

What is Curling?

Curling is a sport where two teams alternate throwing 16 stones across a sheet of ice towards target circles, called the house. Score is calculated at each end of a period, called an end, by the number of stones closest to the center.





Figure 1. Last position of an end, yellow scores +2.

Problem Description


Pertinent subproblems to this project are:

- Evaluating states accurately and quickly
- Predicting human performance
- Comparing search and machine learning techniques on the above tasks

Input



Score 4-2
Red's Turn



Score 2-2
Red's Turn

Output

52.9% +2 for red
58.9% win for red

Predicted Score 4-4

45.9% -1 for red
41.9% win for red

Predicted Score 3-2

Figure 2. From a state, winning probabilities and predicted score differences are outputted.

Comparison of Approaches

	Search Bot (Richardson)	Logistic Predictor
Method	Simulation	Pattern Recognition
Requires training data	No	Yes
Computation Time	High	Lower (after training)
State Universality	Yes	Best with states similar to training
Probabilities	Yes (sampling)	Yes (directly)

Figure 3. Richardson and logistic perform the same task radically differently.

Algorithm Descriptions

Search Bot (Richardson)

Plays out a state to the end (100x)

Averages W/L and score difference outcomes

Supervised Predictor (Logistic, Linear)

Olympic logs are labelled with W/L and score differences

36 features are computed:
• The immediate score difference
• Center draw area
• # and x,y of guards
• ...

Predictions are combined for all states

Trains on 9/10 of the labelled data, tests the rest

For all data

Prediction Results

Comparison of End Outcome Predictors Predicting 2010/2014 Olympics

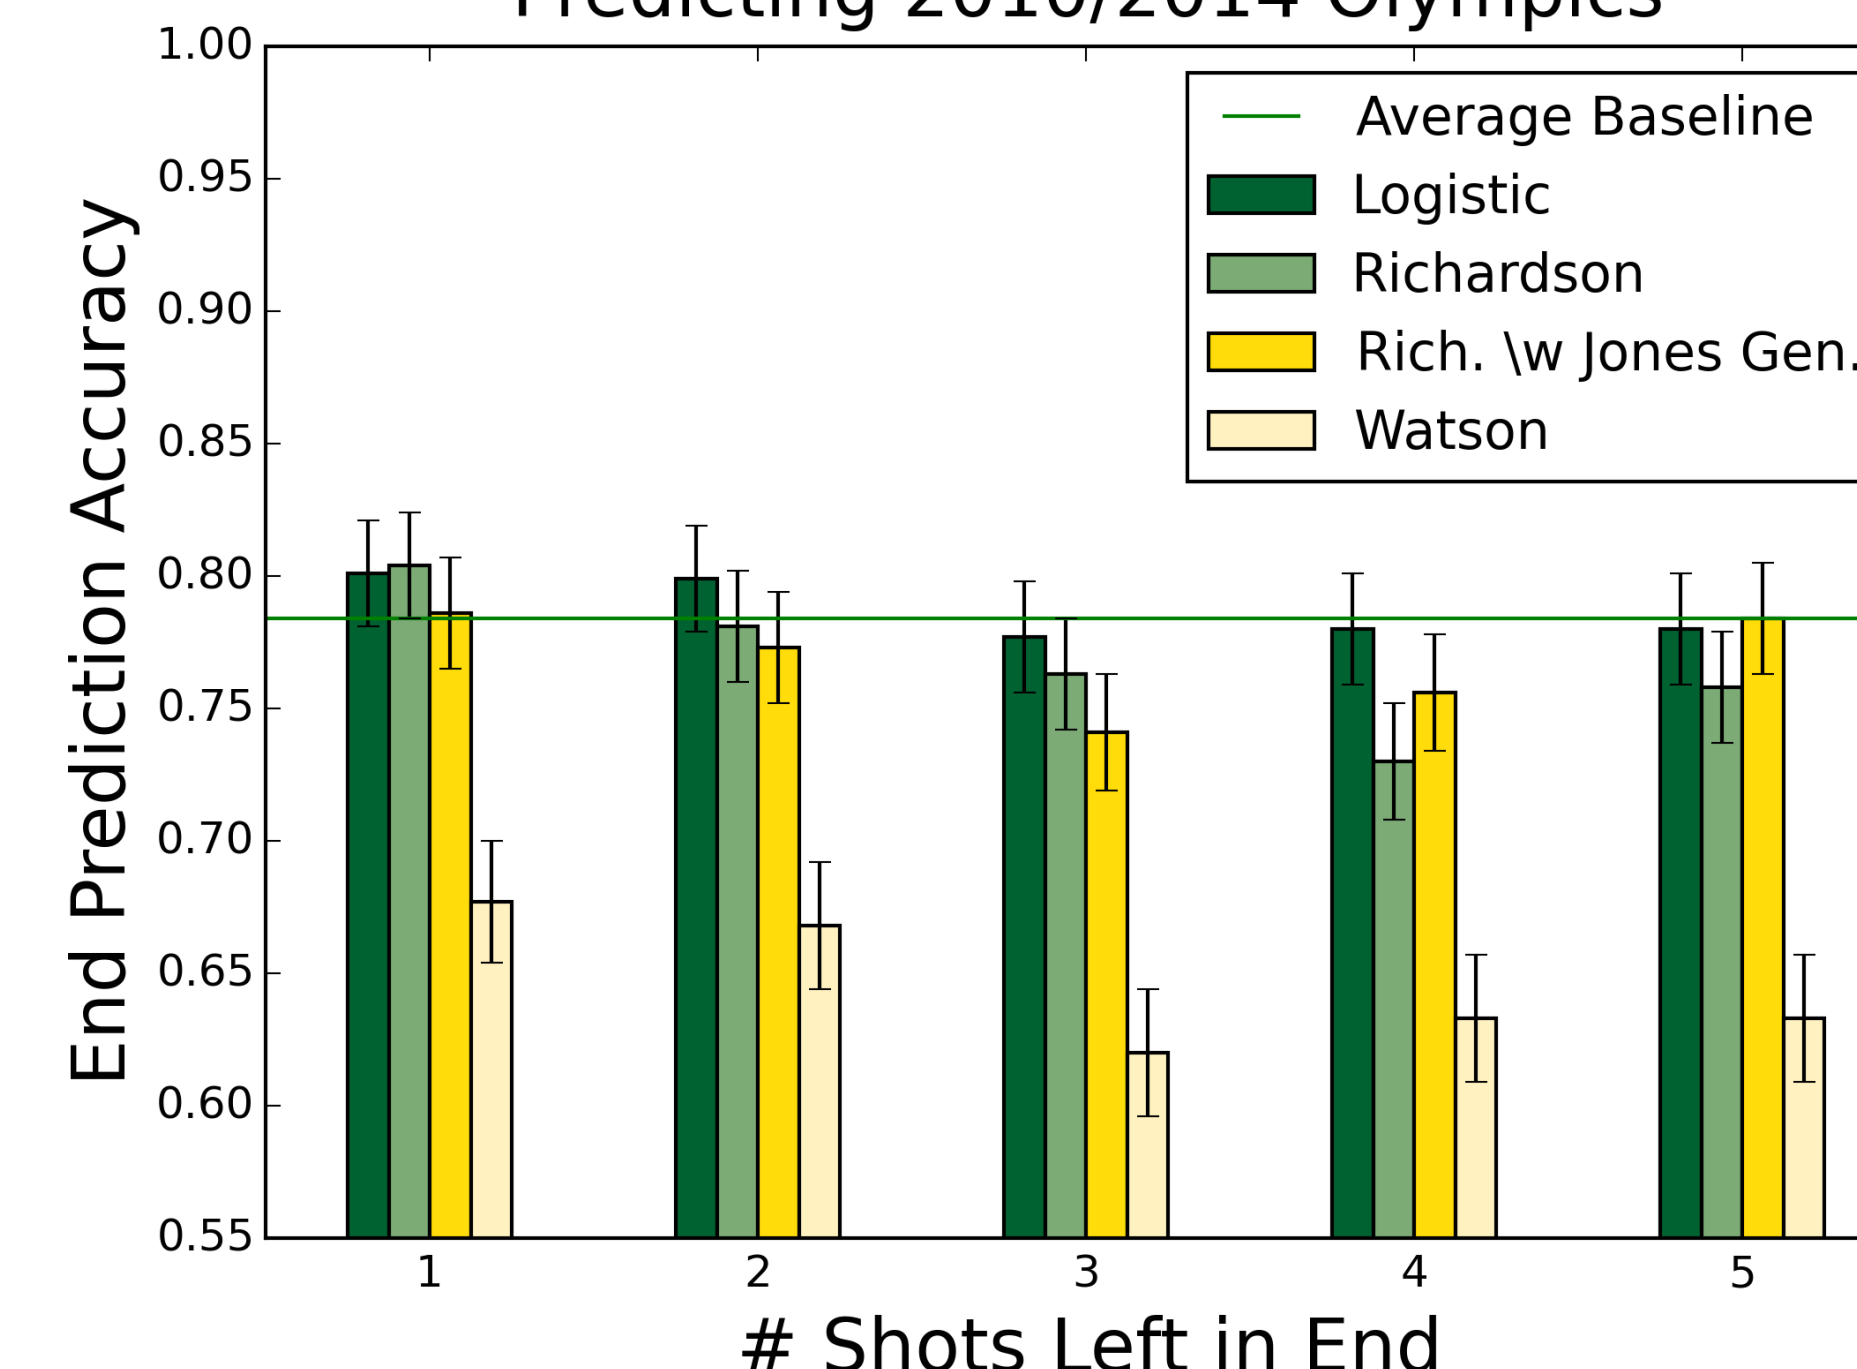


Figure 4. Predicting binary end win/loss

Comparison of End Outcome Predictors Predicting 2010/2014 Olympics Games




Figure 5. Linear regression always outperforms logistic at the score prediction task.

Comparison of End Outcome Predictors Predicting Richardson Games




Figure 6. Richardson can predict his own games better, at least for the last two shots.

State Evaluation Results

Logistic regression predicts the probability of each end score difference. This can be used directly (fig. 7) or converted to a win percentage (fig. 8).

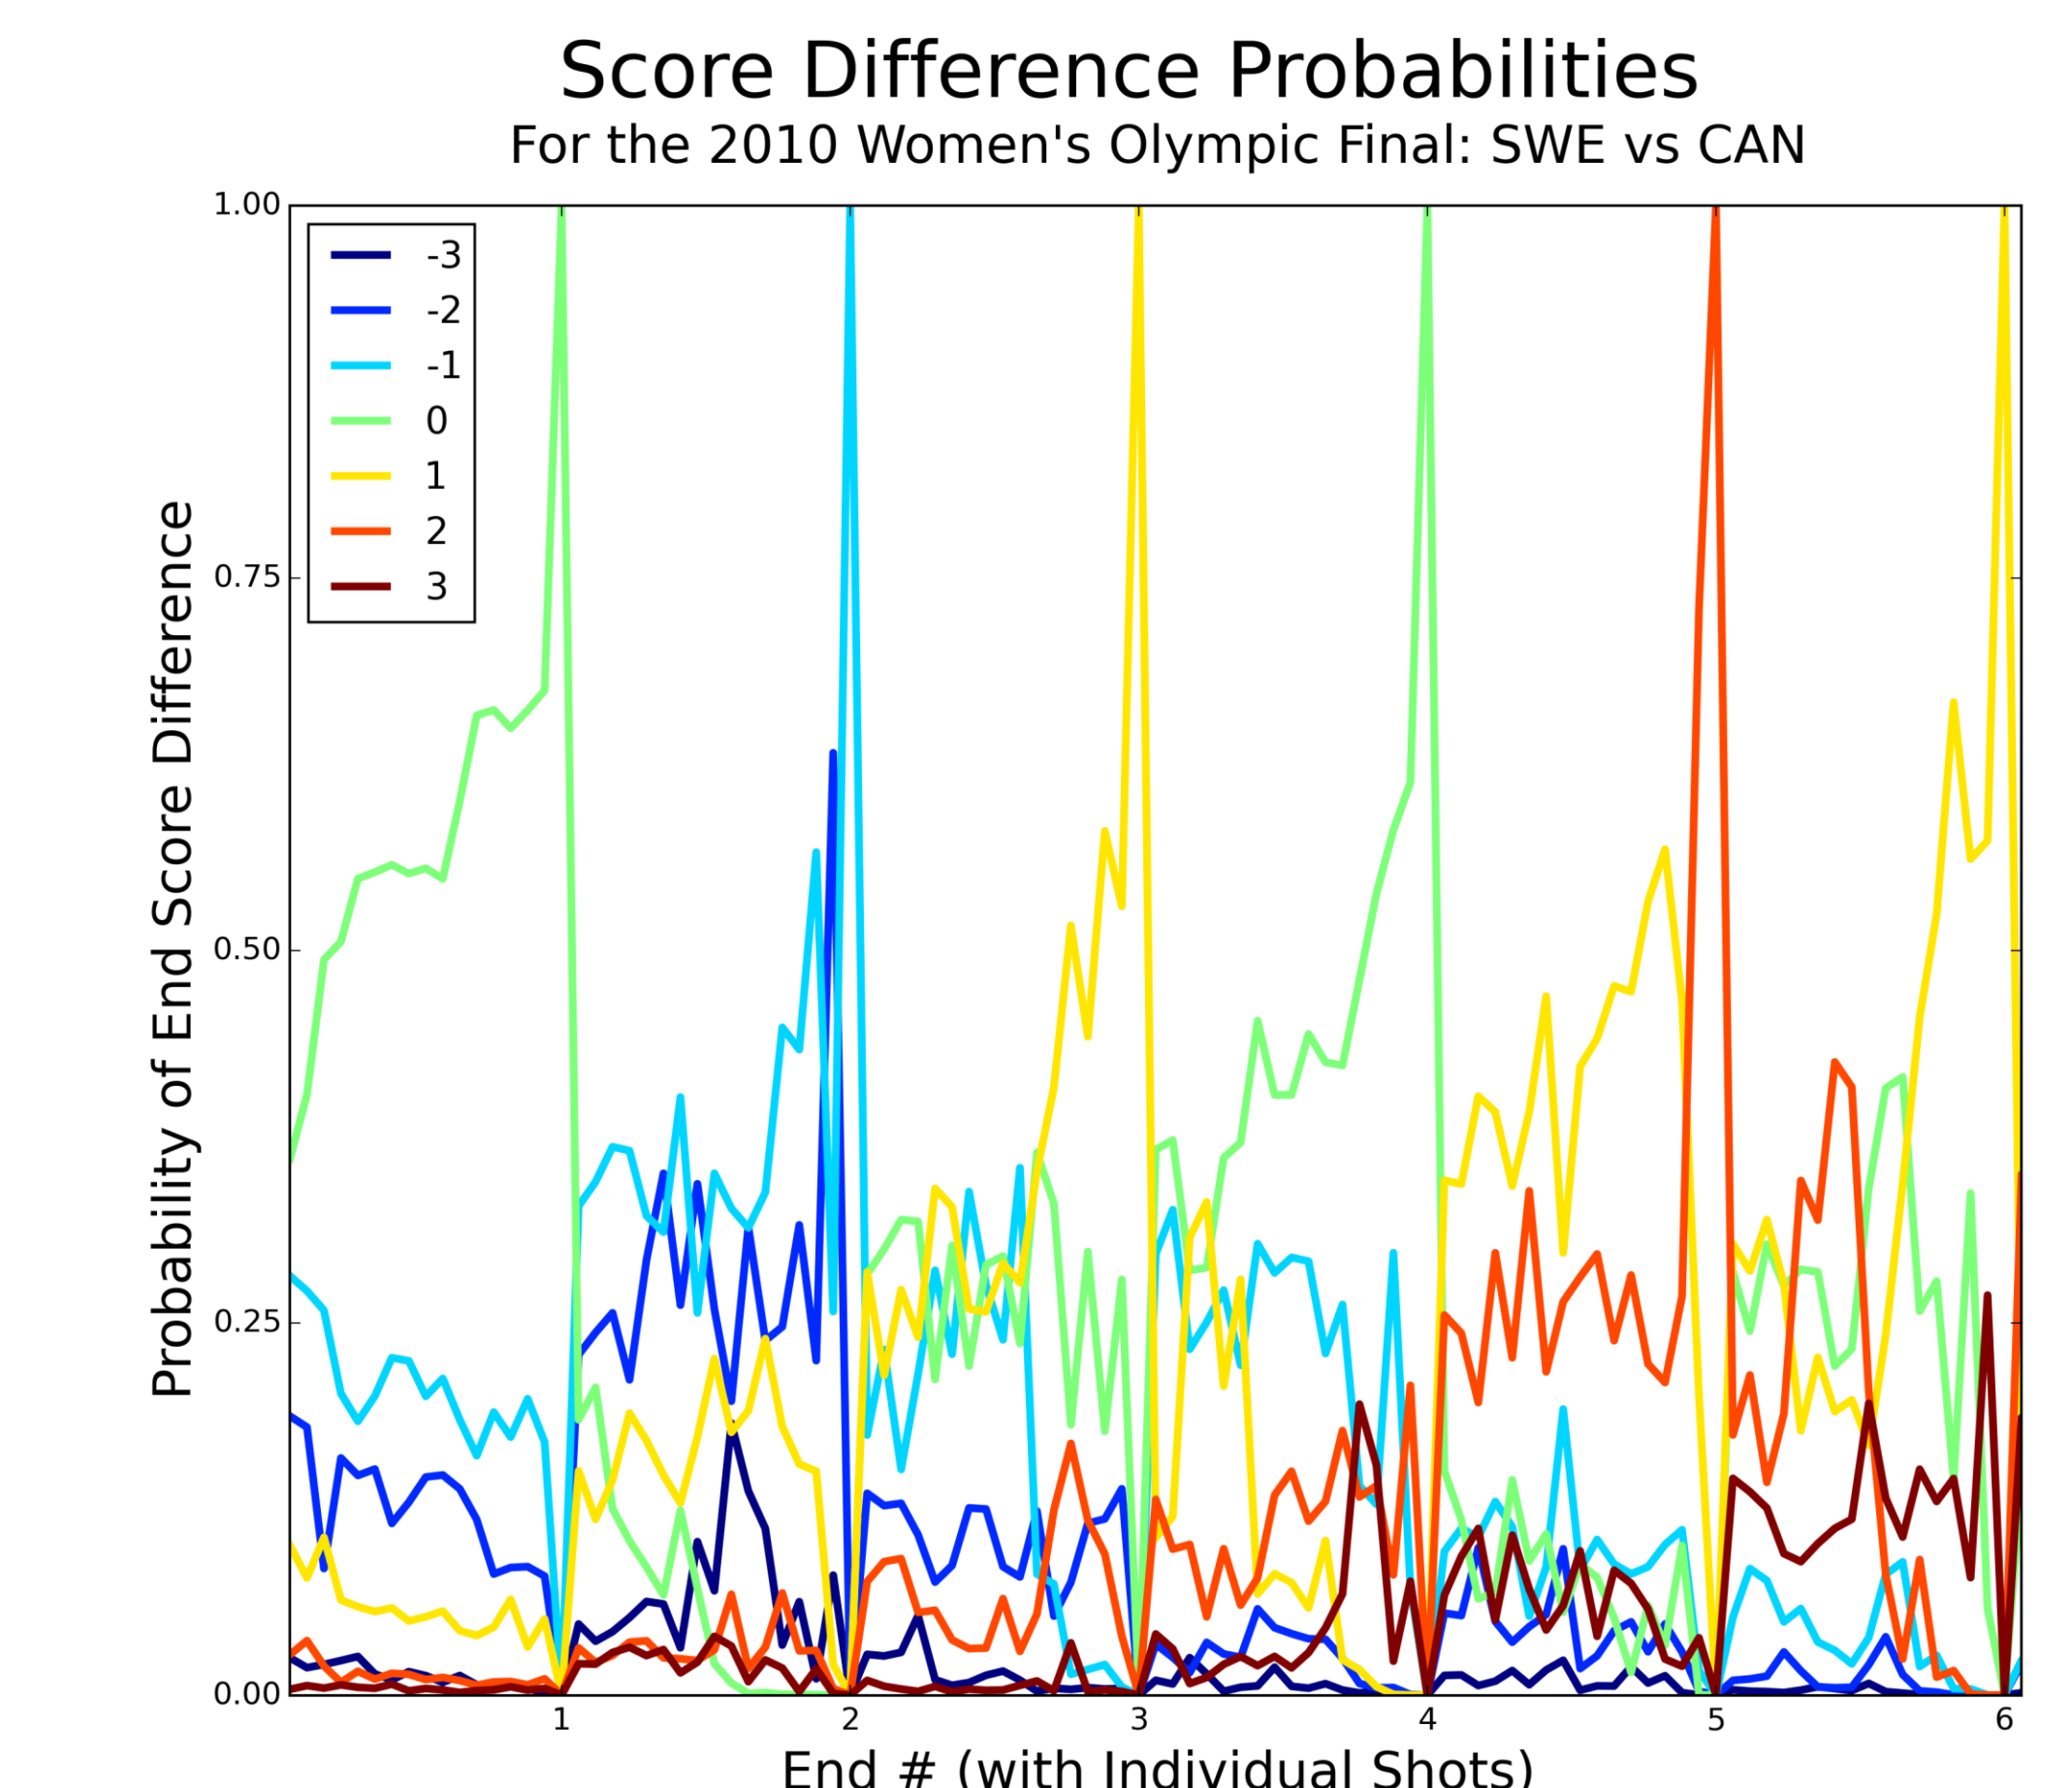


Figure 7. The evolution of score difference over time for an individual game.



Figure 8. Logistic state evaluation yields a significant ($\alpha=.05$) 7% gain over rollout evaluation!

Conclusion

Linear and Logistic regressions do not notably over or under perform compared to Richardson in the prediction task. On the other hand, logistic improves state evaluation.

Acknowledgments

Thank you to Micheal Bowling for supervising the supervised learning project, NSERC CRSNG for their support of undergraduate research, and the curling group for their help!