Playful AI Education

Todd W. Neller
Gettysburg College

Introduction

- Teachers teach best when sharing from the core of their enjoyment of the material.
  - E.g. Those with enthusiasm for graphics should use graphical examples.
- Game playing is one of my favorite hobbies, so I often utilize games in my teaching.
- “Play is our brain's favorite way of learning.” - Diane Ackerman

Outline

- What makes a good game for teaching?
- What are non-game-tree search examples of game use in AI teaching?
- What are some future opportunities and challenges for playful AI education?

Good Games Have Simple Rules

- Breakthrough is better than Chess for teaching game-tree search.
- Chess:
  - Many different piece movement rules
  - En passant
  - Castling
  - 50 move rule
- Breakthrough
  - One piece movement rule

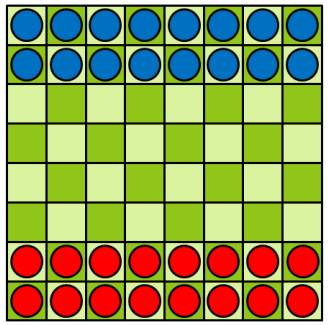

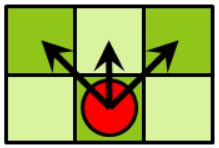

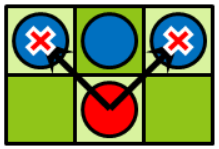

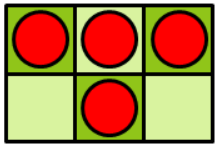

- Unless the goal is to teach complex modeling, prefer games with simpler rules (e.g. prefer Hex to Go as a connection game).

… Simple Rules and Fun Depth

- Tic-Tac-Toe has simple rules, but isn't particularly fun once you see to its shallow depth.
- *Depth* is a term in game design that correlates with the expected range of ELO scores.
- The best games reward players that explore the subtle terrains of these microworlds.
  Learning → winning

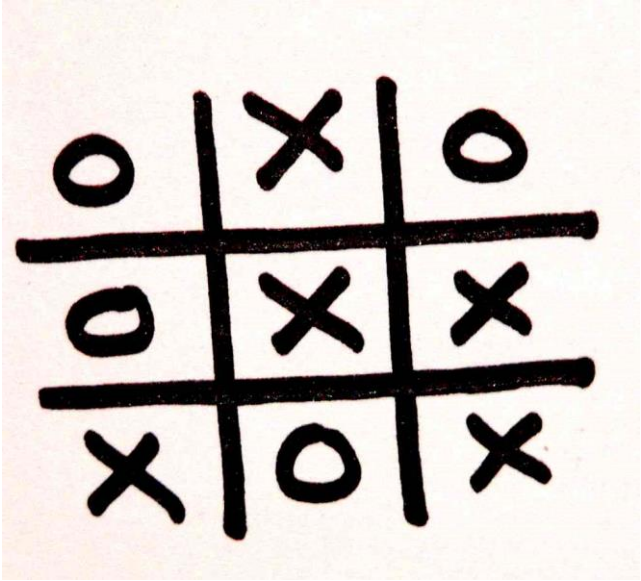

https://kidavalanche.wordpress.com/2010/02/03/tic-tac-toe-magic-trick-prediction/

Beyond Game-Tree Search

- Constraint Satisfaction Problems
- Logical Reasoning
- Planning
- Uncertain Reasoning
- Machine Learning
- Robotics
- Etc.

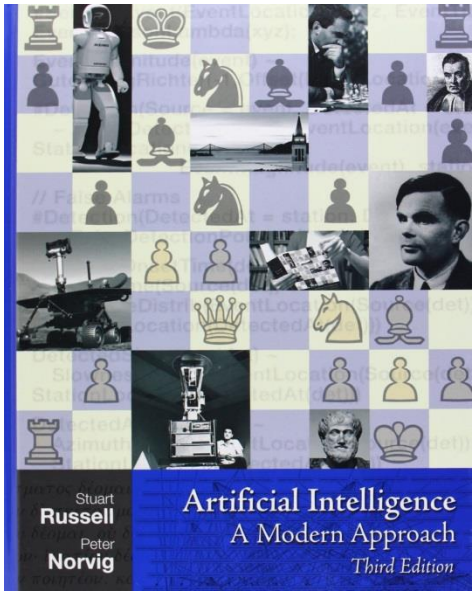

+

Playful Robotics

- FIRST Robotics Competition (FRC) (1992-present)
- RoboCup robot soccer (Pre-RoboCup-96 & RoboCup-97 - present)
- DARPA Grand Challenges: autonomous driving (2004, 2005, 2007)
- Many physical sport games become interesting testbeds for robotics.

Clue/Cluedo *Case* Example

The Game of Clue

- 21 cards: 6 suspects, 6 weapons, 9 rooms
- Case file has unknown, random suspect, weapon, and room (SWR)
- Remaining cards dealt to players
- Player *suggests* SWR, first player clockwise that can refute, must show card
- Each player can make 1 SWR *accusation*
- Correct → win; incorrect → lose (& refute)

More Than Child's Play

- Children mark off each dealt/shown card
- Small notepad + instruction = trivial play
- However, consider this:
  - You know player **A** has 3 cards, two of which are **t** and **u**.
  - Player **A** refuted player **B**'s and **C**'s suggestions of (**v**,**w**,**x**) and (**y**,**z**,**x**), respectively, by showing a card.
  - Therefore, … **A**'s 3rd card must be **x**.

Constraint Satisfaction

- Clue reasoning is constraint satisfaction.
- One formulation: Boolean variables c_p denoting "Card *c* is in place *p*."
- Given CNF representation of Boolean constraints, reason with SAT solver refutations
- See Model AI Assignment for details (http://modelai.gettysburg.edu/2011/clue)

Basic Logic Concepts

- Concepts: sentences, operators, literals, truth assignments, (un)satisfiability, models, validity, tautologies, entailment, logical equivalence, derivation, soundness, completeness, …
- That's covering a *lot* without first-order logic.
- Simple, minimalist, high-utility approach
  - Con: No predicates, unification, FOL generalizations.
  - Pro: Time-efficient, experiential learning.

ClueReasoner

```
ClueReasoner cr = new ClueReasoner();
String[] myCards = {"wh", "li", "st"};
cr.hand("sc", myCards);
cr.suggest("sc", "sc", "ro", "lo", "mu", "sc");
cr.suggest("mu", "pe", "pi", "di", "pe", null);
cr.suggest("wh", "mu", "re", "ba", "pe", null);
cr.suggest("gr", "wh", "kn", "ba", "pl", null);
cr.suggest("pe", "gr", "ca", "di", "wh", null);
cr.suggest("pl", "wh", "wr", "st", "sc", "wh");
cr.suggest("sc", "pl", "ro", "co", "mu", "pl");
cr.suggest("mu", "pe", "ro", "ba", "wh", null);
cr.suggest("wh", "mu", "ca", "st", "gr", null);
cr.printNotepad();
```

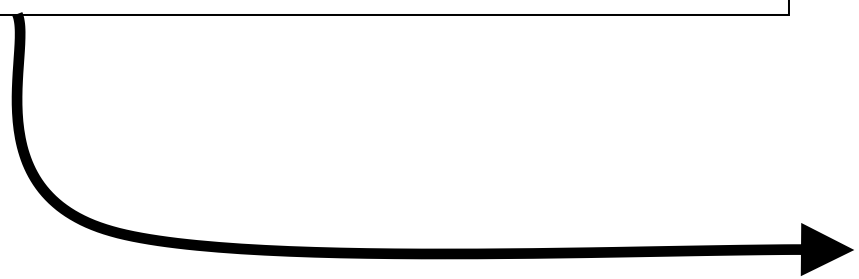

	sc	mu	wh	gr	pe	pl	cf
mu	n	-	-	n	-	-	-
pl	n	Y	n	n	n	n	n
gr	n	n	Y	n	n	n	n
pe	n	-	n	n	-	-	-
sc	n	Y	n	n	n	n	n
wh	Y	n	n	n	n	n	n
kn	n	-	-	-	n	-	-
ca	n	n	n	Y	n	n	n
re	n	-	-	n	-	-	-
ro	n	-	-	-	-	-	-
pi	n	-	n	n	-	-	-
wr	n	-	-	-	-	-	-
ha	n	-	-	-	-	-	-
lo	n	-	-	-	-	-	-
di	n	n	n	n	-	n	-
ki	n	-	-	-	-	-	-
ba	n	-	-	n	n	-	-
co	n	-	-	-	-	-	-
bi	n	-	-	-	-	-	-
li	Y	n	n	n	n	n	n
st	Y	n	n	n	n	n	n

At-Least Constraints

- One constraint omitted from CNF SAT representation: Player *p* has exactly *n* cards.
- A CNF SAT clause is true iff at least **one** literal is true. Generalizing…
- An *at-least constraint* clause is true iff at least ***n*** literals are true.
- Example, for each player in a 6 player game, each player holds exactly 3 cards, so…
  - at least 3 cards are in a player's hand, and
  - at least 18 cards are not in a player's hand.

Gettysburg College Faculty/Student Projects with At-Least Constraints

- Mixed logical and probabilistic estimation:
  - Generalizations of WalkSAT, DPLL, and Knuth's Dancing Links algorithms
  - Empirical study of SLS sampling bias in estimation of card location probabilities
- Explanation generation (auto-generation of natural language proof of case-file contents)
- Study of human vs. computer reasoning patterns

The Dice Game of Pig

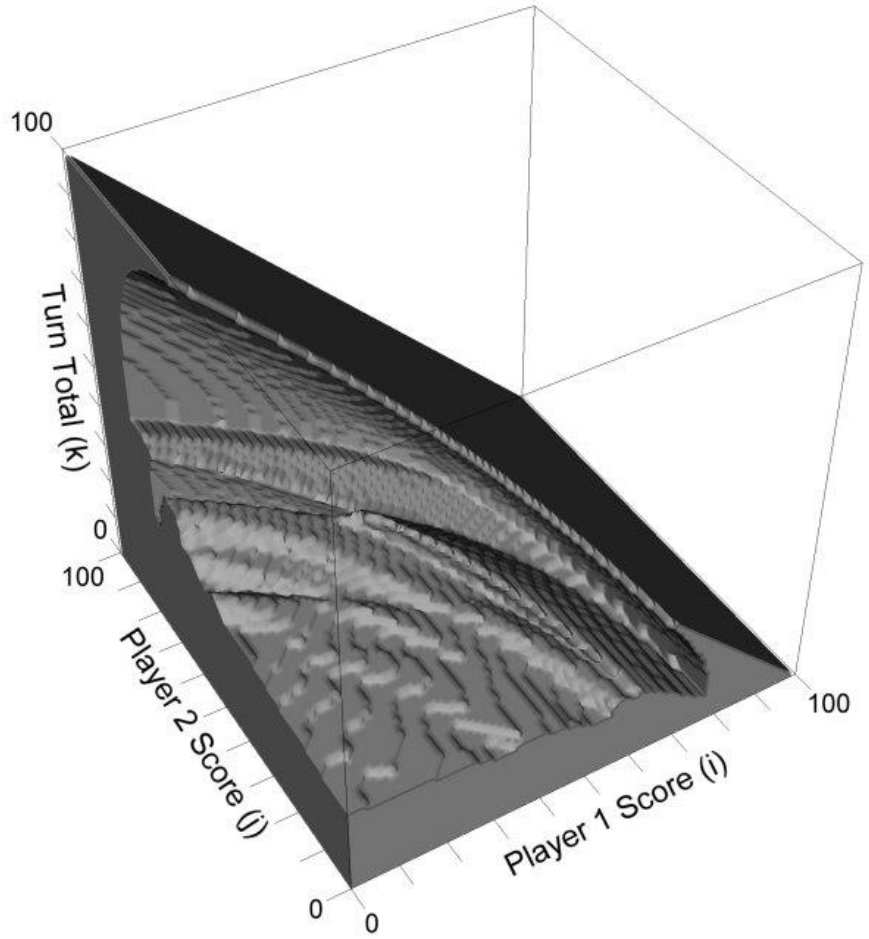

You never sausage a simple and fun game!

Pig: Simple Yet Fun and Highly Useful

- Rules:
  - The first player reaching 100 points wins.
  - On each turn, a player rolls a die as many times as desired until either the player holds and scores the sum of the rolls, or rolls a 1 and scores nothing.
- High Fun-to-SLOC (source lines of code) ratio
- Many uses for teaching probabilities in Mathematics, control structures and incremental build model in CS1, simple GUI development, and in AI…

Pig in AI Education

- Reinforcement Learning:
  - Value iteration for computation of optimal play
  - Reinforcement learning algorithm results compared to value iteration solution (e.g. On-/Off-Policy Monte Carlo, TD-learning, etc.)
- Supervised Learning:
  - Regression fit to collected human roll/hold data
- Unsupervised Learning:
  - Examination of clusters of suboptimal human plays

Games Throughout AI

- Bayesian reasoning? Estimating an opponent's hand in Gin Rummy
- Computer vision? Autonomous driver races, robot soccer, Zachary Dodds' Set gameplay with OpenCV
- Natural Language Processing? Computer play of interactive fiction
- There is nothing in AI we cannot make playful.

Credit: Tim Mossman

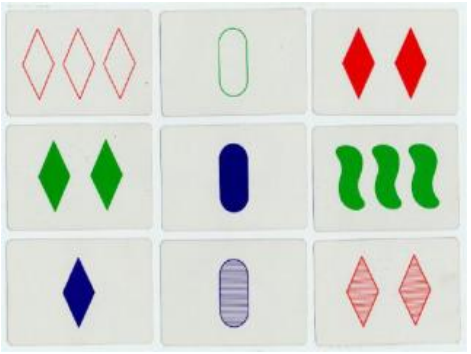

Credit: Zachary Dodds

```
>u
It is pitch dark.  You are likely to be eaten by a grue.
>what is a grue
The grue is a sinister, lurking presence in the dark places of the
earth.  Its favorite diet is adventurers, but its insatiable
appetite is tempered by its fear of light.  No grue has ever been
seen by the light of day, and few have survived its fearsome jaws
to tell the tale.
>_
```

Bringing People Together

- Activities that bring people together in cultures universally:
  - Working together
  - Eating together
  - Playing games together
- So what is a game in its most general sense?

Game Definition

- From diverse definitions, I clustered related terms/concepts:
  - system/activity/form of art/play
  - players/decision-makers/forces/voluntary participants
  - goal/outcome/objective/state of affairs
  - rules/structure/limiting context/resource management
  - voluntary/accept limitations
- My consensus game definition: **“a voluntary activity where players pursue goals according to an agreed-upon set of rules.”**

Generalizing

- In truth, most any activity we enjoy can be gamified, e.g.
  - Protein folding (fold.it)
  - Fitbit step average rankings among friends
- In a general sense, we can bring a playfulness to most any activity.
  - Play can be cooperative rather than competitive as well.

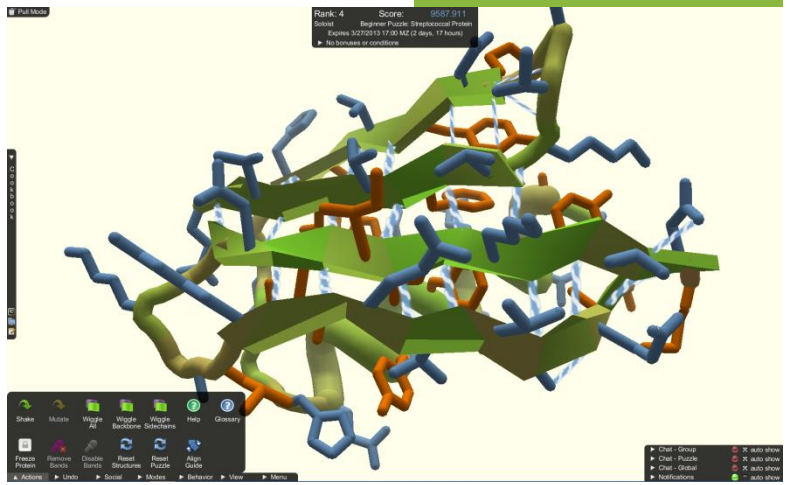

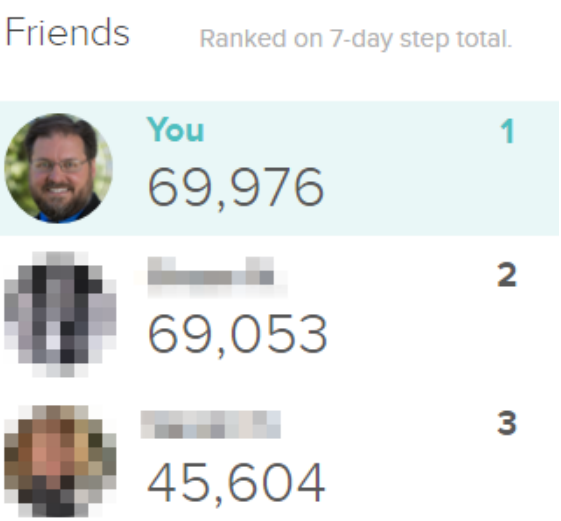

Sharing Inspiration

- Games are not the only means of bringing joy to AI education.
- You (and your colleagues) undoubtedly bring unique enthusiasms, hobbies, experiences, skills, etc. to your own teaching of AI.
- How have you (or a colleague) brought what you enjoy most into your teaching of AI?

Future Opportunities and Challenges

- Beyond “Computers can play X better than humans” there are other opportunities:
  - Computer-aided game design
  - Computer game trainer (human play modeling, coaching/teaching)
  - Computer opponent rank-matching
- And there are challenges:
  - Computer-aided cheating
  - Game AI assignment plagiarism

Conclusion

- I greatly enjoy games and delight in finding ways games can be used to playfully teach AI and CS in general.
- However, teachers teach best when they enjoy what they share.
- Teach to your strengths and enthusiasms!
- Thank you for this award and the joy of sharing playful AI education through the years.

Questions?