What is the value of an action in ice hockey? Deep Reinforcement Learning for Context-Aware Player Evaluation

Oliver Schulte

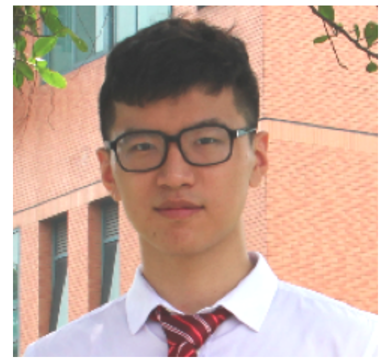

Guiliang Liu

SFU
SIMON FRASER UNIVERSITY

SPORTLOGiQ

Sport Analytics

Growth in Industry

- The Sports Analytics market is expected to grow from USD 123.7 Million in 2016 to USD 616.7 Million by 2021
- Commercial *data providers* include:
  - [Sportlogiq]
  - [Stats]

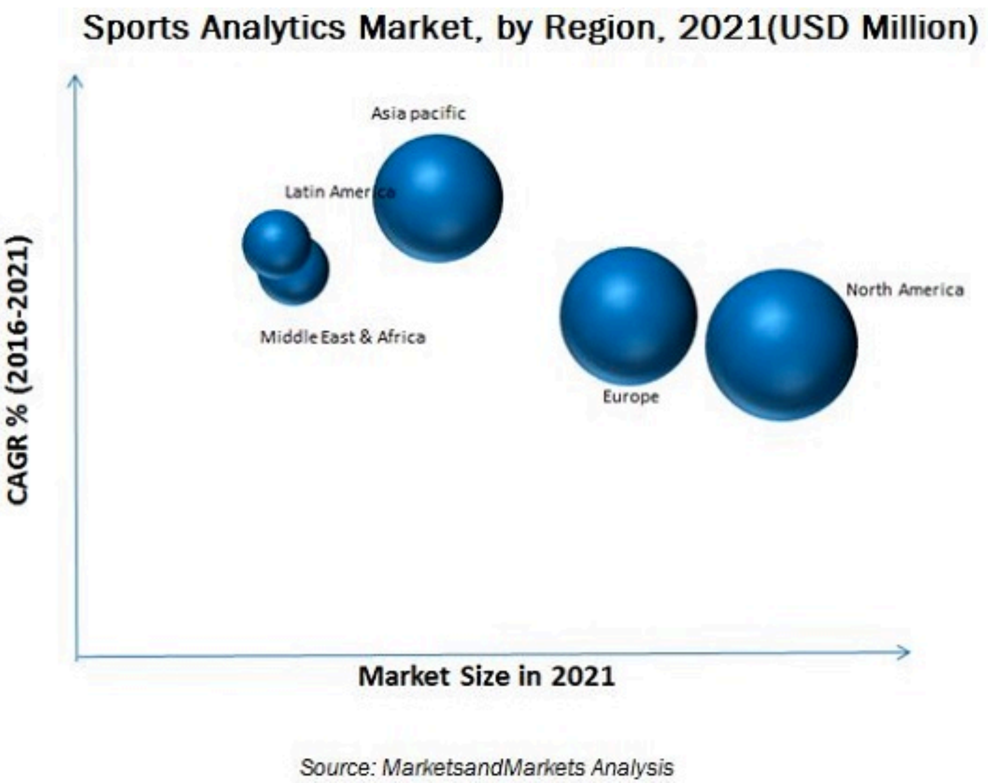

Sport Analytics

Growth in academia

- MIT Sloan Sport Analytics Conference (held every year in Boston since 2007). Research and application papers.
- Journals
  - Journal Quantitative Analysis of Sports
  - Journal of Sports Analytics .
- Sports Analytics Group in SFU.
- Sports Analytics B.Sc. at Syracuse university
- Contributions to AI-related conferences (AAAI, IJCAI, UAI, KDD) in the recent years.

AI Meets Sports Analytics

AI

- modelling and learning game strategies
- multi-agent systems
- structured data (space, time)
- decision support for coaches, players, teams
  - identifying strengths and weaknesses (“gap analysis”)
  - suggesting and identifying tactics

The Big Picture

Our Approach: Sports Analytics as a major application area for Reinforcement Learning

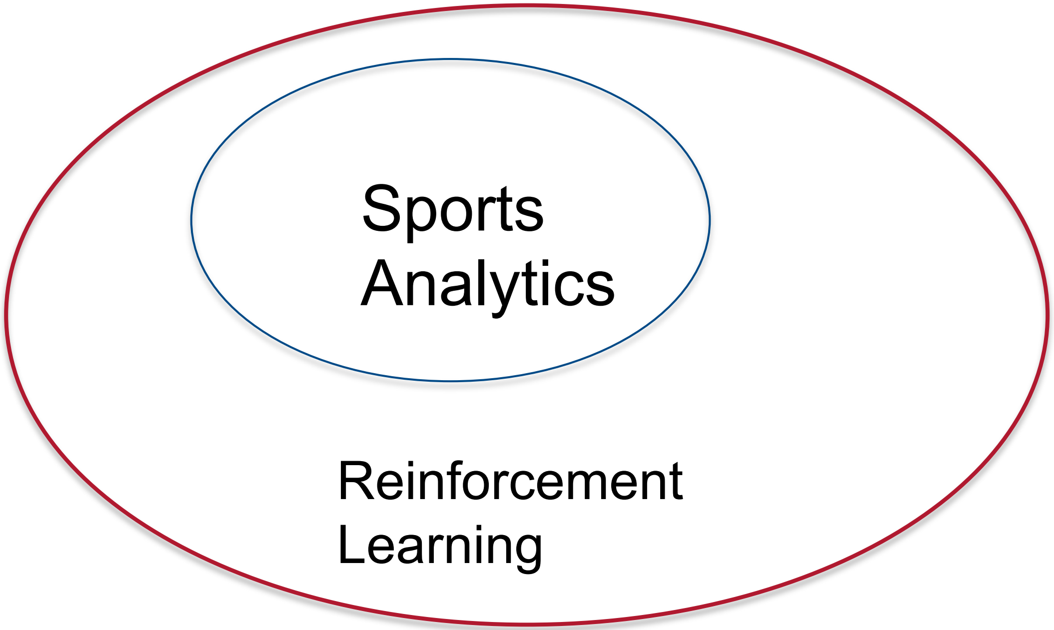

Sports Analytics

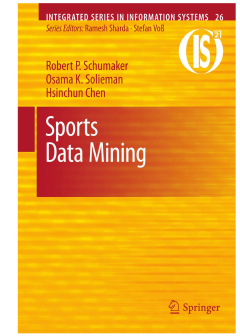

Application

- Tracking
- Evaluation
- Prediction

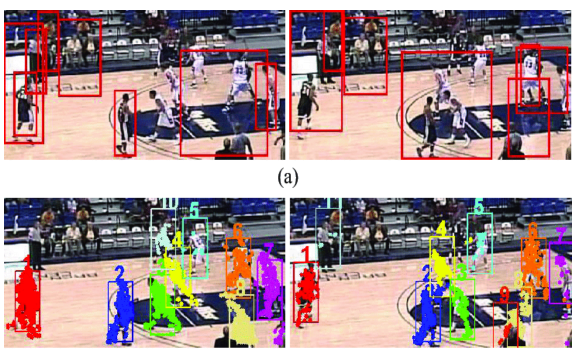

(a)

(b)

Performance Evaluation: A Reinforcement Learning Approach

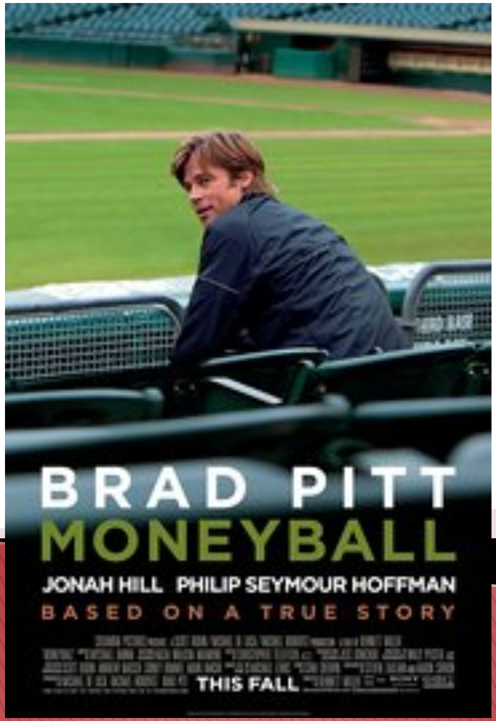

PROBLEM

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

Previous Approaches

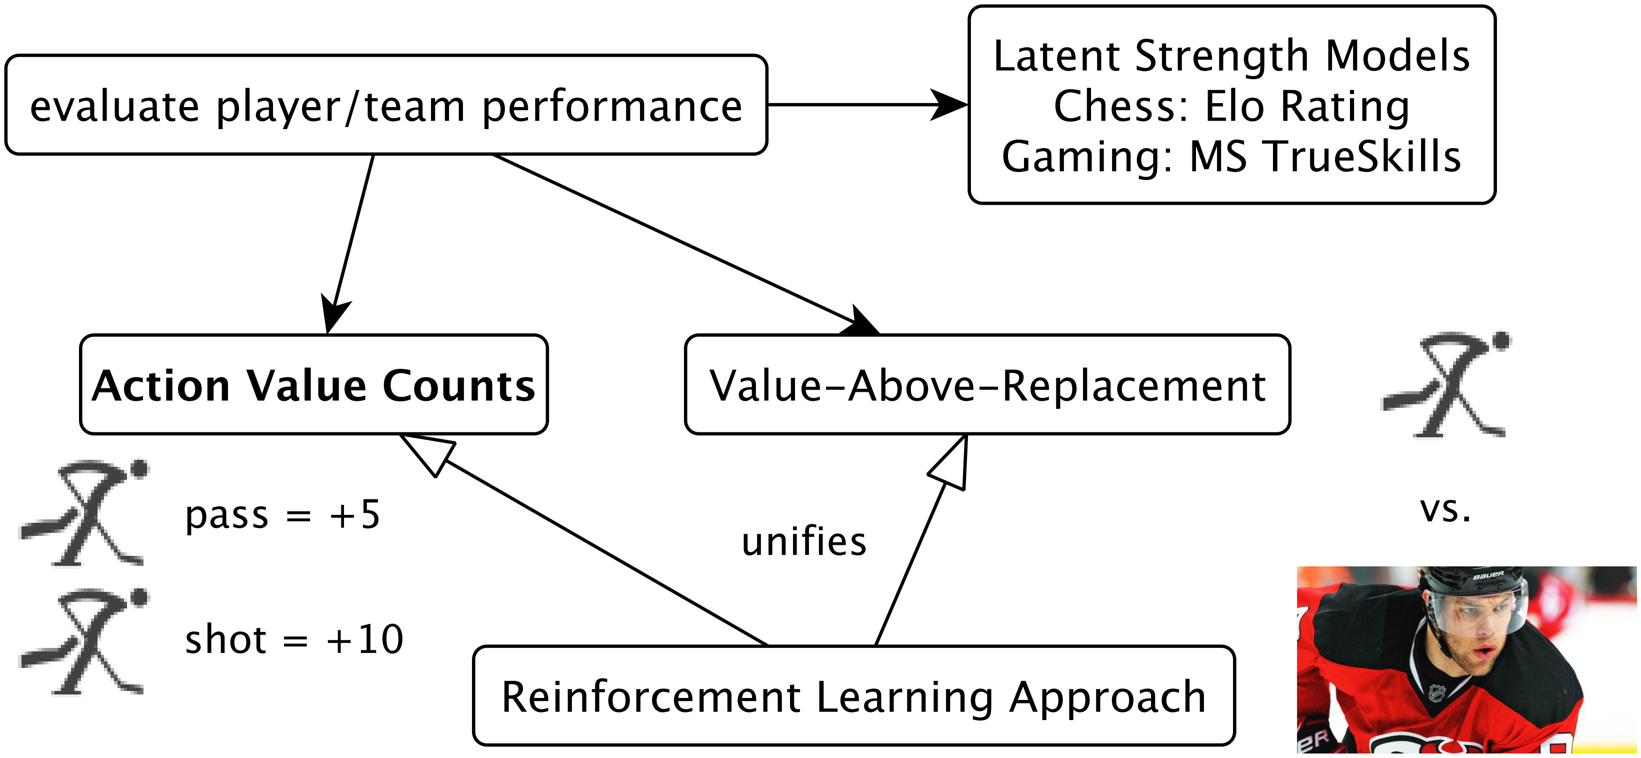

Action Values: Current Approaches

- Like KPIs
- Baseball Statistics
- +/- Score in ice hockey
  - nhl.com
  - Advanced Stats

Problems with Action Counts

- How to combine cou
different actions into
number?
    - e.g. passes + sho
- Ignores context
    - e.g. goal at end o
    more valuable
- Does not capture medium-
term impact: no look-ahead
- Illustration:
[Olympics 2010 Golden Goal]

Solutions for Action Counts

- How to combine counts for different actions into a single number?
  - Use expected utility as measurement scale
- Ignores context
  - Make action value function of *current match state*
- Does not capture medium-term impact: no look-ahead
  - Estimate expected utility with respect to *all future trajectories*

The Q-function

- **The <u>action-value function</u>** in reinforcement learning is just what we need.
- Called Q-function.
- Incorporates
  - context
  - lookahead
- **Familiar in AI, very new in sports analytics!**
- [David Poole's Value Iteration Demo]
-  Q values for actual NHL play, <u>not</u> optimal policy.

OVERVIEW OF METHOD

Framework of Deep Reinforcement Learning (DRL) model

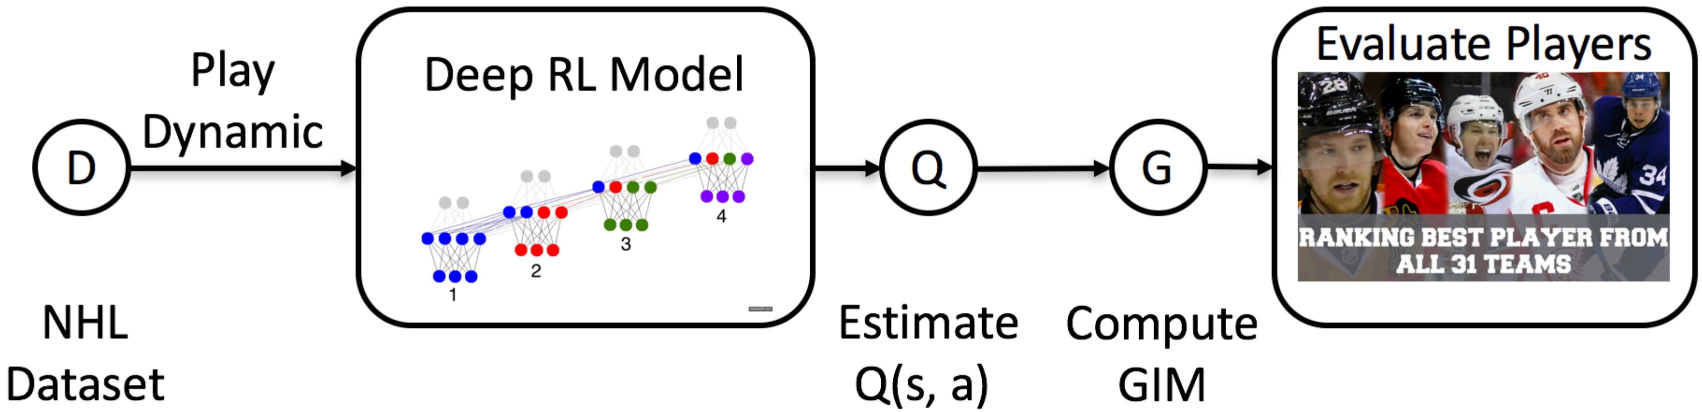

1) Extract play dynamic from NHL dataset.
2) Estimate the Q(s, a) with DRL model.
3) Define a novel Goal Impact Metric (GIM) to value each player.

A Markov Game Model for the NHL




Markov Game Model

- Transition graph with 5 parts:
  - Players/Agents P
  - States ***S***
  - Actions ***A***
  - Transition Probabilities ***T***
  - Rewards ***R***
- **Transitions, Rewards depend on state and *tuple* of actions, one for each agent.**

Markov Game Model: Action Types

13 Action Types

Action Types
Blocked Shot
Faceoff
Giveaway
Goal
Hit
Missed Shot
Shot
Takeaway
...

STATE SPACE

- At each time, we observe the following features
- Model also captures match history (more below)

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}

Example State Trajectory on Rink

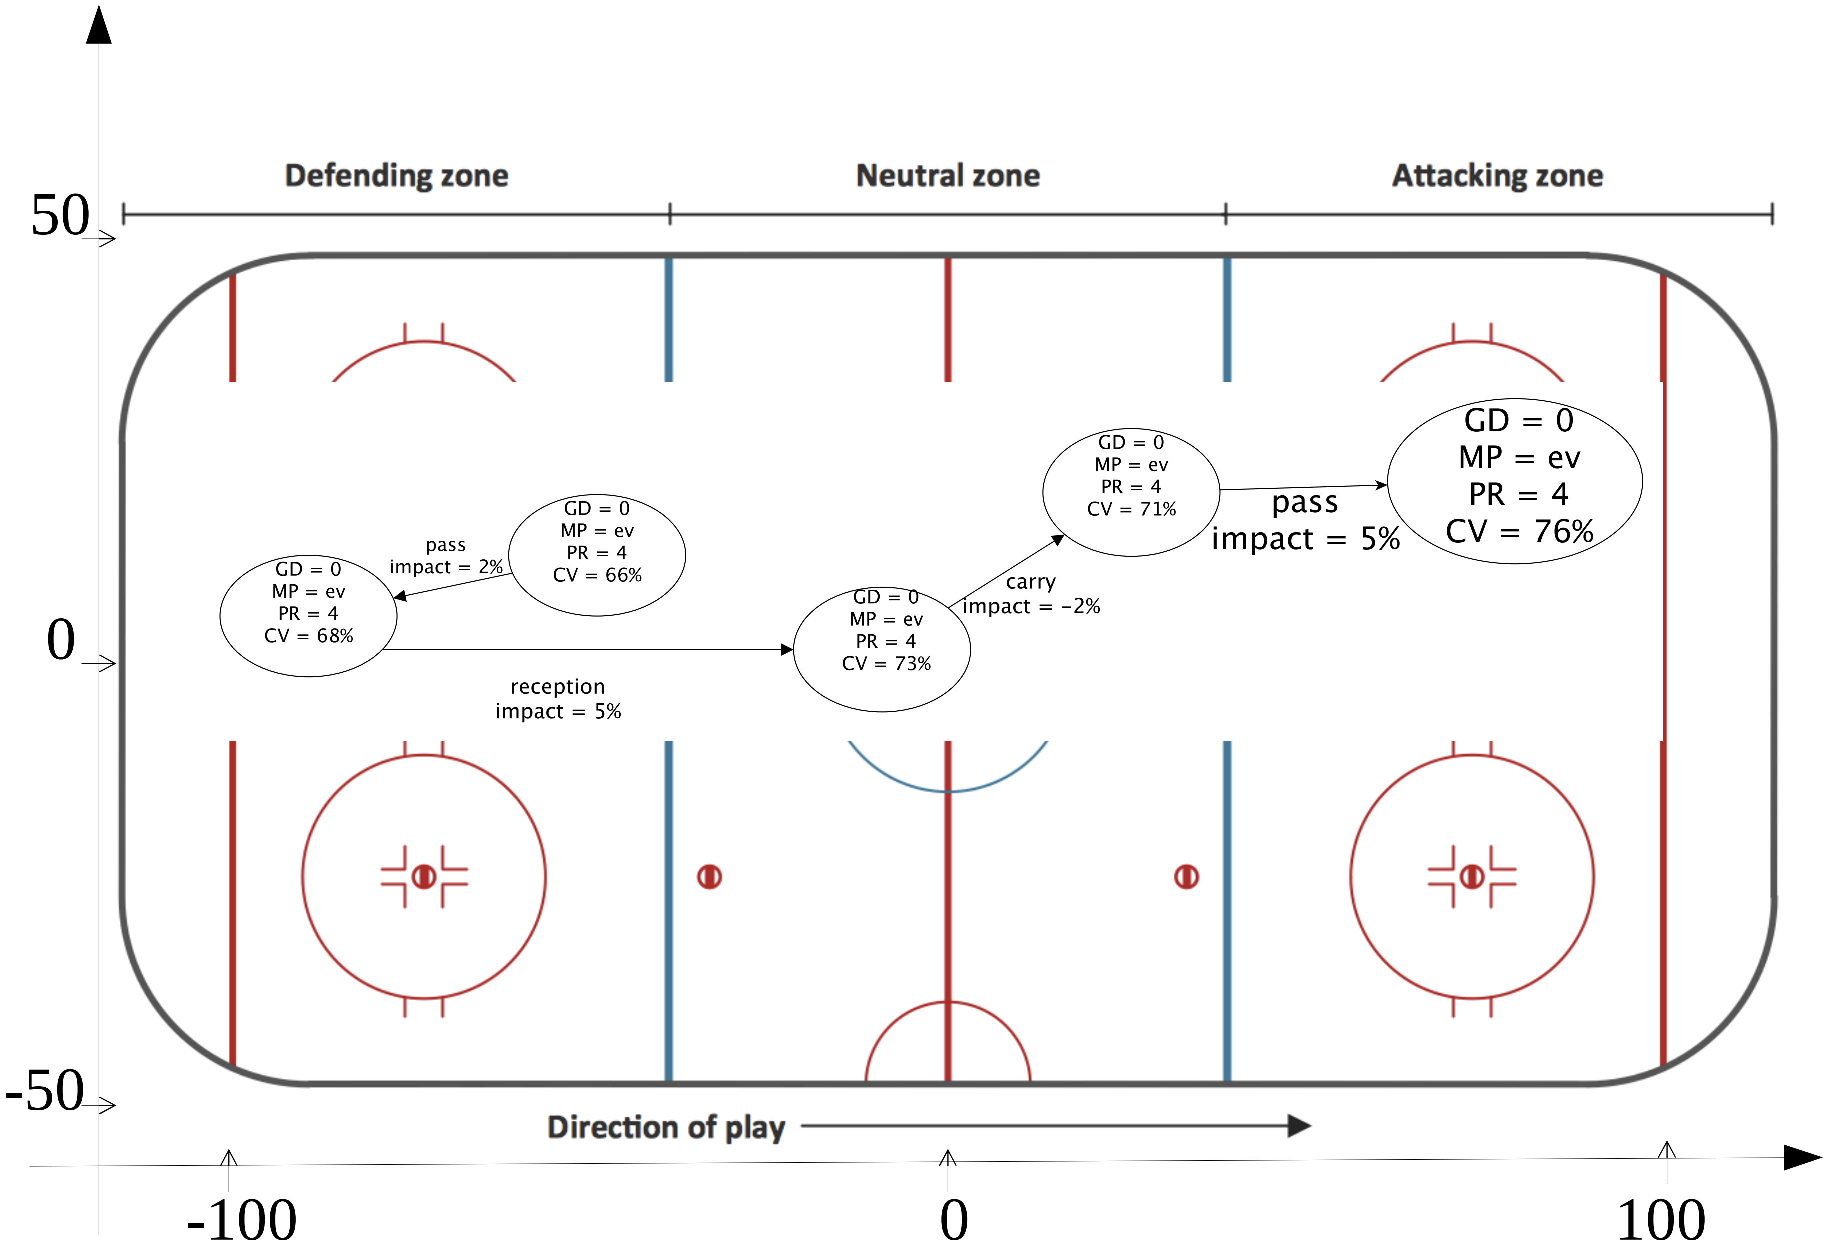

Rewards

options for reward functions

episode

NBA: **Points** from Possession
Cervone et al. 2014
NFL: **Points** from Possession
nflscrapR Yurko et al. 2018
Chan and Puterman 2019
NHL: **Next Goal** <u>(our work)</u>

episode

NHL: **penalties**
Routley and Schulte 2015

final outcome

win probabilities
AlphaGo
Canadian Tire
Hockey:
Pettigrew 2015
Schulte et al. 2017

Learning an Action-Value Function for the NHL

PIPELINE

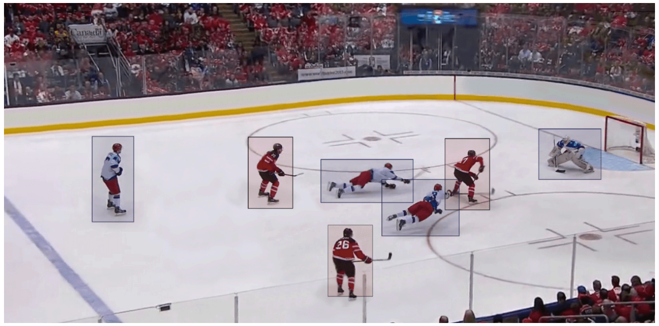

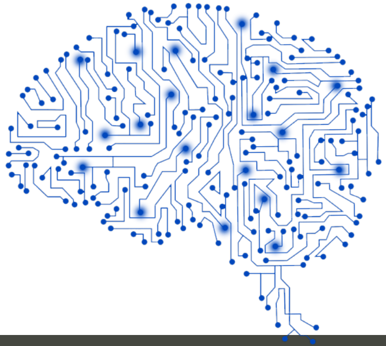

- Computer Vision Techniques: Video tracking
- Play-by-play Dataset
- Large-scale Machine Learning

Sports Data Types

- **Complete Tracking**: which player is where when. Plus the ball/puck. ★
- **Box Score**: Action Counts.
- **Play-By-Play**: Action/Event Sequence.

Tracking Data

- Basketball SportsVU since 2011
- New for NFL Next Gen Stats
- Coming to the NHL?
- Holy Grail: Tracking from Broadcast Video
- Sportlogiq, Stats

Box Score

[Oilers vs. Canucks]

Play-By-Play

- [Successive Play Sequences]
- [nhlscraper, nflscraper]

Our Play-By-Play Data

- Source: SportLoqig
- 2015-16
- **Action Locations**

SportLogiq	
Teams	31
Players	2,233
Games	1,140
Events	3M+

DRL MODEL

- Recurrent LSTM network
- Dynamic trace back to previous possession change

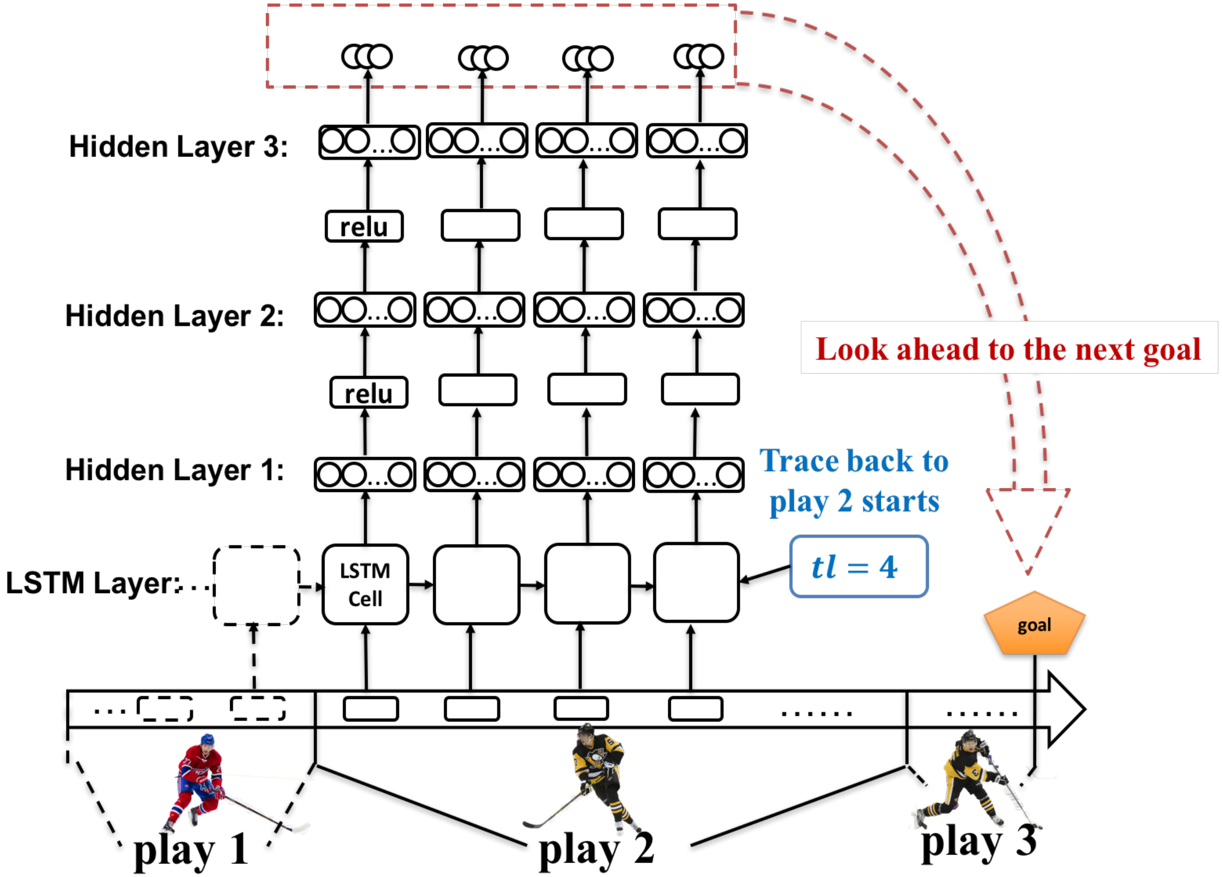

Value Ticker: Temporal Projection

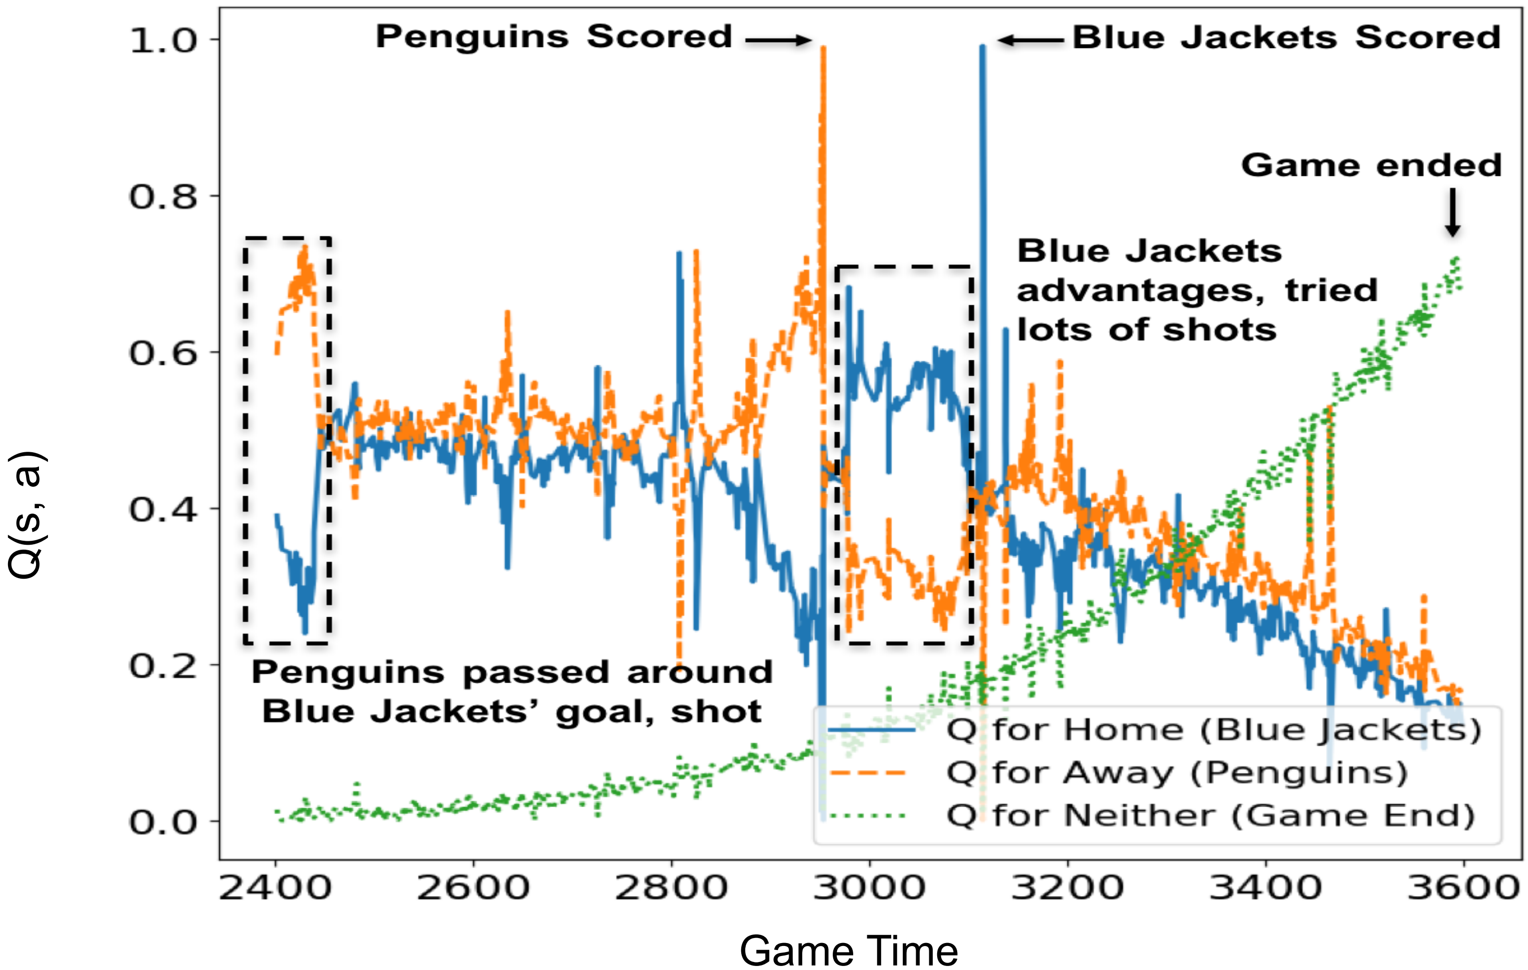

Spatial Projection

Q-value for the action “shot” action over the rink.

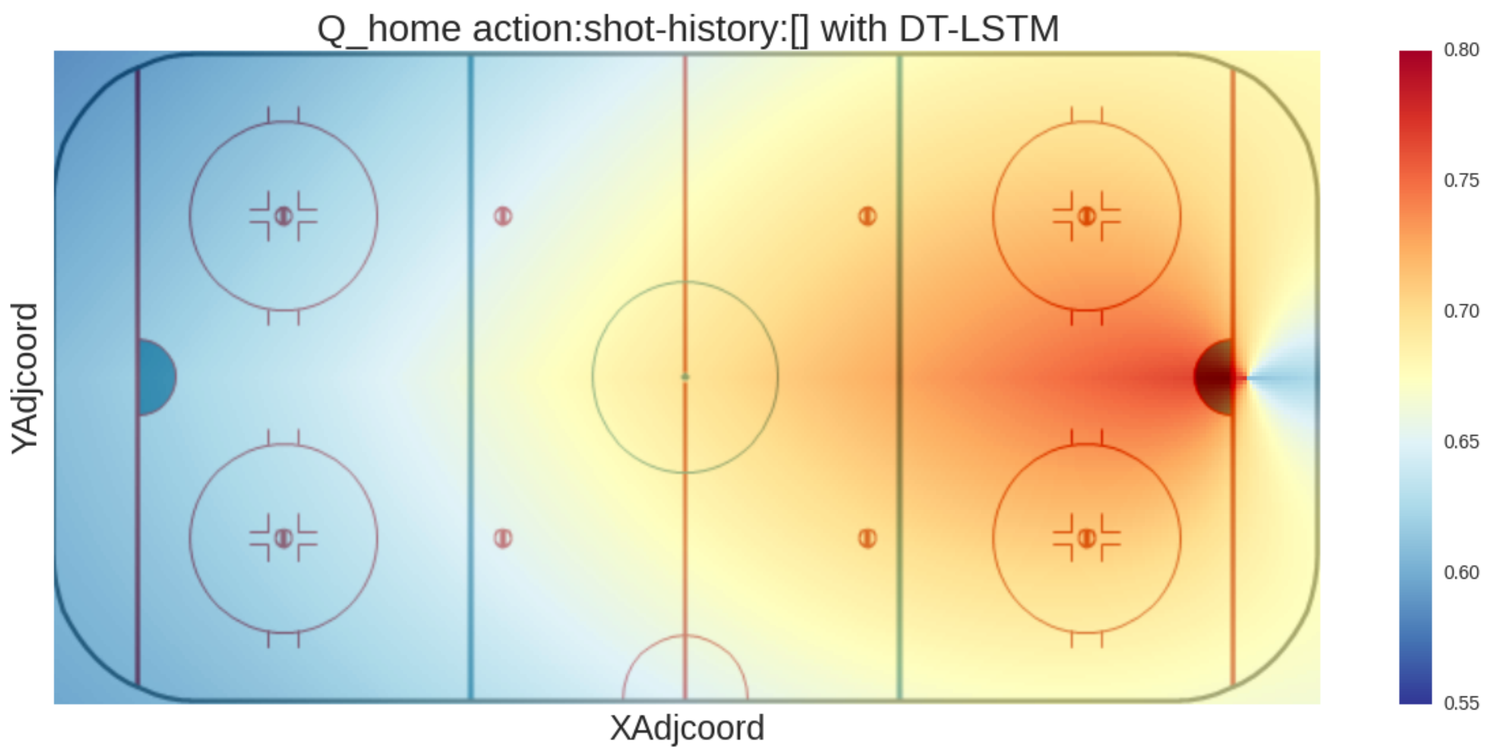

Evaluating Player Performance

The Impact of an Action

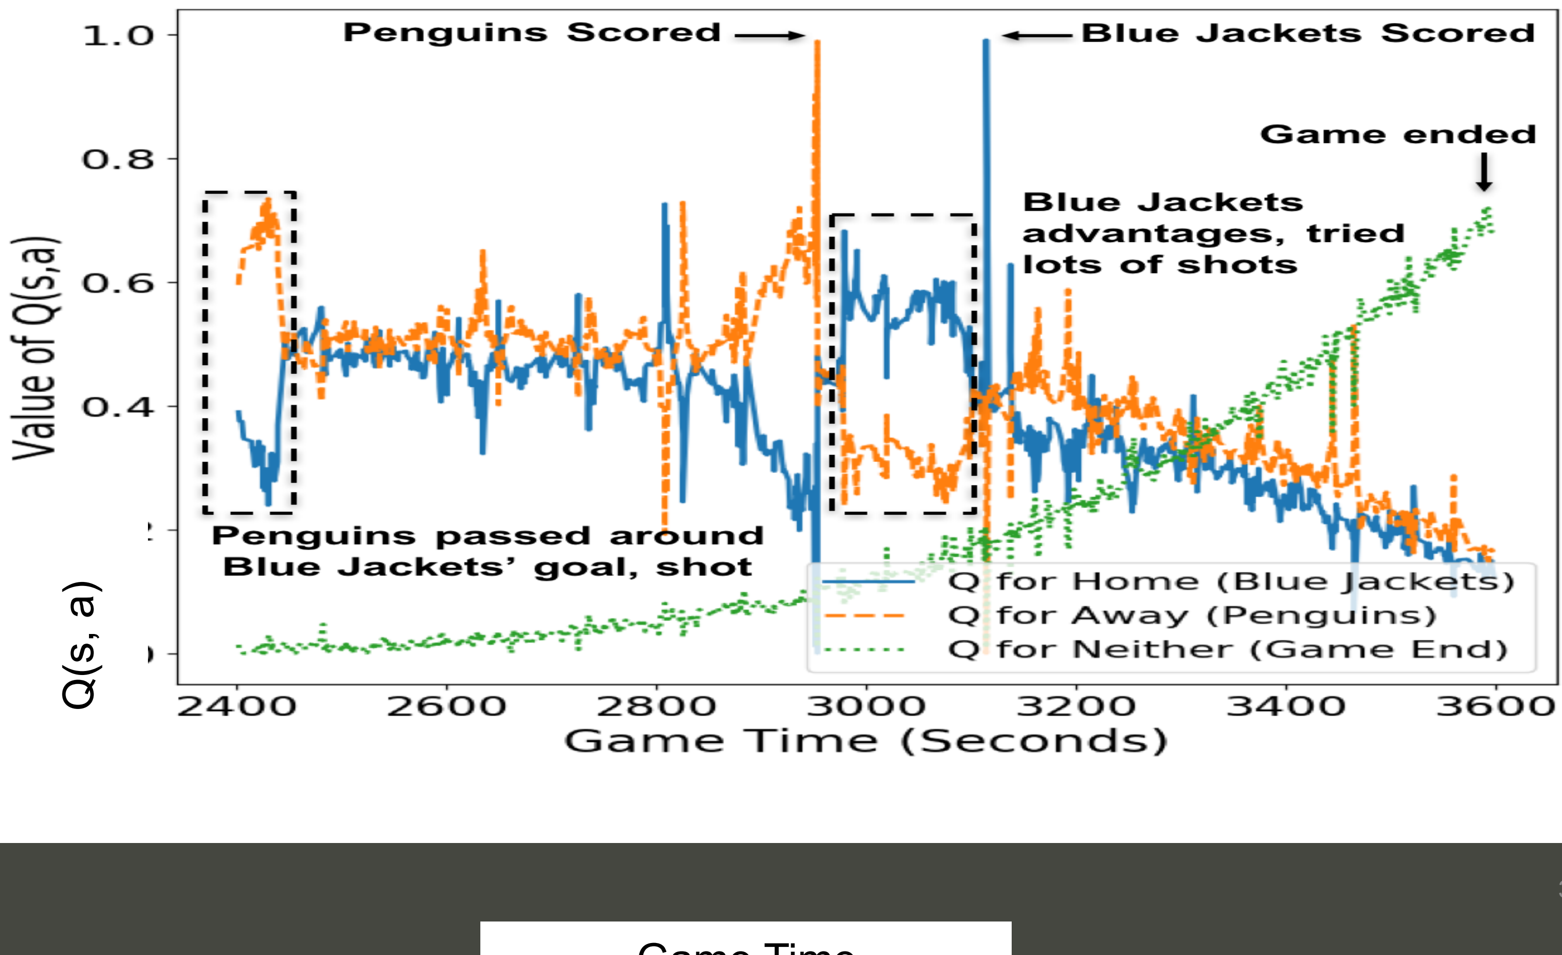

Goal Impact Metric

1. Apply the impact of an action to the player performing the action
2. Sum the impact of his actions over a game to get his net game impact.
3. Sum the net game impact of a player over a single season to get his net season impact.

Evaluation

- No ground truth for player ranking
- Compare with success metrics known to be relevant
- Other desiderata (consistency, predictive power) Franks et al. 2016

Franks, A. M.; D'Amour, A.; Cervone, D. & Bornn, L. (2016), 'Meta-analytics: tools for understanding the statistical properties of sports metrics', *Journal of Quantitative Analysis in Sports* ***12(4), 151--165.***

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

OTHER SUCCESS METRICS

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**

PREDICTIVE POWER, CONSISTENCY

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**.

PREDICTIVE POWER, CONSISTENCY

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 computed over the entire season

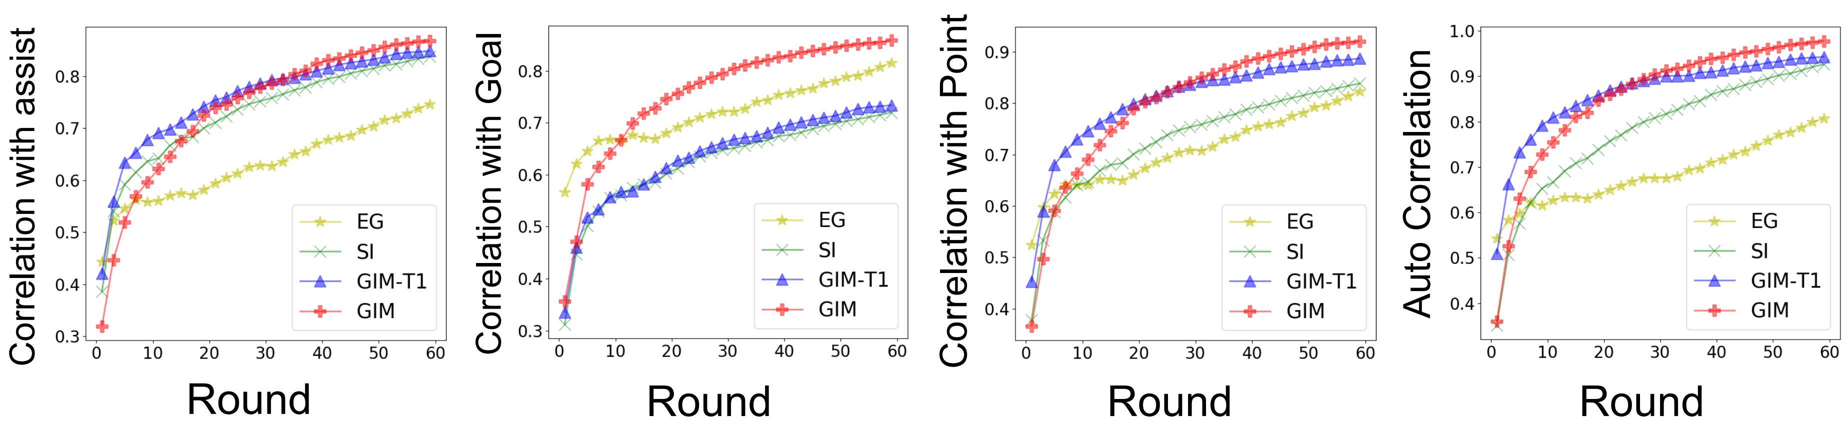

GOAL IMPACT AND SALARY

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).

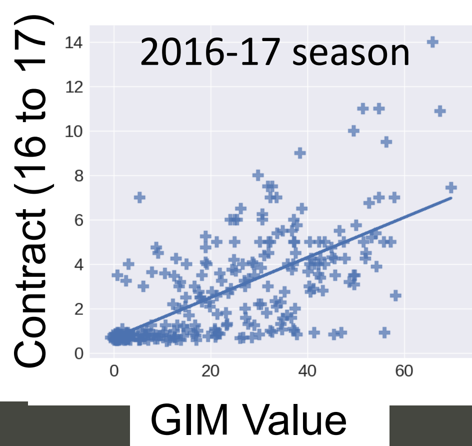

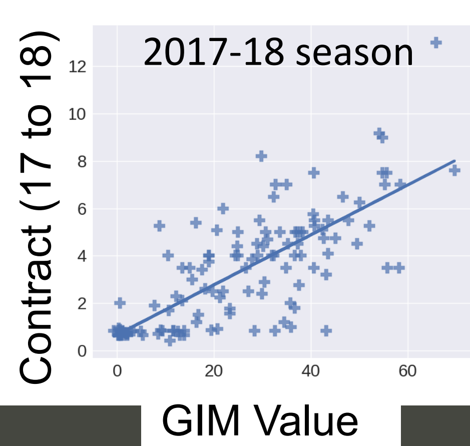

RELATED WORK

Markov Value Function Based Players Evaluation

Year	Venue	Authors	Name	Sports
2019	MIT Sloan	Javier Fernández, Luke Bornn, et.al	Decomposing the Immeasurable Sport: A deep learning expected possession value framework for soccer	Soccer
2018	IJCAI	Guiliang Liu and Oliver Schulte	Deep reinforcement learning in ice hockey for context-aware player evaluation	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

More on the Value Function

- "We assert that most questions that coaches, players, and fans have about basketball, particularly those that involve the offense, can be phrased and answered in terms of EPV [i.e. the value function]." Cervone, Bornn et al. 2014.
- We have seen how the action-value function can be used to rank players
- Can also be ranked to give decision advice to coaches (e.g. Wang et al. 2018)

Wang, J.; Fox, I.; Skaza, J.; Linck, N.; Singh, S. & Wiens, J. (2018), The Advantage of Doubling: A Deep Reinforcement Learning Approach to Studying the Double Team in the NBA, *in 'MIT Sloan Sports Analytics Conference'.*

Future Work

Supported by a Strategic Project Grant with SportLogiq

Pascal Poupart
Waterloo

Greg Mori
SFU

Luke Bornn
SFU, Sacremento Kings

Increasing Realism and Accuracy

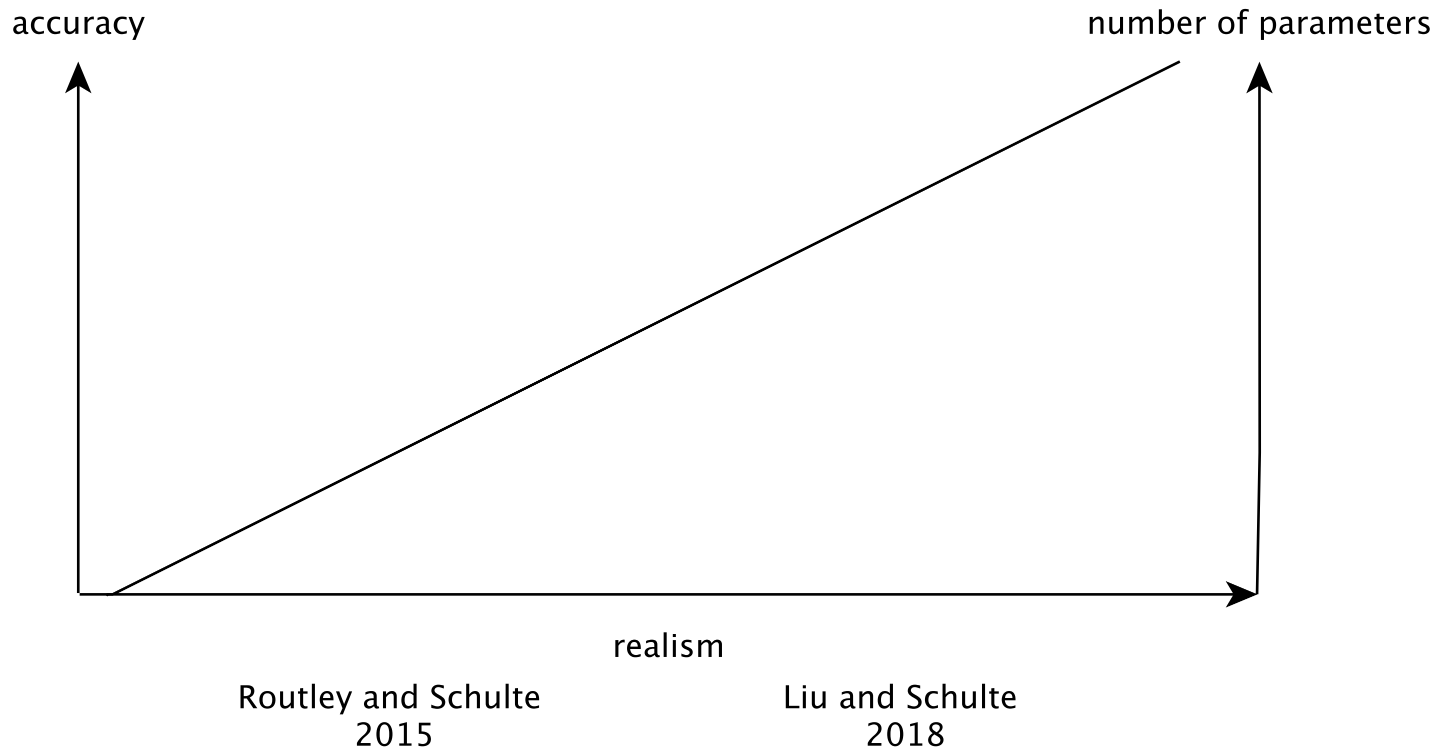

Increasing Realism and Accuracy: Hierarchical Models

- Current Model pools data from all players and teams ➔ average team/player
- How can we capture patterns specific to players/teams?
- Current sports analytics: Use a **hierarchical model**
  - aka shrinkage, multi-level, random effects
- How can we represent individual patterns in a decision process model?
  - In a deep decision process model?

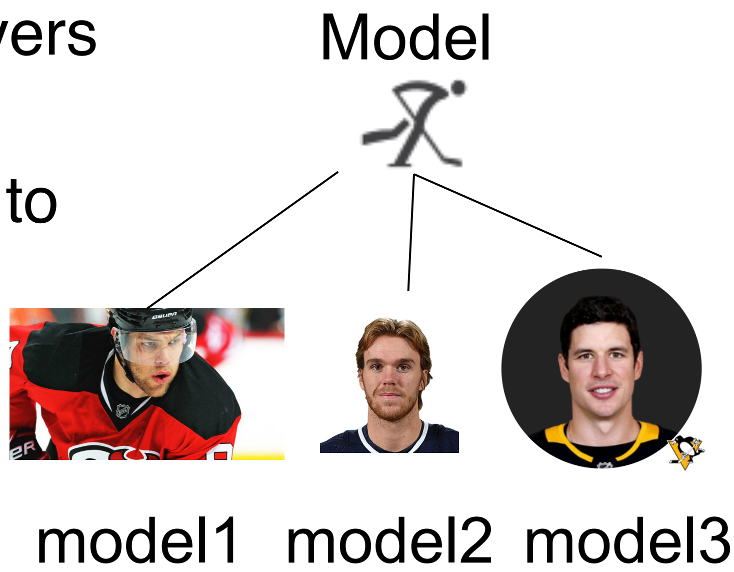

Yurko, R.; Ventura, S. & Horowitz, M. (2018), 'nflwar: A reproducible method for offensive player evaluation in football', *arXiv preprint arXiv:1802.00998.*

Interpretation

- Goal: Explain why the neural net assigns high/low values to some states

1. Mimic Learning (Liu and Schulte 2018)
2. Likely Future Trajectories (Khan, Poupart et al. 2011)

hat-if scenarios?

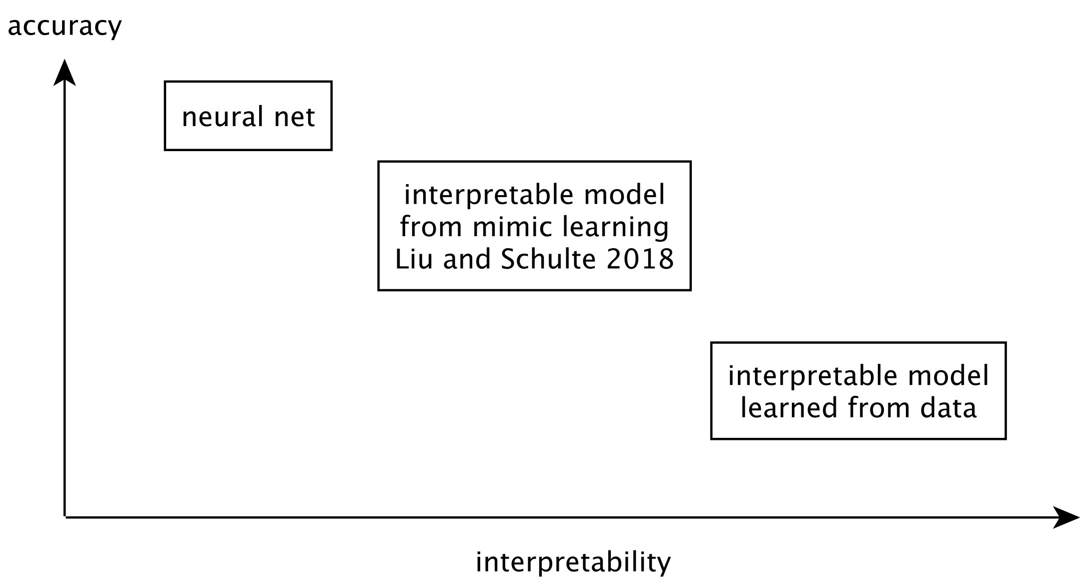

Liu, G.; Schulte, O.; Zhu, W. & Li, Q. (2018), Toward Interpretable Deep Reinforcement Learning with Linear Model U-Trees, *in European Conference on Machine Learning and Knowledge Discovery in Databases (ECML)', pp. 1—16.*

Learning at Higher Scales

- Intuitively, players and coaches think in terms of **plays** (maneuvers).
- Related to RL concepts
  - Options
  - Task hierarchies
- Common Example in Sports Analytics: Trajectory Clustering

Jonsson, A. & Barto, A. G. (2001), Automated state abstraction for options using the U-tree algorithm, *in 'Advances in neural information processing systems', pp. 1054--1060.*

NFL Example: Route Types as Higher-Scale Options

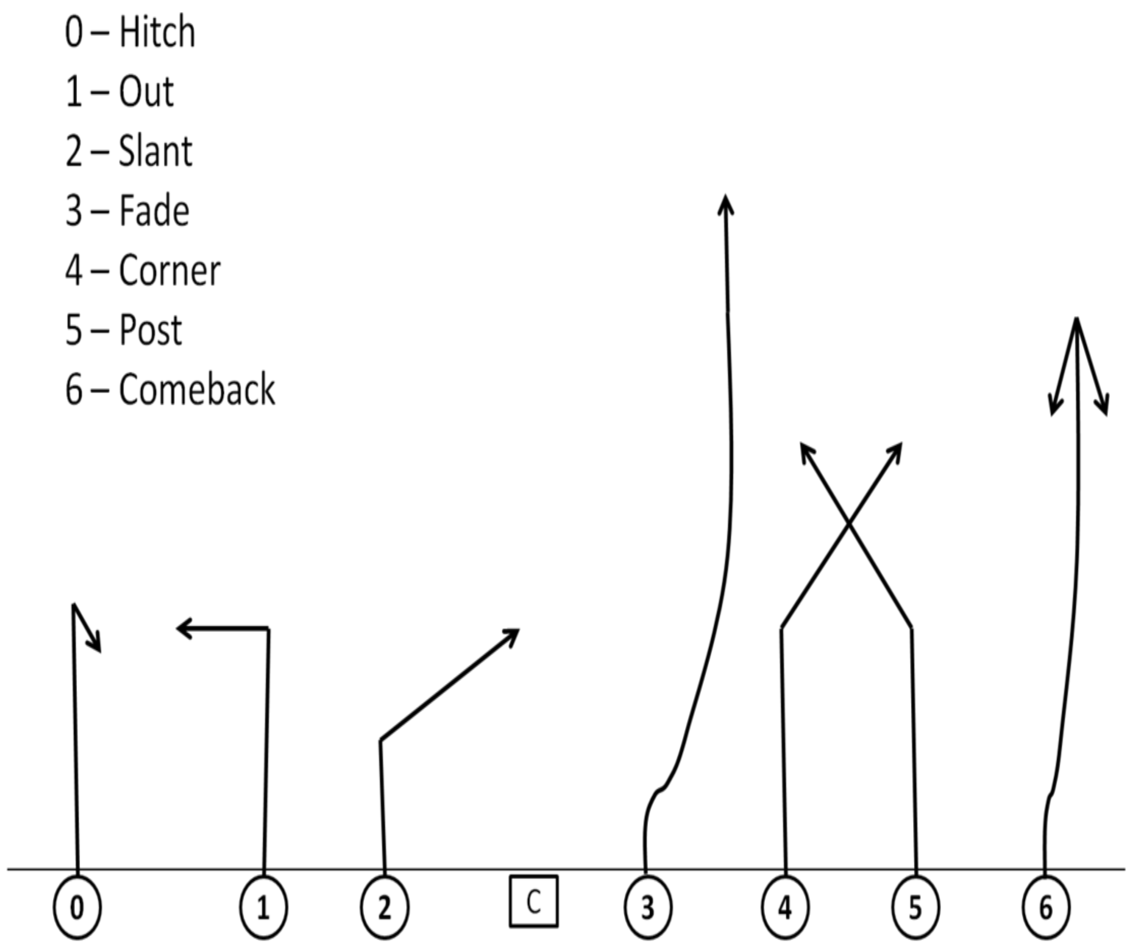

Figure due to Chu et al. 2019

Conclusion

- Modelling ice hockey dynamics in the NHL
- A new context-aware method for evaluating actions and players
- A configurable and scalable Markov Game model that incorporates context and long-term effects of *all* actions
- Learning an action-value function is a powerful AI-based approach to supporting decisions in sports

THANK YOU!

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