Deep Reinforcement Learning in Ice Hockey for Context-Aware Player Evaluation

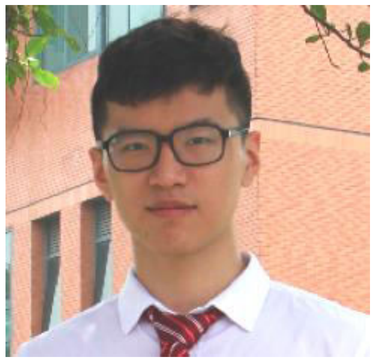

Guiliang Liu

Oliver Schulte

SFU
SIMON FRASER UNIVERSITY

SPORTLOGiQ

IJCAI ECAI 2018

PROBLEM

Evaluate players in the largest ice hockey league: National Hockey League (NHL)

RELATED WORK

Action-Wise Players Evaluation

Year	Venue	Authors	Name	Sports
2018	Arxiv	Tom Decroos, Lotte Bransen, et al.	Actions speak louder than goals: Valuing player actions in soccer.	Soccer
2017	MIT Sloan	Oliver Schulte, Zeyu Zhao, et al.	Apples-to-apples: Clustering and ranking NHL players using location…	Ice Hockey
2015	UAI	Kurt Routley and Oliver Schulte.	A Markov game model for valuing player actions in ice hockey.	Ice Hockey
2014	MIT Sloan	Dan Cervone , Alexander, et al.	Pointwise: Predicting points and valuing decisions in real time …	Basket ball

MOTIVATION: DATASET

- Computer Vision Techniques: Video tracking

- Play-by-play Dataset

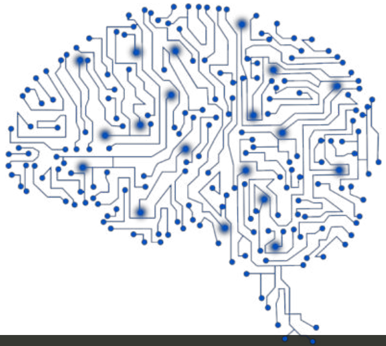

- Large-scale Machine Learning

MOTIVATION: GAME COMPLEXITY

- Model complex game *context*
- Partial Observability

OVERVIEW OF METHOD

- Framework of Deep Reinforcement Learning (DRL) model

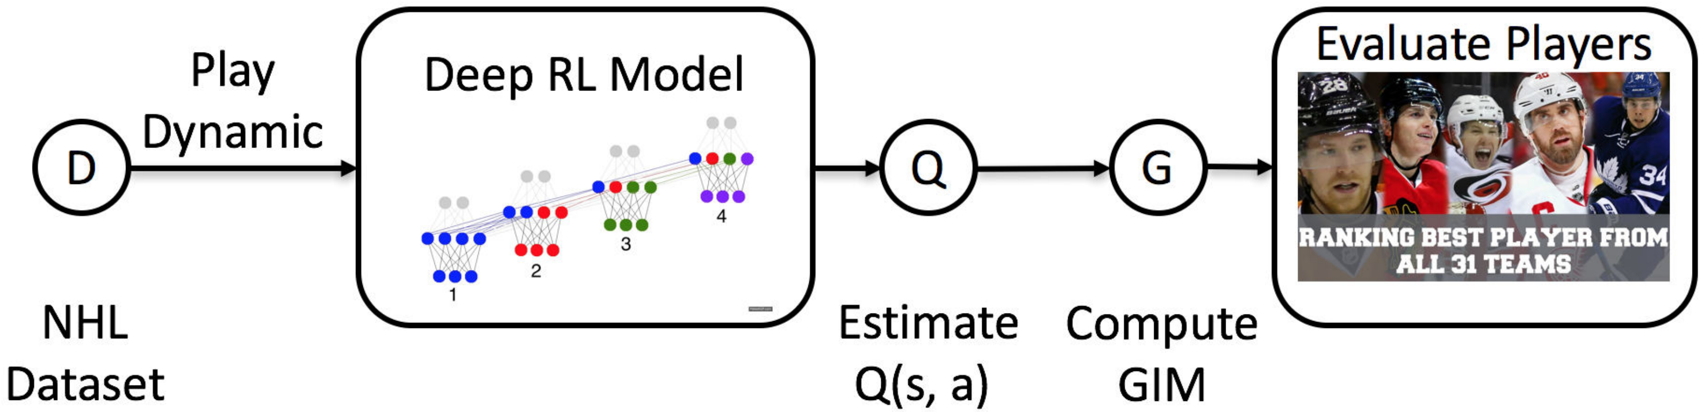

1) Extract play dynamic from NHL dataset.
2) Estimate the Q(s, a) with DRL model.
3) Define a novel Goal Impact Metric (GIM) to value each player.

PLAY DYNAMICS

Play-by-play NHL Dataset

- Contain 3M events.
- Cover 30 teams, 1,140 games and 2,233 players.

PLAY DYNAMICS

Play-by-play NHL Dataset

- Contain 3M events.
- Cover 30 teams, 1,140 games and 2,233 players.

Constructing a Reinforcement Learning Environment

- Action a_t: players' action, including shot, block, assist, etc.
- State s_t: sequence of observations and actions (Q_t, a_{t-1}, Q_{t-1} ...)
- Reward r_t: a one-hot goal vector specifies which team scores.
- Q function Q_{team}: the probability of scoring the next goal:

$$Q^{team}(s, a) = P(goal^{team} = 1 | s_t = s, a_t = a)$$

- Recurrent network with dynamic trace length LSTM

Model Structure

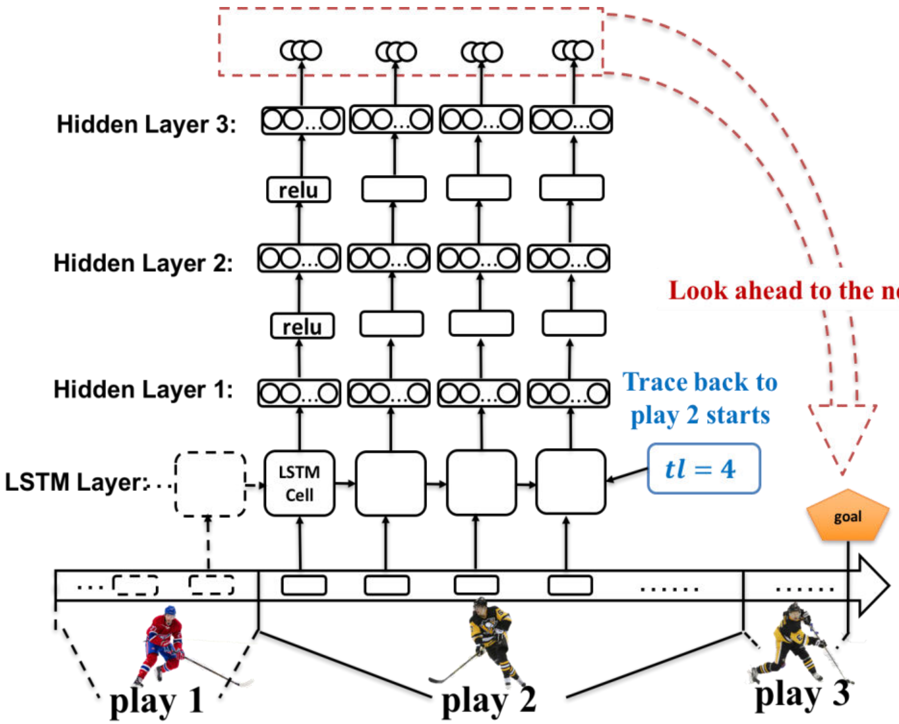

Value Ticker: Temporal Projection of learned Q functions（NN output）

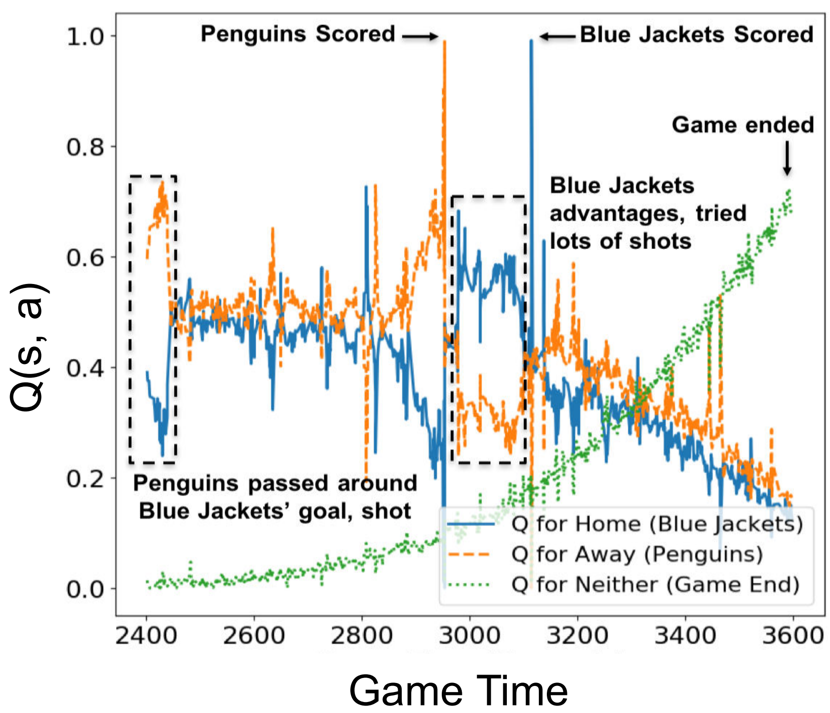

GOAL IMPACT METRIC

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

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

GOAL IMPACT METRIC

- $\boldsymbol{Impact(s_t, a_t)}$ 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.

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

PLAYER RANKING

Rank players by GIM and identify undervalued players

Name	GIM	Assists	Goals	Points	Team	Salary
Taylor Hall	96.40	39	26	65	EDM	$6,000,000
Joe Pavelski	94.56	40	38	78	SJS	$6,000,000
Johnny Gaudreau	94.51	48	30	78	CGY	$925,000
Anze Kopitar	94.10	49	25	74	LAK	$7,700,000
Erik Karlsson	92.41	66	16	82	OTT	$7,000,000
Patrice Bergeron	92.06	36	32	68	BOS	$8,750,000
Mark Scheifele	90.67	32	29	61	WPG	$832,500
Sidney Crosby	90.21	49	36	85	PIT	$12,000,000
Claude Giroux	89.64	45	22	67	PHI	$9,000,000
Dustin Byfuglien	89.46	34	19	53	WPG	$6,000,000
Jamie Benn	88.38	48	41	89	DAL	$5,750,000
Patrick Kane	87.81	60	46	106	CHI	$13,800,000
Mark Stone	86.42	38	23	61	OTT	$2,250,000
Blake Wheeler	85.83	52	26	78	WPG	$5,800,000
Tyler Toffoli	83.25	27	31	58	DAL	$2,600,000
Charlie Coyle	81.50	21	21	42	MIN	$1,900,000
Tyson Barrie	81.46	36	13	49	COL	$3,200,000
Jonathan Toews	80.92	30	28	58	CHI	$13,800,000
Sean Monahan	80.92	36	27	63	CGY	$925,000
Vladimir Tarasenko	80.68	34	40	74	STL	$8,000,000

- Mark Scheifele drew salaries **below** what his GIM rank would suggest.
- Later he received a $5M+ contract in 2016-17 season

EMPIRICAL EVALUATION

Comparison Metric:

- Plus-Minus (+/-)
- Goal-Above-Replacement (GAR)
- Win-Above-Replacement (WAR)
- Expected Goal (EG)
- Scoring Impact (SI)
- GIM-T1

EMPIRICAL EVALUATION

Comparison Metric:

- Plus-Minus (+/-)
- Goal-Above-Replacement (GAR)
- Win-Above-Replacement (WAR)
- Expected Goal (EG)
- Scoring Impact (SI)
- GIM-T1

Correlations with standard Success Measures:

- Compute the correlation with 14 standard success measures:

methods	Point	SHP	PPP	FOW	P/GP	TOI	PIM
+/-	0.237	0.159	0.089	-0.045	0.238	0.141	0.049
GAR	0.622	0.226	0.532	0.16	0.616	0.323	0.089
WAR	0.612	0.235	0.531	0.153	0.605	0.331	0.078
EG	0.854	0.287	0.729	0.28	0.702	0.722	0.354
SI	0.869	0.37	0.707	0.185	0.655	**0.955**	**0.492**
GIM-T1	0.902	0.384	0.736	0.288	0.738	0.777	0.347
GIM	**0.93**	**0.399**	**0.774**	**0.295**	**0.749**	0.835	0.405

methods	Assist	Goal	GWG	OTG	SHG	PPG	S
+/-	0.236	0.204	0.217	0.16	0.095	0.099	0.118
GAR	0.527	0.633	0.552	0.324	0.191	0.583	0.549
WAR	0.516	0.652	0.551	0.332	0.192	0.564	0.532
EG	0.783	0.834	0.704	0.448	0.249	0.684	0.891
SI	0.869	0.745	0.631	0.411	0.27	0.591	0.898
GIM-T1	0.873	0.752	0.682	0.428	0.291	0.607	0.877
GIM	**0.875**	**0.878**	**0.751**	**0.465**	**0.345**	**0.71**	**0.912**

EMPIRICAL EVALUATION

Round-by-Round Correlations:

- How *quickly* a metric acquires predictive power for the season total.
- For a metric (EG, SI, GIM-T1, GIM), measure the *correlation* between
  a) Its value computed over the **first n round**.
  b) The value of the three main success measures, assists, goals, points and its value computed over the **entire season**.

EMPIRICAL EVALUATION

Round-by-Round Correlations:

- How *quickly* a metric acquires predictive power for the season total.
- For a metric (EG, SI, GIM-T1, GIM), measure the *correlation* between
  a) Its value computed over the **first n round**.
  b) The value of the three main success measures, assists, goals, points and its value computed over the **entire season**.

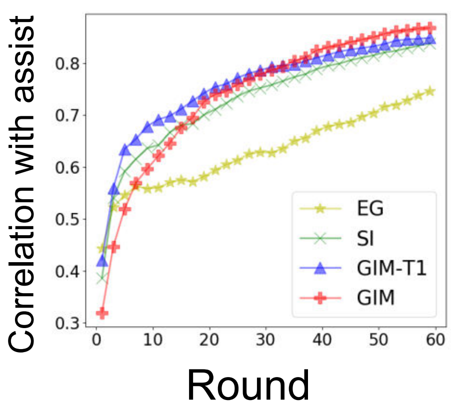

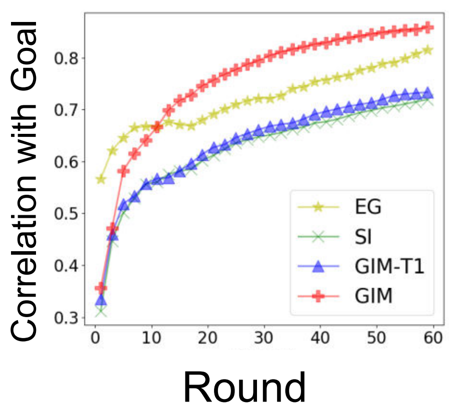

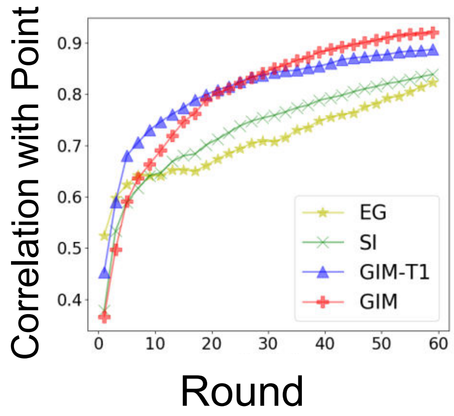

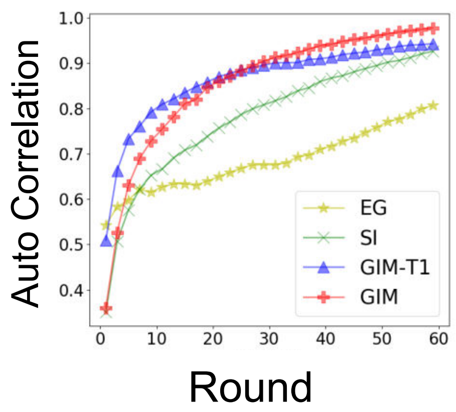

EMPIRICAL EVALUATION

Predicting Players' Salary:

- A good metric is positively related to players' future contract.

methods	2016 to 2017 Season	2017 to 2018 Season
Plus Minus	0.177	0.225
GAR	0.328	0.372
WAR	0.328	0.372
EG	0.587	0.6
SI	0.609	0.668
GIM-T1	0.596	0.69
GIM	**0.666**	**0.763**

EMPIRICAL EVALUATION

Predicting Players' Salary:

- A good metric is positively related to players' future contract.

methods	2016 to 2017 Season	2017 to 2018 Season
Plus Minus	0.177	0.225
GAR	0.328	0.372
WAR	0.328	0.372
EG	0.587	0.6
SI	0.609	0.668
GIM-T1	0.596	0.69
GIM	**0.666**	**0.763**

- Many underestimated players in 16-17 season. (high GIM, low salary).
- This percentage decreases in 17-18 season. (from 32/258 to 8/125).

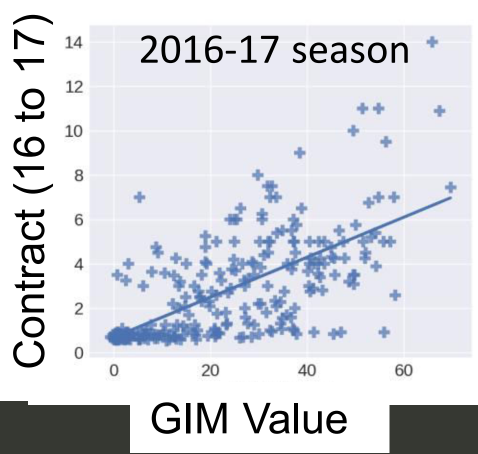

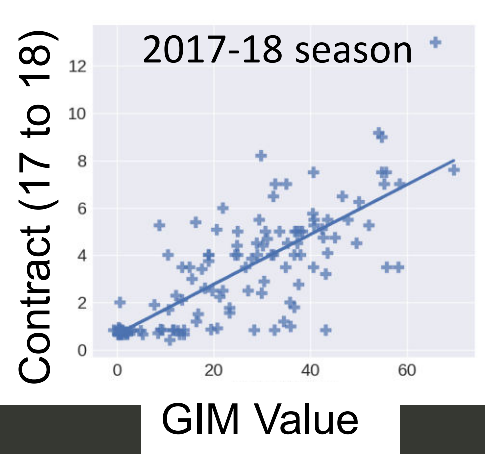

THANK YOU!

For more information:

Poster: #2177

Github link: https://github.com/Guiliang/DRL-ice-hockey

My homepage: http://www.galenliu.com/