


Optimal, Approximately Optimal, and Fair Play of the Fowl Play Card Game




Todd W. Neller, Marcin Malec,
Clifton G. M. Presser, and Forrest Jacobs

Gettysburg
COLLEGE
Computer Science


Outline

- Introduction to the Fowl Play jeopardy card game
- Computation of optimal play
- Neural network approximation of optimal play
- Red Light: computer-aided redesign of Fowl Play optimizing fairness



Fowl Play

- Designed by Robert Bushnell and published by Gamewright in 2002
- Object: Be the first of 2+ players to reach a given goal score.
 - We consider 2 players, 50 point goal case.
- Materials: Shuffled 48 card deck with
 - 42 chicken cards (each incrementing turn total)
 - 6 wolf cards (each causing loss of turn total)



Fowl Play (cont.)

- On each turn, draw one or more cards until:
 - you draw a wolf and score no points, or
 - you hold and score the number of chickens you've drawn.
- After the last (6th) wolf is drawn, reshuffle all cards.



Optimizing Score

- 5 variables: Let i and j be the current player and opponent scores, respectively. Let k be the current turn total. Let w and c be the number of wolves and chickens drawn since the last shuffle, respectively.
- Let c_{rem} and w_{rem} be the number of chickens and wolves remaining, respectively.
- Balancing reward and risk, draw iff:

$$\frac{c_{\text{rem}}}{w_{\text{rem}} + c_{\text{rem}}} > \frac{w_{\text{rem}}}{w_{\text{rem}} + c_{\text{rem}}} k$$

- Simplifying:

$$c_{\text{rem}} > w_{\text{rem}} k$$

Optimizing Win Probability

- Assuming optimal play, non-terminal win probabilities are defined as:

$$P_{i,j,k,w,c} = \begin{cases} P_{i,j,k,w,c,\text{draw}} & k = 0, \text{ and} \\ \max(P_{i,j,k,w,c,\text{draw}}, P_{i,j,k,w,c,\text{hold}}) & \text{otherwise.} \end{cases}$$

where $P_{i,j,k,w,c,\text{draw}}$ and $P_{i,j,k,w,c,\text{hold}}$ are the probabilities of winning if one draws and holds, respectively. These probabilities are given by:

$$P_{i,j,k,w,c,\text{draw}} = \frac{c_{\text{rem}}}{w_{\text{rem}} + c_{\text{rem}}} P_{i,j,k+1,w,c+1} + \frac{w_{\text{rem}}}{w_{\text{rem}} + c_{\text{rem}}} \begin{cases} 1 - P_{j,i,0,w+1,c} & w < w_{\text{total}} - 1, \text{ and} \\ 1 - P_{j,i,0,0,0} & \text{otherwise.} \end{cases}$$

$$P_{i,j,k,w,c,\text{hold}} = 1 - P_{j,i+k,0,w,c}$$

- Equations are solved using a generalization of value iteration.

Neural Network Approximation

- The storage of the optimal play would be problematic on modern mobile and embedded systems.
- There are a number of possible design choices for the function approximation task.
- Multi-layer feed-forward neural networks showed the greatest promise for closely approximating optimal play with minimal memory requirements.

Neural Network Structure

- **Network Input:** In addition to the input features i, j, k, w_{rem} , and c_{rem} , the features $c_{rem}/(w_{rem} + c_{rem})$ and $(w_{rem} \times k)/(w_{rem} + c_{rem}) \times k$, the probability of drawing chicken and wolf cards respectively, scaled to range $[-1, 1]$ aided in the function approximation.
- **Network Output:** The single output unit classifies *draw* or *hold* actions according to whether the output is above or below 0.5, respectively. For training purposes, target outputs are 1 and 0 for *draw* and *hold*, respectively.
- **Hidden Layer:** Optimizing the number of hidden units empirically so as to boost learning rate and generalization, the single hidden layer worked best with 13 units.
- **The activation function** for all hidden and output units is the logistic function $f(x) = 1 / 1 + e^{-x}$

Neural Network Training

- Network weights are initialized using the Nguyen-Widrow algorithm.
- Online learning is used with the training cases generated through simulations of games with optimal play.
- Backpropagation learning algorithm is then applied with learning rate $\alpha = 0.1$ to update network weights.
- After 100,000 games the performance of the network is evaluated through network play against an optimal player for 1,000,000 games. If the win rate is better than 49.5% and the highest win rate so far, we re-evaluate using policy iteration with $\varepsilon = 1 \times 10^{-6}$. We then continue to alternate between training and evaluation for total of 400 iterations.

Optimal Komi

- Assuming optimal player play, we have computed the probability of winning with the second player starting with any number of points.
- We can thus find the komi (compensation points) that optimize fairness for Fowl Play:
 - current: 2nd player starts with 0 points → 1st player wins 52.42% of games.
 - best: 2nd player starts with 1 point → 1st player wins 50.16% of games.

Optimizing Fowl Play for Fairness

- Further, we can
 - vary the number of good (chicken) and bad (wolf) outcomes,
 - compute optimal play for each game,
 - compute komi optimizing fairness, and
 - thus optimize the game design for fairness.
- With our optimized game of Red Light, the first player wins 50.001% assuming optimal play.

Red Light

- Object: Be the first of 2 players to reach 50 points. The 2nd player begins with 1 point.
- Materials: 28 poker chips in a bag with
 - 24 green “green light” chips (each incrementing turn total)
 - 4 red “red light” chips (each causing loss of turn total)



Red Light (cont.)

- On each turn, draw one or more chips until:
 - you draw a red light and score no points, or
 - you hold and score the number of green lights you've drawn.
- After the last (4th) red light is drawn, all chips are returned to the bag and reshuffled.
- Alternate materials:
Standard ("French") deck of playing cards using Ace (red light) and 2-7 (green lights) of each suit.



Conclusion

- Fowl Play is a simple yet non-trivial jeopardy card game. We have:
 - computed optimal play for 2-players, 50 point goal
 - computed a good NN approximation of optimal play
 - optimized the game for fairness, creating the variant Red Light
- Red Light presents
 - a challenge for identifying features for good human play, and
 - an accessible, minimalist game of chance for research and teaching.