“Do I Hear 3NT?”: Learning a Bridge Bidder

Greg White (gregw@stanford.edu)

Predicting

- This project attempts to learn the optimal classification for two important sub-problems of Bridge’s bidding phase
  - Given a 13-card hand, which of 36 possible opening bids should the player use?
  - Given two 13-card hands, in which of 36 possible final contracts should the team end?
- An opening bid and a final contract can each be one of Pass or [1-7][NT|S|H|D|C]
- Bidding in Bridge is a nuanced problem that takes human players a long time to learn: possibly extensible to state-based or logic-based AI with learning underpinning

Data

- Data is from tournament hands played by human experts and other computer agents available from a variety of sources that publish results of their tournaments
- Opening bid data starts with 13 card hands for each of 4 players and sequence of bids
  - Translate to hand-bid labeled data for each player up to and including first non-Pass bid
  - Possible multiple labels: only training and testing if no disagreement, reporting two error figures
- Final contract data requires processing to label ground truth (solving full-information game tree problem, uses 3rd party library)

Models and Features

- Softmax regression using two types of features. Maximize log-likelihood:

$$l(\theta) = \sum_{i=1}^{m} \log \prod_{l=1}^{36} \left(\frac{e^{\theta_l^T x^{(i)}}}{\sum_{j=1}^{36} e^{\theta_j^T x^{(i)}}} \right)^{1\{y^{(i)}=l\}}$$

  using stochastic gradient ascent:

$$\theta := \theta + \alpha \left(\begin{bmatrix} 0 \\ \vdots \\ 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix} - \frac{1}{1+\sum_{j=1}^{35} e^{\theta_j^T x^{(i)}}} \begin{bmatrix} e^{\theta_1^T x^{(i)}} \\ \vdots \\ e^{\theta_{y^{(i)}-1}^T x^{(i)}} \\ e^{\theta_{y^{(i)}}^T x^{(i)}} \\ e^{\theta_{y^{(i)}+1}^T x^{(i)}} \\ \vdots \\ e^{\theta_{35}^T x^{(i)}} \end{bmatrix} \right) (x^{(i)})^T$$

- Two feature extractors:
  - Raw card indicators: for all 52 cards
  - Limited domain knowledge indicators: # cards per suit, number of high card points per suit and total

Results

- Opening bid problem: low error, model generalizes well, and misclassifications are on borderline decisions for human players (and many of those classify correctly)

	Train Error		Dev Error	
	All	Agreed Labels	All	Agreed Labels
m	16.6K	9.6K	4.7K	2.7K
Raw feat.	16.0%	10.4%	16.9%	11.1%
Domain feat.	10.3%	5.3%	12.5%	**6.6%**

- Final contract problem has not had positive results yet (see discussion and future work)

Discussion

- Opening bid results with simple domain features are better than expected
- Raw card features generalize better than expected
- Problem may be simpler than seems
  - Fewer than 36 classes with large probabilities
  - Relatively rote decision for human players, more nuance introduced with sequential bids
- Final contract model oversimplifies and predicts few classes, with additional time refining would be a key next step

Future Work

- Final contract problem: additional features, debugging model/algorithm
- Sequential bidding agent: incorporate state- and/or logic-based model
- Learning vs. programming: in competition, must follow explainable conventions

References

- Mernagh, Michael, “Learning a Double Dummy Bridge Solver”, 2016 CS229 Final Project
- Haglund, Bo, “DDS Double Dummy Solver”, https://github.com/dds-bridge/