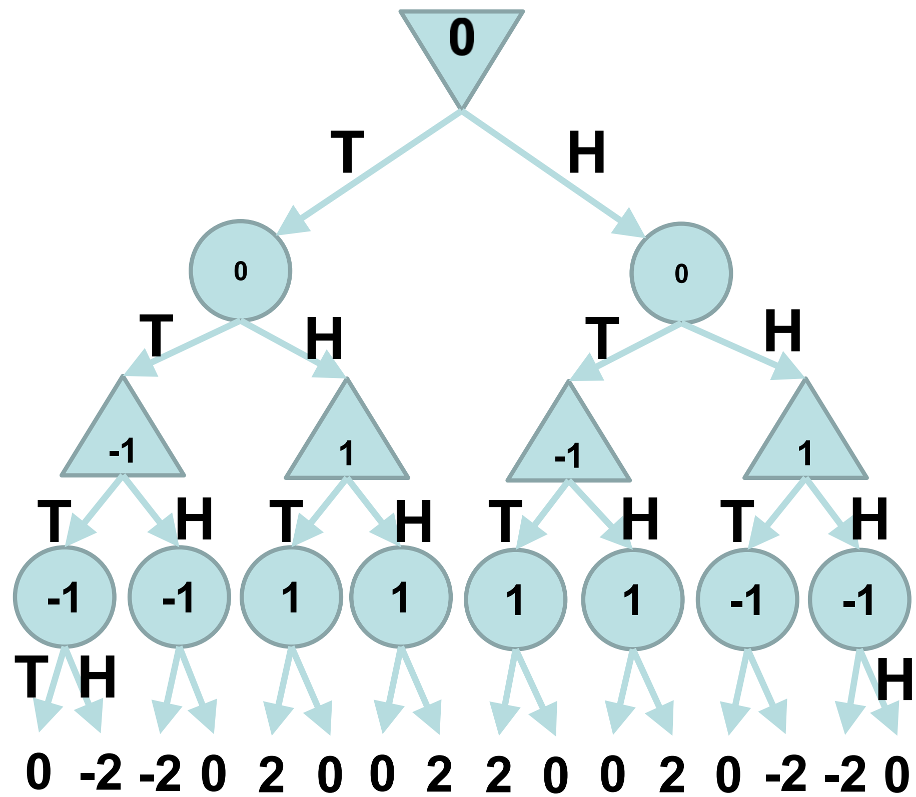

Non-Zero-Sum Games

- Dominant strategy
  - a strategy that's optimal for one player, regardless of what the other player does
  - Not all games have dominant strategies
- Nash equilibrium
  - an outcome (one action by each player) such that, knowing the other player's action, each player has no reason to change their own action
  - Every game with a finite set of actions has at least one Nash equilibrium, though it might be a mixed-strategy equilibrium.
- Pareto optimal
  - an outcome such that neither player would be able to win more without simultaneously forcing the other player to lose more
  - Every game has at least one Pareto optimal outcome. Usually there are many, representing different tradeoffs between the two players.
- Mixed strategies
  - A mixed strategy is optimal only if there's no reason to prefer one action over the other, i.e., if $0 \leq P_A \leq 1$ and $0 \leq P_B \leq 1$ such that:

$$(1-P_B)r_A(d,d) + P_B r_A(d,c) = (1-P_B)r_A(c,d) + P_B r_A(c,c)$$

$$(1-P_A)r_B(d,d) + P_A r_B(d,c) = (1-P_A)r_B(c,d) + P_A r_B(c,c)$$

Adversarial Learning and Mechanism Design

- Simultaneous gradient ascent:

$$\begin{bmatrix} z_A \\ z_B \end{bmatrix} \leftarrow \begin{bmatrix} z_A \\ z_B \end{bmatrix} + \eta \begin{bmatrix} \frac{\partial E[r_A]}{\partial z_A} \\ \frac{\partial E[r_B]}{\partial z_B} \end{bmatrix}$$

- Nash equilibrium: $\frac{\partial E[r_A]}{\partial z_A} = \frac{\partial E[r_B]}{\partial z_B} = 0$
- Every game has a Nash equilibrium, but not every game has a stable Nash equilibrium!
- Mechanism design: Adjust $\boldsymbol{R}_A$ and $\boldsymbol{R}_B$ to get desired player behavior

How to study

Recommended method:

- Do both the required and optional problems from daily quizzes
- Do each one a few times, to make sure you've seen all of the variants!
- All exam questions will be from the quizzes, so this is the best guaranteed method to do well on the exam.

Other possibilities:

- There is an older sample exam on the web page with other variants of some problems
- Machine problems study the same problems in other ways

Outline

- Mechanics
- Topics
- Course eval

Course evaluation

- Course evaluations are useful for future course design!
- Some types of homework have been eliminated, others have been created, some have had their point totals reduced, in response to student evaluations
- Some lecture topics have been expanded, others reduced
- Please participate, tell us what you think!

https://ices.citl.illinois.edu/