Because It's the Cup: Predicting the Stanley Cup Playoffs

Mason Swofford, Shuvam Chakraborty, Vineet Kosaraju

Backround

The Stanley Cup playoffs have long been known for their drama: most games are close, upsets are common, and teams not considered one of the best can win the cup. However, current predictions are based on traditional NHL statistics, which are not indicative of success due to luck influenced outcomes. Our project has two main goals: **1)** predict regular season and playoff game results, and **2)** construct a gambling agent to optimize returns.

Data Collection

1. Playoff Results **2. Daily Team Stats** **3. Betting Odds**

Two main datasets:

- Regular season & playoff games
- Training set amalgamated from data sources #2 (stats: features) and #3 (game results: labels)
- Each training example is 1 game, label = winning team (0=A, 1=H)
- Features for each training example consists of team statistics averaged over past games in season (see right)

Baseline Predictions

Team Type	Accuracy
Both Good Teams	0.524
Both Bad Teams	0.547
One Good Team, One Bad	0.572

Features

"Basic" Features

Shot attempt/quality features.

Feature	Description
CF	Corsi For, shot attempts for a team, including blocked shots, and shots not on goal
CA	Corsi Against, shot attempts against a team, including blocked and not on goal shots.
GF	Goals For
GA	Goals Against
xGF	Expected Goals For, based on quality of shot attempts for
xGA	Expected Goals Against, based on quality of shot attempts against
PENT	Penalties Taken
PEND	Penalties Drawn

"Advanced" Features

Includes shot attempt features, and adds overall team-based metrics.

Feature	Description
PDO	Shooting%+Save% (rough measure of luck)
FF	Fenwick For (unblocked shot attempts)
FA	Fenwick Against
SF	Shots For (on goal)
SA	Shots Against
xPDO	Expected PDO
dPDO	PDO difference
OZS	Offensive Zone Starts
DZS	Defensive Zone Starts
NZS	Neutral Zone Starts
ZSR	Zone Start Ratio
FOW	Faceoffs Won
FOL	Faceoffs Lost
GVA	Giveaways
TKA	Takeaways
HF	Hits For
HA	Hits Against
% Win	Winning Percentage

Models

Regarding **Goal 1**, classification models attempted include:

- Logistic, softmax regression
- SVM (rbf, linear, poly, sigmoid)
- ANNs (varying hidden layers, activation functions)

Features chosen using basic feature selection and PCA. Predicting if team A wins a playoff is done with a binomial distribution, where p is the prob. A wins a game:

$$= \sum_{n=4}^{7} \binom{n-1}{3} p^4 (1-p)^{(n-4)}$$

Regarding **Goal 2**, the gambling problem was formulated as a Markov Decision Process.

State: (currentMoney, game).

Start state: (initialMoney, 0).

Action: (money, team). Can bet up to current money on *Home/Away;* betting amounts discretized.

T(s, a, s'): Probabilities of transitions are given by our ML model.

isEnd(s): If we run out of money, or we have reached the last game.

R(s, a, s'): 1 if we have reached an end state and have greater than or equal to Desired Amount and 0 otherwise

Discount: Set to 1.

User Parameters:

Payoff: A number greater than 1 that corresponds to how much you get back for each dollar bet

Bucket Size: Discretization size for betting.

Desired Amount: minimum money we want to finish with

Results & Error Analysis

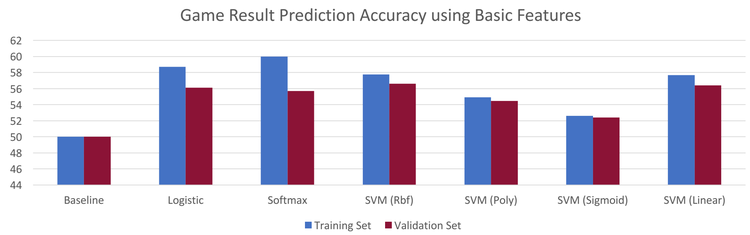

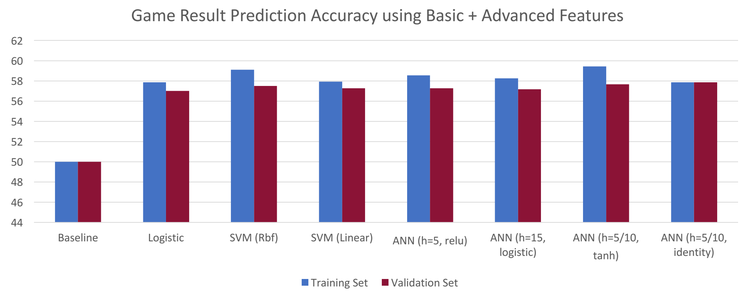

Figure 1, 2: Training and validation accuracies reported using 10-fold cross validation. For the best model, the accuracy on the test set of playoffs was 54.66% (for reference, ESPN experts were ~51% accurate).

Ablative analysis of basic features demonstrated that more advanced ones were needed, however even advanced features didn't help. The literature mentions a theoretical limit for predicting the result of a single hockey game due to luck/variability:

This limit of 60-63% was confirmed in a Monte Carlo simulation, running 1000 trials, suggesting games can't be directly predicted (right). This model is similar to those used in the NFL.

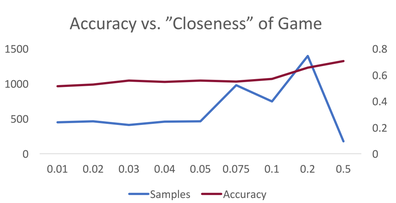

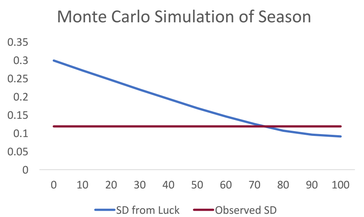

Figure 3: Results are accounted for by 73% luck, so when making predictions we can accurately predict 27% and guess with 50% accuracy on the 73%, which gives us a 63.5% ceiling.

Team Type	Number Games	Accuracy
Both Good	1088	0.5588
Both Bad	1306	0.5628
Good & Bad	3146	0.5950

MDP: Example

Total: $1475 in 3 games.

Conclusions

- Hockey is a very challenging sport to predict due to variability inherent to sport.
- Perfect stats could allow reaching the theoretical accuracy limit, but incremental progress needed.
- Reached 70% using SVM on playoff data, so model could be fine-tuned.
- Applications in other leagues (other than NHL) or sports (baseball, etc).

References

Pischedda, Gianni. *Predicting NHL Match Outcomes with ML Models*. citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.735.795&rep=rep1&type=pdf.

Weissbock, Joshua, et al. *Use of Performance Metrics to Forecast Success in the National Hockey League*. ceur-ws.org/Vol-1969/paper-06.pdf.