Learning a Double Dummy Bridge Solver

Michael Mernagh (mernagh@stanford.edu)

Predicting

In the game of Contract Bridge, agents need to predict how many tricks they will win, in both the bidding and trick-taking stages of the game. We found a repository of 300,000 professional games, and used that data to predict the final score in other professional games, using ensemble regression models and softmax regression. We obtained passable accuracies (around 87%) for the most difficult-to-predict games. This is a significant improvement over attempts to use neural networks for the same task (around 35% accuracy) [1]!

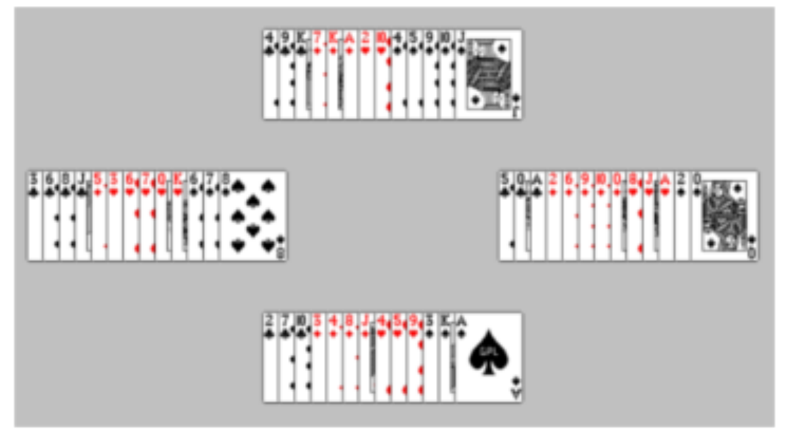

Goal: Given a hand like this, predict the final score.

Data

The original data was scattered through files in hundreds of zip files, where each file contained thousands of professional games, stored in PBN format (a text format that describes the entire game). We built a bridge game simulator that replayed each bridge game, updating the state of the game with each card played, and then outputting a customized hand representation like

```
4 6.J6.A82.QT98 AQ.AQ.Q53.J76 K987.97.KT76. JT.KT8.J94.AK
```

Results

Softmax regression had very good variance, but incredibly poor bias. The ensemble methods in general had more variance, but also had much better bias. The random forest methods did not generalize as well as gradient boosting regression. The regression methods decreased in accuracy as the hand sizes decreased. This is likely due to several factors: the ease of predicting the result actually decreased with the hand size, since results that were easily predicted by humans were excluded from the original data set, which also influenced the second reason: the size of the data set decreased with the hand size.

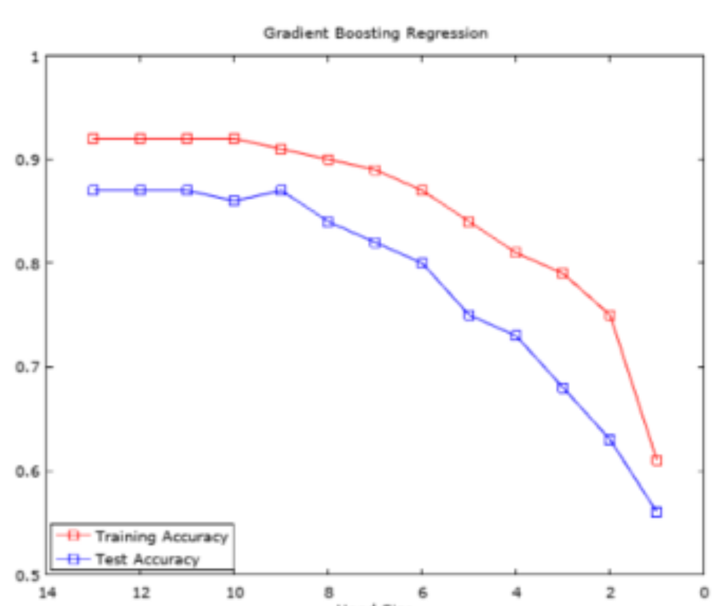

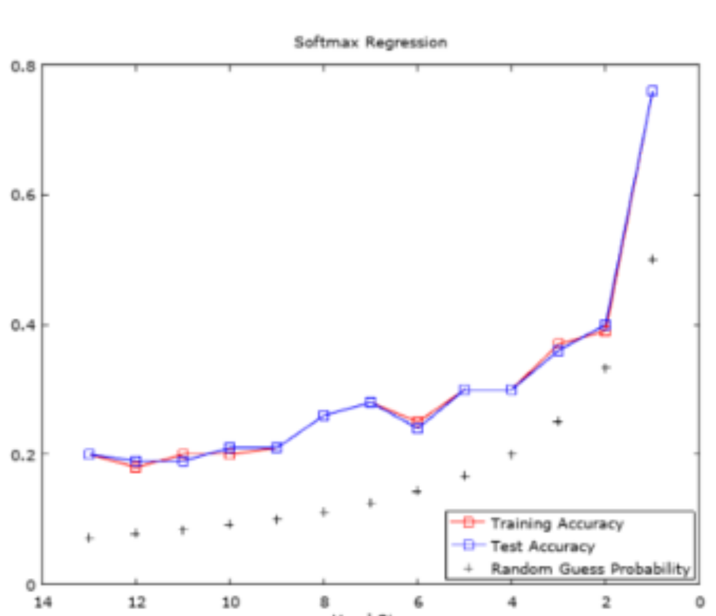

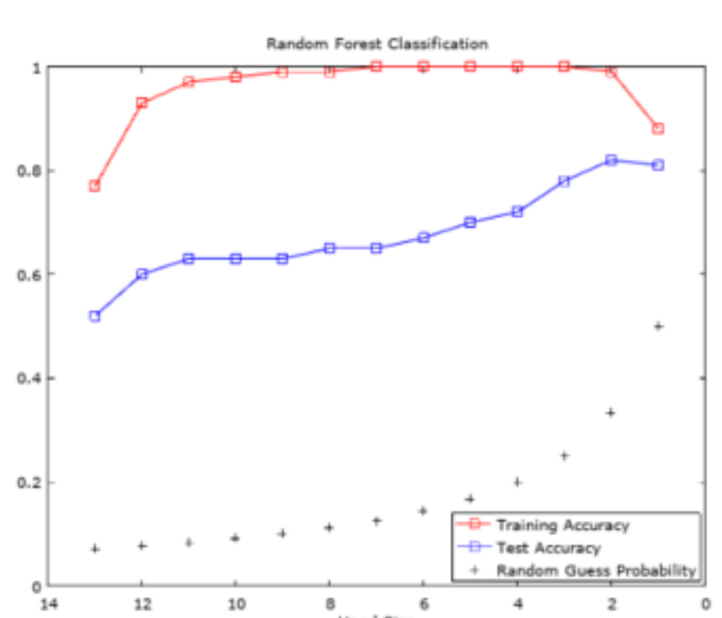

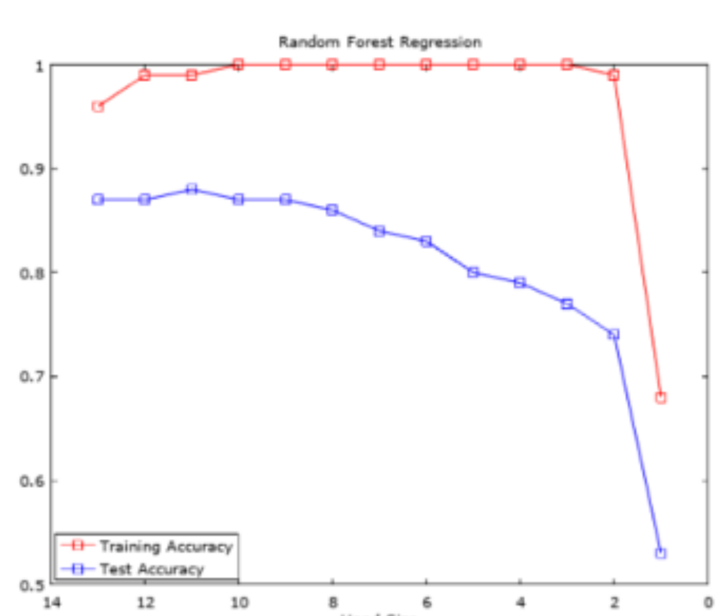

Features

For each player we used:

- The number of cards they held in each suit
- The rank (ace, king, etc) of the highest card in each suit (normalized for already played cards)
- The sum of the ranks of the cards they held in each suit, normalized for already played cards
- Stopper ranks for each suit. A contrived feature that indicates the rank of the stopper card in each suit, where a stopper card would stop an opponent from a run in that suit.
- Sure winners in each suit. The number of cards that player could successively play and be guaranteed to win, in a given suit.

These features are related to indicators that expert human players use in conjunction to determine a strategy of cards to play in order to know how many tricks they will win. Our feature vector was of size 80.

Models

We started with Softmax Regression, where we used gradient ascent to maximize

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

We then used ensemble methods (Random Forests and Gradient Boosting). In Random Forests we take the mean prediction of n decision trees learned on random subsets of the data. In Gradient Boosting, we seek to minimize

$$\sum_{i=1}^{n} \left[g_i f_t(x_i) + \frac{1}{2} h_i f_t^2(x_i) \right] + \Omega(f_t)$$

where $g_i = \frac{\partial}{\partial y_i^{(t-1)}} (y_i - y_i^{(t-1)})^2, h_i = \frac{\partial^2}{\partial y_i^{(t-1)}} (y_i - y_i^{(t-1)})^2$

by successively fitting decision trees to optimize this value, and updating our predictions based on the leaf values.

Future Work

Better data for the smaller hands should be generated or collected, and increasing the amount of data might help too. It may also be worth attempting to model the data as a hidden Markov model, since this may capture the fact that professionals may be imperfectly playing the hands. Given the nature of the data, experimenting with a Q-learner and a minimax player may also generate fairly accurate predictions.

References

[1] Mossakowski, Krzysztof, and Jacek Mandziuk. "Learning without human expertise: a case study of the double dummy bridge problem." IEEE Transactions on Neural Networks 20.2 (2009): 278-299.