Interpreting Deep Sports Analytics: Valuing Actions and Players in the NHL

Guiliang Liu, Wang Zhu, Oliver Schulte
Machine Learning Lab

SIMON FRASER UNIVERSITY
ENGAGING THE WORLD

ECML-PKDD 2018 Presentation

Problem

Understand the Deep Reinforcement Learning (DRL) Model in National Hockey League (NHL)

DRL Model: Previous Work

Liu and Schulte IJCAI 2018

Dataset

- *Game events* and *player actions* for the 2015-2016 NHL season.
- Augment the data with derived features (red lines).

Table 3: Complete Feature List. Values for the feature Manpower are EV=Even Strength, SH=Short Handed, PP=Power Play.

Name	Type	Range
X Coordinate of Puck	Continuous	[-100, 100]
Y Coordinate of Puck	Continuous	[-42.5, 42.5]
Velocity of Puck	Continuous	(-inf, +inf)
Time Remaining	Continuous	[0, 3600]
Score Differential	Discrete	(-inf, +inf)
Manpower	Discrete	{EV, SH, PP}
Event Duration	Continuous	[0, +inf)
Action Outcome	Discrete	{successful, failure}
Angle between puck and goal	Continuous	[−3.14, 3.14]
Home/Away Team	Discrete	{Home, Away}

- Divide NHL games into **goal-scoring episodes** that
  - Begin at the beginning of the game, or after a goal.
  - Terminate with a goal, or the end of the game.

DRL Model

- Estimate chance that team scores the next goal given current match state and action = $Q_{team}(s,a)$.
- Recurrent network with dynamic trace length LSTM.

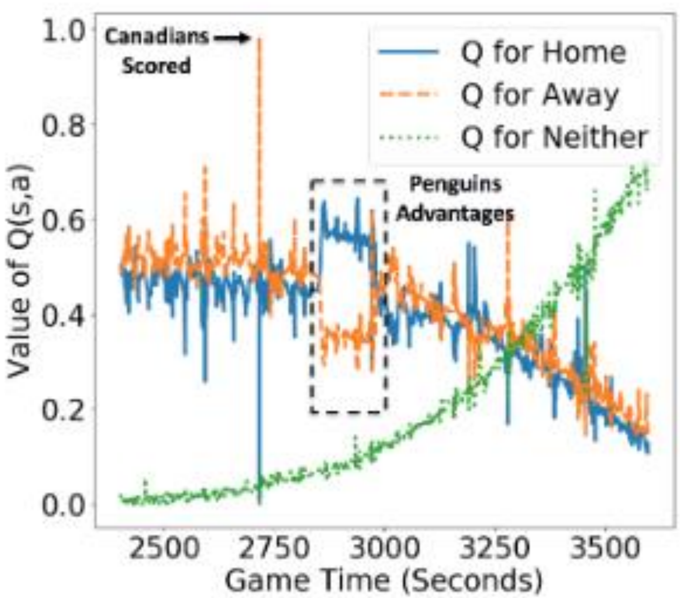

Temporal Projection: evolution of scoring probabilities for the next goal

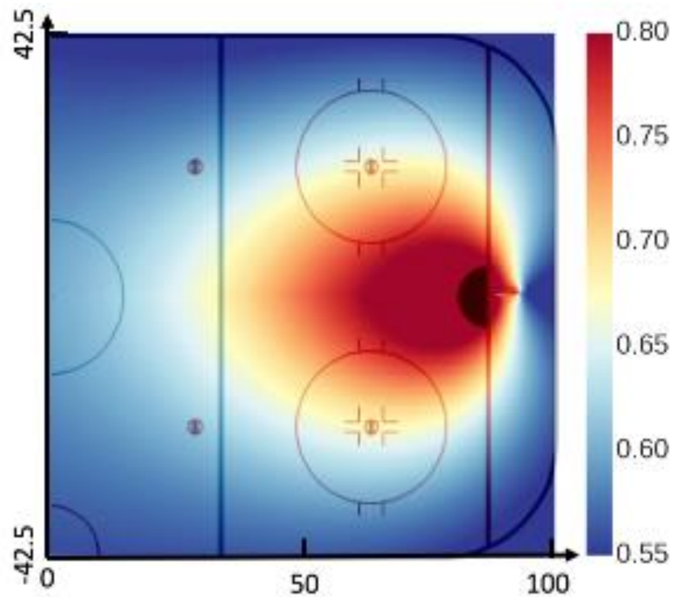

Spatial Projection (for shot): The probability that the home team scores the next goal after taking a shot at a rink location

Goal Impact Metric

- $\boldsymbol{Impact}(\boldsymbol{s_t}, \boldsymbol{a_t})$ measures the quality of action a_t by how much it changes the expected return of a player's team.

$$impact^{team}(s_t, a_t) = \underbrace{Q^{team}(s_t, a_t) - Q^{team}(s_{t-1}, a_{t-1})}_{\text{Difference of consecutive Q values}}$$

- Define **Goal Impact Metric (GIM)** of player i by the total impact of a player in entire game season dataset D.

$$GIM^i(D) = \sum_{s,a} n_D^i(s,a) \times impact^{team_i}(s,a)$$

Goal Impact Metric

- The **Impact** metric passes "eye test".
- Correlates strongly with goals, points, salary, etc. in NHL.
- Consistent between and within seasons.
- All actions including **defensive** and **offensive** actions.

Table 4: 2015-2016 Top-10 Player Impact Scores

Name	GIM	Assists	Goals	Points	+/-	Age	Team	Salary
Taylor Hall	96.40	39	26	65	-4	24	EDM	$6,000,000
Joe Pavelski	94.56	40	38	78	25	31	SJS	$6,000,000
Johnny Gaudreau	94.51	48	30	78	4	22	CGK	$925,000
Anze Kopitar	94.10	49	25	74	34	28	LAK	$7,700,000
Erik Karlsson	92.41	66	16	82	-2	25	OTT	$7,000,000
Patrice Bergeron	92.06	36	32	68	12	30	BOS	$8,750,000
Mark Scheifele	90.67	32	29	61	16	23	WPG	$832,500
Sidney Crosby	90.21	49	36	85	19	28	PIT	$12,000,000
Claude Giroux	89.64	45	22	67	-8	28	PHI	$9,000,000
Dustin Byfuglien	89.46	34	19	53	4	31	WPG	$6,000,000

Interpreting the DRL Model

Model

Mimic Learning Framework for **General Model**:

- Mimicking Q functions and impact separately.
- History Window of last 10 observations.
- A Multi-variate Regression Tree (MRT) trained with CART method.

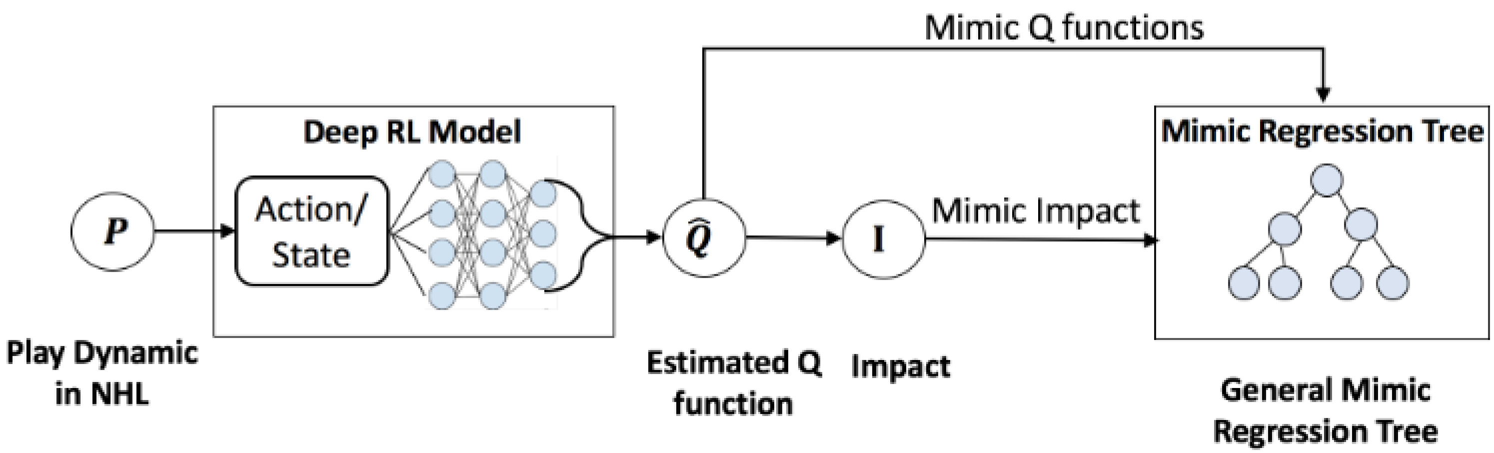

Fig. 3: Interpretable Mimic Learning Framework

Model

Player Specific Model:

- Inherit the tree structure of the *general model*.
- Use the target player data to prune the *general model.*
- Use the same player data to expand the tree.

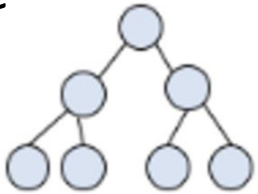

General tree

initialize

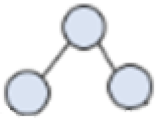

prune

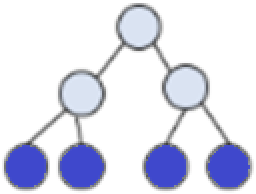

expand

Player tree
e.g. Sidney Croby

Model

Mean Sample Leaf (MSL):

- Control the minimum number of samples at each leaf node.
- Satisfactory performance when **MSL = 20**.

Table 5: Performance of General Mimic Regression Tree (RT) with different Minimum Samples in each Leaf node (MSL). We apply ten-fold cross validation and report the regression result with format: Mean Square Error (Variance)

model	Q_home	Q_away	Q_end	Impact
RT-MSL1	3.35E-04 (1.43E-09)	3.21E-04(1.26E-09)	1.74E-04(2.18E-09)	1.33E-03(5.43E-09)
RT-MSL5	2.59E-04(1.07E-09)	2.51E-04(0.89E-09)	1.35E-04(1.87E-10)	9.84E-04(2.72E-09)
RT-MSL10	2.38E-04(1.02E-09)	2.30E-04(0.89E-09)	1.25E-04(2.30E-10)	8.66E-04(2.17E-09)
RT-MSL20	**2.31E-04(0.92E-09)**	**2.22E-04(0.82E-09)**	**1.23E-04(2.05E-10)**	7.92E-04(1.45E-09)
RT-MSL30	2.35E-04(0.98E-09)	2.27E-04(0.85E-09)	1.27E-04(2.32E-10)	7.67E-04(1.16E-09)
RT-MSL40	2.39E-04(0.96E-09)	2.30E-04(0.85E-09)	1.29E-04(2.19E-10)	**7.58E-04(1.10E-09)**

Feature Importance

Rank feature by average variance reduction:

- Find the top 10 important features using *general model*.
- The impact function recognizes **shooting, successful** actions.
- History Window is necessary.

Table 6: Top 10 features for Q values (left) and Impact (right). The notation $T - n : f$ indicates that a feature occurs n time steps before the current time.

Feature Name	Frequency	Importance
T: Time Remaining	12,524	0.817431
T-1: Manpower	93	0.070196
T-1: Team Identifier	57	0.020504
T: Manpower	346	0.017306
T: Shot	31	0.011159
T: Score Differential	3,229	0.009568
T: X Coordinate	11,797	0.006968
T-1: X Coordinate	3,406	0.006963
T-2: Manpower	82	0.005045
T: Home/Away Team	135	0.003755

Feature Name	Frequency	Importance
T: Goal	1	0.160595
T: Shot-on-Goal	1	0.099482
T: X Coordinate	7,142	0.077410
T: X Velocity	8,087	0.041903
T-1: X Coordinate	3,591	0.041847
T: Angle to Goal	7,525	0.041607
T: Time Remaining	8,669	0.036289
T: Duration	7,411	0.028831
T: Home/Away Team	378	0.027177
T: Y Coordinate	6,890	0.027597

Partial Dependence

Partial Dependence plot:

- Use *general model* to interpret Q functions and impact.
- Select Time Remaining, X Coordinate and X Velocity to visualize.

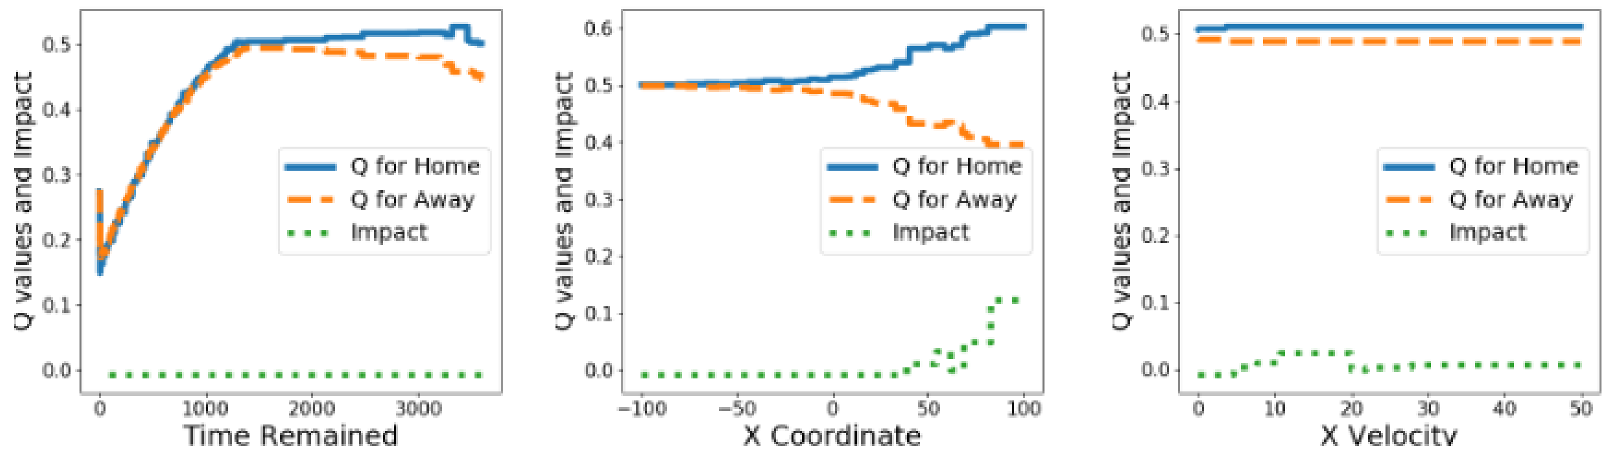

Fig. 4: Partial Dependence Plot for Time Remaining (left), X Coordinate (middle) and X Velocity (right)

Exceptional Players

Find the most unusual players:

- Use *player specific model* to compare the whole dataset (general data) and player specific data.

Table 7: Exceptional Players Based on Tree Discretization

Player	Q_home	Q_away	Q_end	Impact
Taylor Hall	1.80E-04	2.49E-04	2.28E-04	6.66E-05
Joe Pavelski	**4.64E-04**	**2.90E-04**	**3.04E-04**	1.09E-04
Johnny Gaudreau	2.12E-04	1.96E-04	1.43E-04	6.77E-05
Anze Kopitar	2.58E-04	2.00E-04	2.43E-04	8.28E-05
Erik Karlsson	2.97E-04	1.89E-04	1.86E-04	**2.00E-04**

- Joe Pavelski scored the most in the 2015-2016 game season.
- Erik Karlsson had the most points (goal+assists).

Thank You!

For more information:

Poster: #xxx

Homepage: http://www.galenliu.com/

Exceptional Players

How to find the most unusual player:

- Focus on top players
- For each *player specific model*:
  - For each leaf in the tree
    - There is an original value (e.g. I_l^P)
    - Learn a value based on the whole dataset (I_l)
    - Weight by the percentage of cases that get to the leaf (n_l/n_D)
    - Sum over squared differences $n_l/n_D\left(I_l^P - I_l\right)^2$

Table 7: Exceptional Players Based on Tree Discretization

Player	Q_home	Q_away	Q_end	Impact
Taylor Hall	1.80E-04	2.49E-04	2.28E-04	6.66E-05
Joe Pavelski	**4.64E-04**	**2.90E-04**	**3.04E-04**	1.09E-04
Johnny Gaudreau	2.12E-04	1.96E-04	1.43E-04	6.77E-05
Anze Kopitar	2.58E-04	2.00E-04	2.43E-04	8.28E-05
Erik Karlsson	2.97E-04	1.89E-04	1.86E-04	**2.00E-04**