Deep Soccer Analytics: Learning An Action-value Function For Evaluating Soccer Players

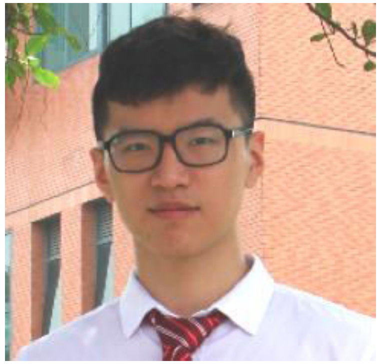


Guiliang Liu   Yudong Luo   Oliver Schulte   Tarak Kharrat

GHENT, BELGIUM, 14-18 SEPTEMBER
ECML PKDD 2020

SFU SIMON FRASER UNIVERSITY

SPORT LOGiQ

UNIVERSITY OF LIVERPOOL

Overview

Learning an Action-Value Q Function for soccer player evaluation:

- Modeling *play dynamics* based on a Markov Game Process (s,a,r).
- Build a *Deep Reinforcement Learning* (DRL) model to compute action-value Q function.
- Compute a *Game Impact Metric (GIM).*
- *Rank player* and evaluate their performance.
- *Examine* the model with a Multi-League play-by-play dataset.

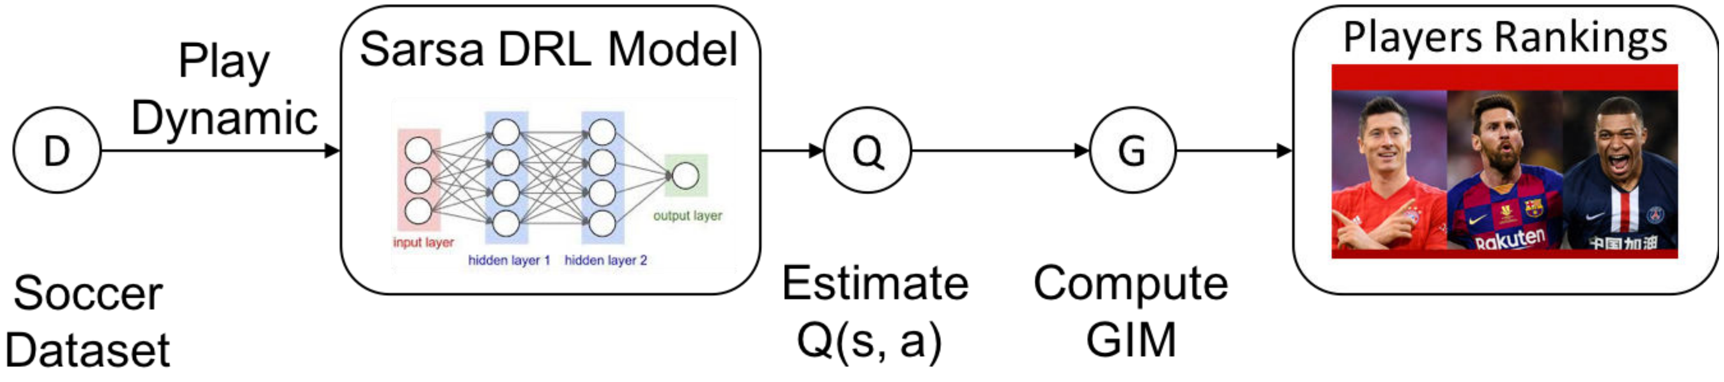

Motivation

Why Deep Reinforcement Learning (DRL):

Previous Model-based methods [1,2,3]:

- Explicitly construct a Markov Model.
- Model building and function learning are *independent*.
- Infeasible for large dataset.

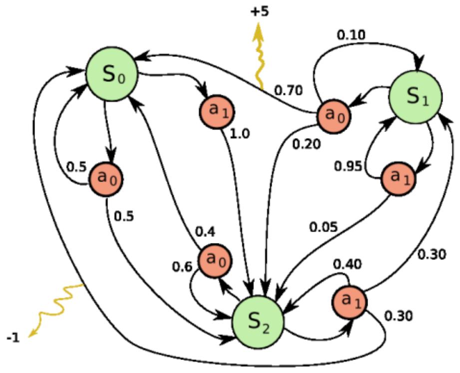

- Requires discretizing the continuous features.
- Huge state numbers (e.g., 10 features each with 10 dimension indicates 10^{10} states).
- Complex transitions.

Our Sarsa DRL model:

- Model-Free RL (no pre-built models).
- An end-to-end model (no data pre-processing, no intermediate model).
- Generalize to large dataset (mini-batch gradient descent fits dataset with any size).

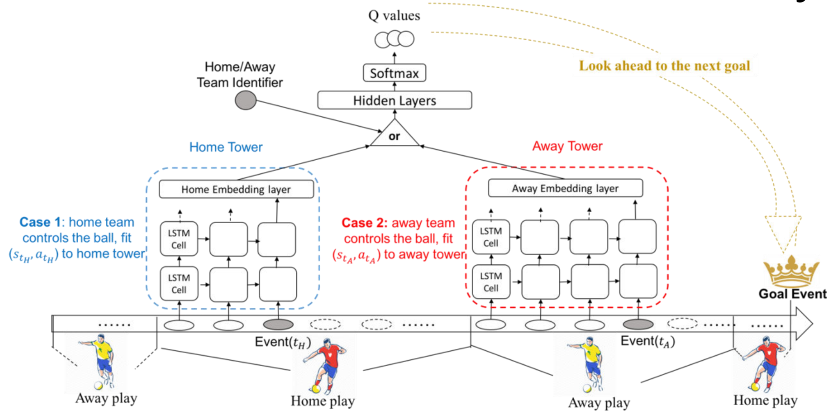

[1] Routley, Kurt, and Oliver Schulte. "A Markov Game model for valuing player actions in ice hockey." Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence. 2015.
[2] Schulte, Oliver, et al. "A Markov Game model for valuing actions, locations, and team performance in ice hockey." Data Mining and Knowledge Discovery 31.6 (2017): 1735-1757.
[3] Cervone D, D'Amour A, Bornn L, Goldsberry K (2016) A multiresolution stochastic process model for predicting basketball possession outcomes. J Am Stat Assoc 111(514):585–599

Preliminary Result

Visualizing the Q functions learned by DRL:

Temporal Projection

- Q values for a game between Fulham (Home) and Sheffield Wednesday (Away), which has happened on Aug. 19th, 2017.
- Q functions represents the probability of home/away team score the next goal or nobody score.

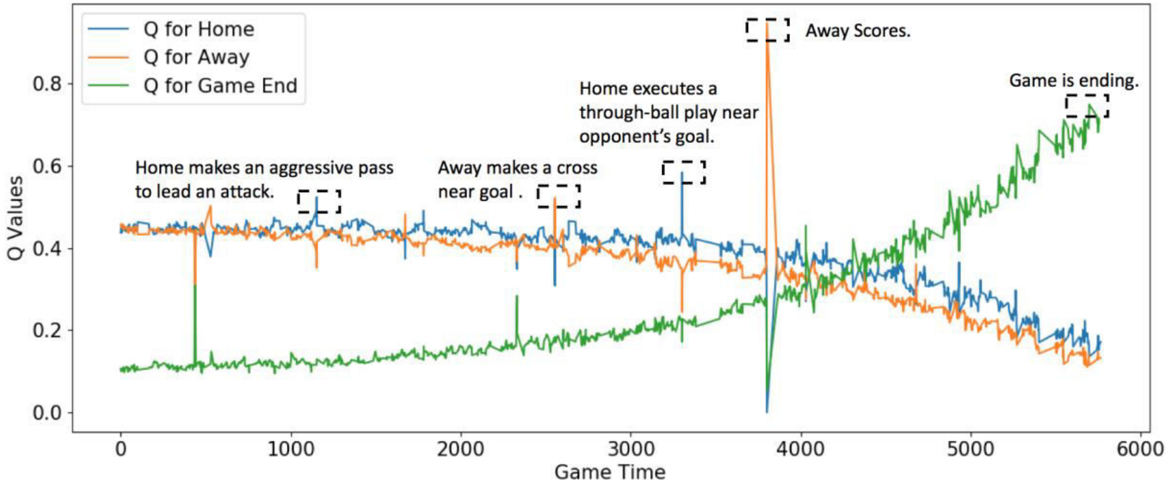

Spatial Projection

- Q functions for actions: shots and tackles.
- Q function (learned by DRL) generalizes from observed states and actions to those that have not occurred.

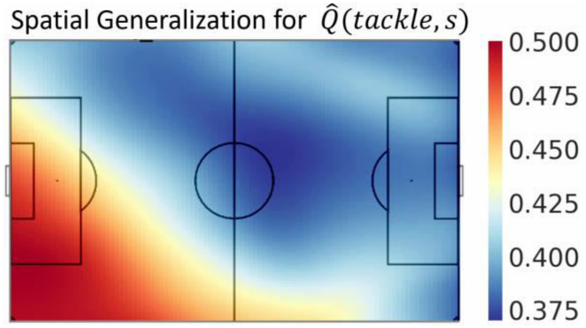

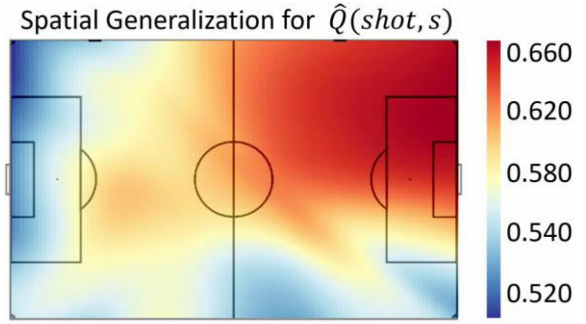

THANK YOU!

Dataset and Preprocessing

A play-by-play soccer *dataset* for sports analytic

- Records the actions of on-the-ball players and the *spatial* and the *temporal* context features.
- *Multiple leagues*, multiple teams and players.

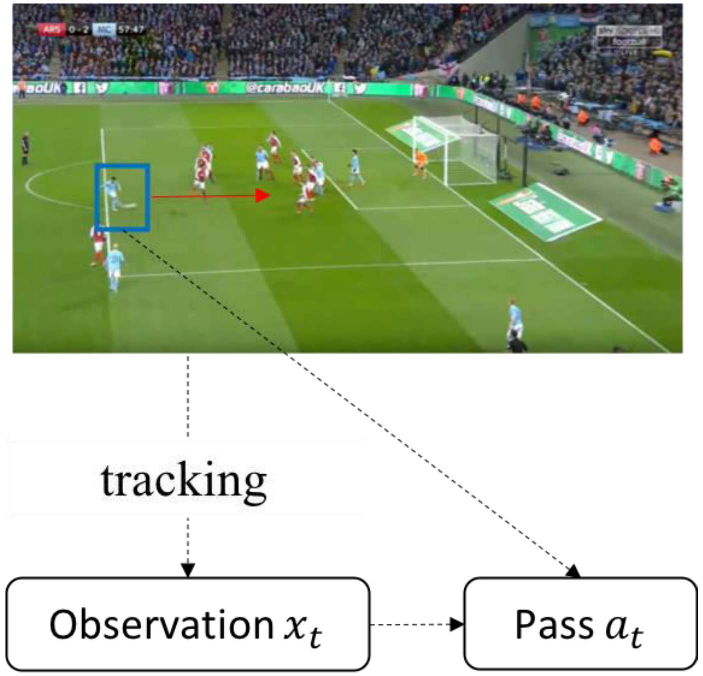

Name	Type	Range
Game Time Remaining	Continuous	[0, 100]
X Coordinate of ball	Continuous	[0, 100]
Y Coordinate of ball	Continuous	[0, 100]
Manpower Situation	Discrete	[-5, 5]
Goal Differential	Discrete	$(-\infty, +\infty)$
Action	Discrete	one-hot representation
Action Outcome	Discrete	{success, failure}
Velocity of ball	Continuous	$(-\infty, +\infty)$
Event Duration	Continuous	$[0, +\infty)$
Angle between ball and goal	Continuous	$[-\pi, +\pi]$
Home or Away Team	Discrete	{Home, Away}

Table 2: Complete feature list. For the feature manpower situation, negative values indicate short-handed, positive values indicate power play.

Dataset	F24
Events	4,679,354
Players	5,510
Games	2,976
Teams	164
Leagues	10
Season	2017-18
Place	Europe

Table 3: Dataset statistics. The basic unit of this dataset is *event*, which describes the game context and the on-the-ball action of a player at a time step.

The *dataset* utilizes adjusted spatial coordinates

- Both the X-coordinates and Y-coordinates are adjusted to [0, +100].
- We reverse the coordinates when the team in possession attacks towards the left
- The play flows from left to right for either team on the adjusted soccer pitch.

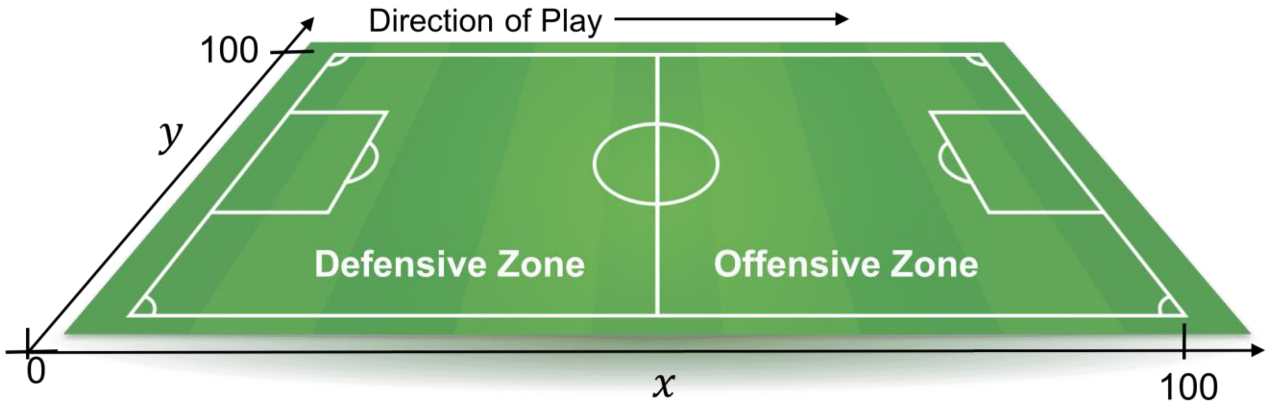

MP=Manpower, GD=Goal Difference, OC = Outcome, S=Succeed, F=Fail, H=Home, A=Away, T=Team who performs action, GTR = Game Time Remain, ED = Event Duration

GTR	X	Y	MP	GD	Action	OC	Velocity	ED	Angle	T	Reward
35m44s	87	26	Even	1	simple pass	S	(2.2, 1.7)	11.0	0.19	H	[0,0,0]
35m42s	90	17	Even	1	standard shot	F	(1.5, -4.5)	2.0	0.11	H	[0,0,0]
35m42s	99	44	Even	1	save	S	(0, 0)	0.0	0.06	A	[0,0,0]
35m9s	100	1	Even	1	cross	S	(0.0, -1.3)	33.0	0.0	H	[0,0,0]
35m7s	85	56	Even	1	simple pass	S	(-7.3, 27.6)	2.0	0.39	H	[0,0,0]
35m5s	92	67	Even	1	simple pass	S	(3.6, 5.4)	2.0	0.28	H	[0,0,0]
35m4s	97	50	Even	1	corner shot	S	(5.1, -16.2)	1.0	1.74	H	[0,0,0]
35m4s	100	50	Even	1	goal	S	(0, 0)	0.0	0.0	H	[1,0,0]
....	...	...	....	...	....	...	....	...	.....	.	......
3m41s	62	96	Even	2	long ball	F	(4.5, 9.3)	9.0	0.08	A	[0,0,0]
3m39s	19	89	Even	2	clearance	S	(-21.5, -3.2)	2.0	0.07	H	[0,0,0]
3m35s	24	100	Even	2	throw in	S	(1.3, 2.7)	4.0	0.09	A	[0,0,0]
3m33s	27	96	Even	2	simple pass	S	(1.1, -2.2)	2.0	0.1	A	[0,0,0]
3m31s	12	95	Even	2	cross	S	(-7.5, -0.5)	2.0	0.07	A	[0,0,0]
3m28s	6	46	Even	2	simple pass	S	(-1.7, -16.3)	3.0	0.79	A	[0,0,0]
3m26s	14	48	Even	2	standard shot	S	(3.8, 1.3)	2.0	0.44	A	[0,0,0]
3m26s	0	50	Even	2	goal	S	(0, 0)	0.0	0.0	A	[0,1,0]

Table 1: A data sample featuring team scoring: a sequence of events where home team scores and then away team scores. The rewards [1,0,0] and [0,1,0] indicate the scoring event of home team and away team respectively (see Section 4.1). We skip some events in the middle due to space issues.

Play Dynamic in Soccer

A Markov model for soccer games.

- Two **agents**: Home and Away
- An **action** a_t (one-hot representation) denotes the movements of players who control the ball.
- An **observation** is a feature vector x_t specifying a value of the features.
- A game **state** records the complete sequence $s_t \triangleq (x_t, a_{t-1}, x_{t-1}, \ldots, x_0)$.
- The **reward** r_t is a vector of goal values g_t that specifies which team (Home, Away) scores.

An action-value Q function.

- Divide a soccer game into **goal-scoring episodes**. 1) **starts** at the beginning of the game, or immediately after a goal, and 2) **terminates** with a goal or at the end of the game.
- The **next-goal Q-function** represents the probability that the home resp. away team scores the goal at the end of the current goal-scoring episode.

$$Q_{team}(s, a) = P(goal_{team} = 1 | s_t = s, a_t = a)$$

Two-Tower Dynamic Play LSTM (TTDP-LSTM)

- **Three output nodes** at each time step: $\hat{Q}_{home}$, $\hat{Q}_{away}$, and $\hat{Q}_{end}$.
- ***Two towers***: fits home and away data separately.
- ***Dynamic possession-LSTM***: 1) apply a dynamic trace length. 2) trace back to the beginning of a play.
- **Temporal Difference** (TD) Loss:

$$\mathcal{L}(\theta) = \sum_{team \in T} \mathbb{E}\left[(r_{team,t+1} + \hat{Q}_{team}(s_{t+1}, a_{t+1}) - \hat{Q}_{team}(s_t, a_t))^2\right]$$

- **Training settings**:
  1) Stacked (a tow layer) LSTM
  2) Minibatch training.
  3) max trace length is 10.

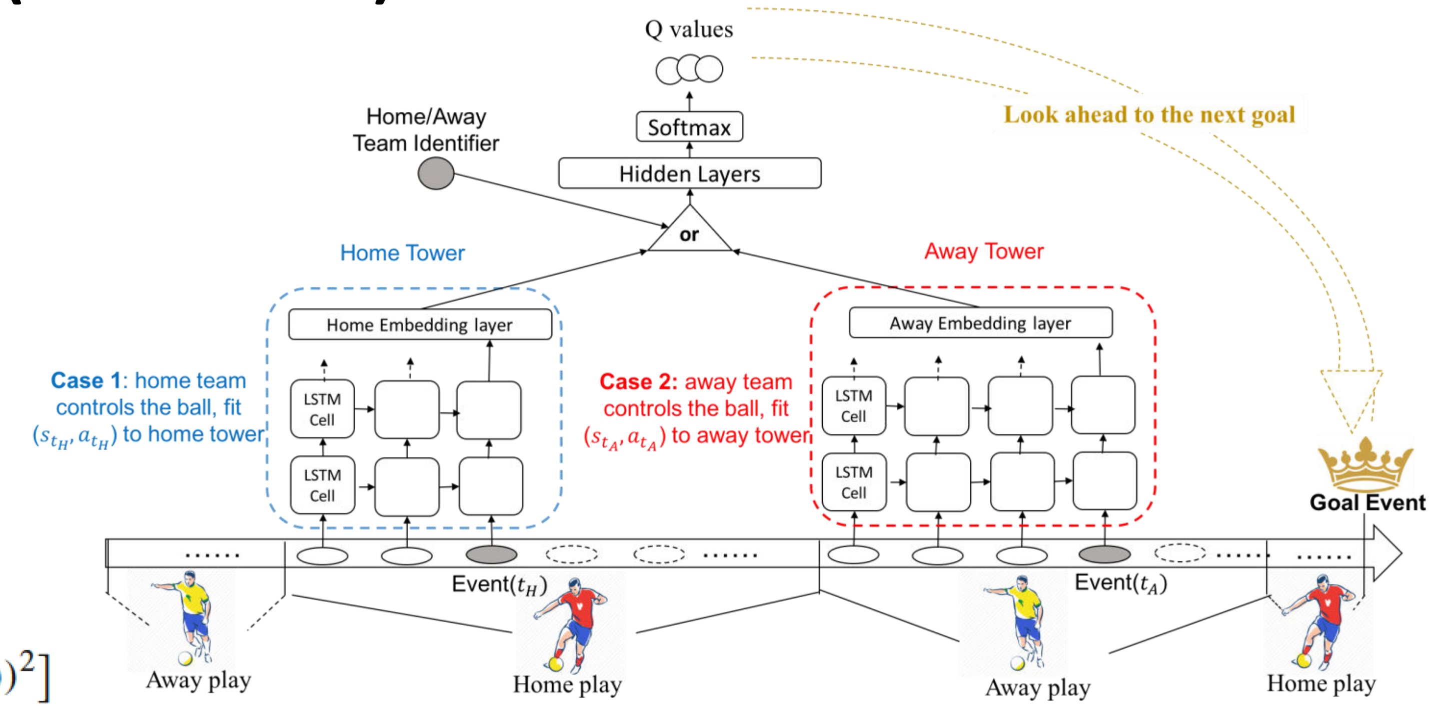

Model Validation: Q Values

Illustration of Temporal and Spatial Projection:

Go back to slide 4

Calibration Quality for the learned Q-function:

- Evaluate how well our learned Q-function fits the observed scoring frequencies.
- Discretized **game context**:
  1) Manpower (Short Handed (SH), Even Strength (ES), Power Play (PP)).
  2) Goal Differential ($\geq$-3, -2, -1, 0, 1, 2, $\geq$ 3).
  3) Period (1 (first half), 2 (second half)).
- Measures (how close they are):

1) Empirical Scoring Probabilities

$$Q^{obs}_{team}(A) = \frac{1}{|A|} \sum_{s \in A} goal^{obs}_{team}(s)$$

2) Estimated Scoring Probabilities

$$\hat{Q}_{team}(A) = \frac{1}{|A|} \sum_{s \in A} \hat{Q}_{team}(s, a)$$

Man.	Goal.	P.	$\|A\|$	TT_Home	TT_Away	TT_MAE	Markov_MAE
ES	-1	1	73176	0.4374	0.4159	0.0052	0.1879
ES	-1	2	96408	0.3496	0.3025	0.0782	0.1783
ES	0	1	356597	0.4437	0.4272	0.026	0.1908
ES	0	2	160080	0.356	0.3077	0.0814	0.1792
ES	1	1	88726	0.4402	0.4128	0.0335	0.1899
ES	1	2	119901	0.3459	0.295	0.077	0.1787
PP	-1	1	876	0.4366	0.4045	0.1752	0.1937
PP	-1	2	3319	0.352	0.2911	0.0668	0.1685
PP	0	1	3183	0.4414	0.403	0.1308	0.187
PP	0	2	7183	0.3579	0.2855	0.0841	0.1804
PP	1	1	1316	0.4391	0.3949	0.115	0.1825
PP	1	2	7676	0.356	0.2862	0.1121	0.1792

Table 4: Calibration Results. TT_Home and TT_Away report the average scoring probability $\hat{Q}_{team}(A)$ estimated by our TTDP-LSTM model. Here we compare only Q values for pass and shot as they are frequent and well-studied actions. TT_MAE is the Mean Absolute Error (MAE) between estimated scoring probabilities from our model and empirical scoring probabilities. For comparison, we also report a Markov_MAE which applies the estimates from a discrete-state Markov model [Schulte et al., 2017b].

Player Evaluation Metric

Goal Impact Metric (GIM):

- Compute the impact of an action by
  1) How much it changes the expected total reward of a player's team.
  2) Or the difference in expected total reward before and after the player acts.

- GIM calculates the total impact of a player's action:

$$impact^{team}(s, a, s', a') \equiv Q_{team}(s', a') - \mathbb{E}_{s',a'}[Q_{team}(s', a')|s, a]$$

$$GIM^{i}(D) \equiv \sum_{s,a,s',a'} n[s, a, s', a', pl' = i; D] \cdot impact^{team}(s, a, s', a')$$

Q_{home}	$impact^{home}$
0.8	0.4
0.4	-0.2
0.6	

Q Value Above Average Replacement (QAAR):

- The QAAR metric compares 1) the expected total future reward given that player i acts next, to 2) the expected total future reward given that a random replacement player acts next:

$$QAAR^{i}(D) \equiv \sum_{s,a} n[s, a, pl' = i; D]\Big(\mathbb{E}_{s',a'}[Q_{team}(s', a'|s, a, pl' = i)] - \mathbb{E}_{s',a'}[Q_{team}(s', a')|s, a]\Big)$$

- *Proposition*: For each player i recorded in our play-by-play dataset D, $\boldsymbol{QAAR^{i}(D) = GIM^{i}(D)}$:

Mimic Decision Tree

Understanding Impact Values with Mimic Decision Tree:

- **Target**: Understand why some actions have large impacts under certain game contexts.
- **Method**: Mimic Decision Tree.
  1) Feed states and actions into a *CART* to fit the impact values via supervised learning.
  2) Compute the *feature importance* with the learned tree.
- Some results (**Top 10** important features for **shot** and **pass**):

Feature	Influence
X distance (t)	0.6632
outcome (t)	0.2275
Y distance (t)	0.0469
Game Time Remain (t)	0.0242
duration (t)	0.0062
X Coordinate (t-1)	0.0059
Game Time Remain (t-1)	0.0035
interrupted (t)	0.0035
X velocity (t)	0.0030
outcome (t-1)	0.0019

Table 5: Feature influence for the impact of shot.

Feature	Influence
X Velocity (t)	0.1355
Distance to Goal(t)	0.1264
Game Time Remain (t-1)	0.1082
Game Time Remain (t)	0.0816
Outcome (t)	0.0773
Outcome (t-1)	0.0760
Distance to Goal (t-1)	0.0411
Angle (t)	0.0373
Angle (t-1)	0.0298
X Velocity (t-1)	0.0174

Table 6: Feature influence for the impact of pass.

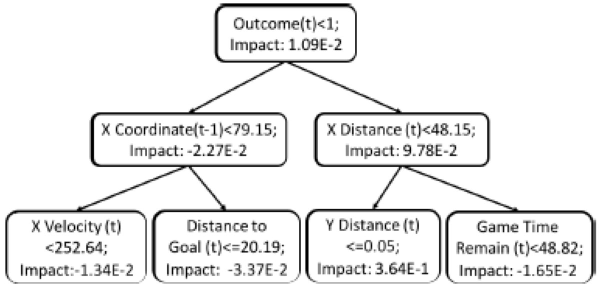

Fig. 6: Regression tree for the impact of shot.

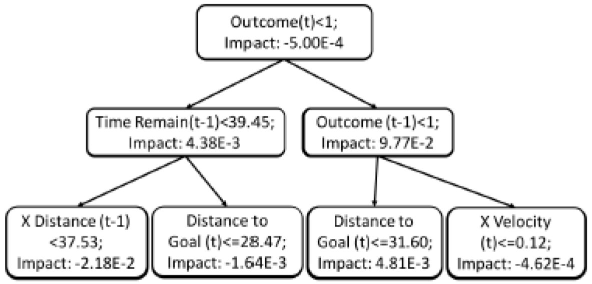

Fig. 7: Regression tree for the impact of pass.

Some findings:

- Shot impact significantly increases as a player *approaches the goal*.
- Passing impact increases with game velocity.

Player Ranking: Case Study

Fine-Tuning:

- **Motivation:** Different leagues have their competition level, season length, and playoff agenda.
- **Approaches:** (EFL Championship games)
  1) Train a general model to evaluate actions in European soccer.
  2) Fine-tune the weight values from the general model to a league specific model.

All-Actions Assessment:

name	team	GIM	Goals	Assists
Matej Vydra	Derby	18.017	21	4
Leon Clarke	Sheffield United	17.785	19	5
Lewis Grabban	Sunderland	16.045	12	0
Bobby De Cordova-Reid	Bristol	15.976	19	7
Diogo José Teixeira da Silva	Wolverhampton	15.707	17	5
Tom Cairney	Fulham	15.24	5	5
Ivan Cavaleiro	Wolverhampton	14.979	9	12
Stefan Johansen	Fulham	13.565	8	8
James Maddison	Norwich	13.23	14	8
Gary Hooper	Sheffield Wednesday	11.953	10	3

Table 7: 2017-2018 season top-10 Player Impact Scores for players in EFL Championship game season.

- *Matej Vydra tops* our 2017-2018 season ranking.

Action-Specific Assessment:

- Top shot players

name	GIM	Goal
Matej Vydra	4.747	21
Leon Clarke	4.024	19
Lewis Grabban	3.775	12
Kouassi Ryan Sessegnon	3.657	15
Harry Wilson	3.135	7
Famara Diedhiou	3.015	13
Sean Maguire	2.5	10
Joe Garner	2.44	10
Jarrod Bowen	2.408	14
Callum Paterson	2.29	10

Table 8: Top-10 players with largest shot impact in 2017-2018 EFL Championship game season.

- Top passing players

name	GIM	Assist
Leon Clarke	8.05	5
Matej Vydra	5.957	4
Bobby De Cordova-Reid	5.134	7
Chris Wood	4.732	1
Gary Hooper	4.694	3
Ivan Cavaleiro	4.533	12
Diogo José Teixeira da Silva	4.283	5
Gary Madine	4.202	2
Tom Cairney	4.123	5
Conor Hourihane	4.042	2

Table 9: Top-10 players with largest pass impact in 2017-2018 EFL Championship game season.

- Top shot players lead the goal scoring.
- Top passing players *may not have leading assists*.

Player Ranking: Empirical Evaluation

Comparison Player Evaluation Metrics:

- **Goal-based Metrics :**
  1) Plus- Minus (**PM**): measures how much the presence of a player influences the goals of his team.
  2) Expected Goal (**XG**): weights each shot by its chance of leading to a goal.
- **All-Action Metrics** :
  1) Valuing Actions by Estimating Probabilities **(VAEP)** applies the difference of action values to compute the impact of on-the-ball actions.
  2) Scoring Impact **(SI)**: based on a Markov model with pre-discretized spatial and temporal features.
  3) **M-GIM**: merges our home/away towers and fits all the states and actions with a single-layer network.

Correlations with Standard Success Measures (all players) :

Methods	Goals	Assists	SpG	PS%	KeyP	Yel	Red
PM	0.284	0.318	0.199	**0.288**	0.218	0.001	-0.069
VAEP	0.093	0.290	0.121	-0.111	0.116	0.024	0.133
XG	0.422	0.173	0.328	0.164	0.278	0.534	0.034
SI	0.585	0.153	0.438	-0.140	0.052	0.114	-0.089
M-GIM	0.648	0.367	0.573	0.153	0.417	-0.110	-0.145
GIM	**0.844**	**0.498**	**0.596**	0.16	**0.562**	-0.181	-0.137

- GIM achieves *promising correlation* with most success measures.
- Our model correctly recognizes that a penalty reduces the scoring probability, influencing the overall player GIM.

Player Ranking: Empirical Evaluation

Correlations with Standard Success Measures (EFL Championship players) :

Methods	Goals	Assists	SpG	PS%	KeyP	Yel	Red
PM	0.262	0.223	0.122	0.155	0.112	0.033	-0.046
VAEP	0.08	0.26	0.116	-0.126	0.137	-0.015	0.215
XG	0.420	0.165	0.394	0.149	0.254	0.578	-0.021
SI	0.574	0.124	0.408	-0.144	0.054	0.084	-0.147
M-GIM	0.629	0.309	0.551	**0.171**	0.388	-0.039	-0.132
GIM	0.638	0.382	0.553	-0.053	0.468	-0.026	-0.105
FT-GIM	**0.736**	**0.585**	**0.569**	0.082	**0.592**	-0.110	-0.171

- Championship League players' correlations generally decrease.
- it is more *severe* for our GIM metric.
- Fine-tuning (FT-GIM) addresses this issue.

Round-by-Round Correlations: Predicting Future From Past Performance :

- All players
- Players in the EFL Champion leagues

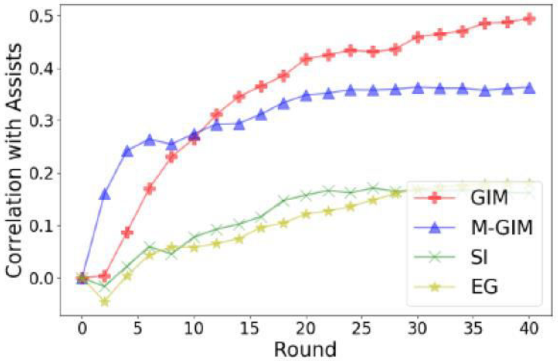

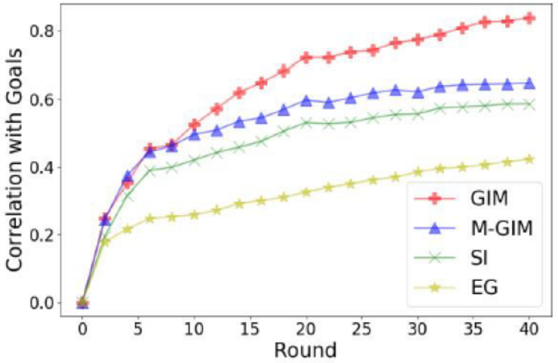

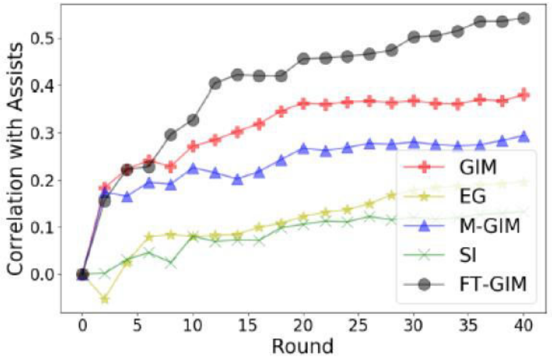

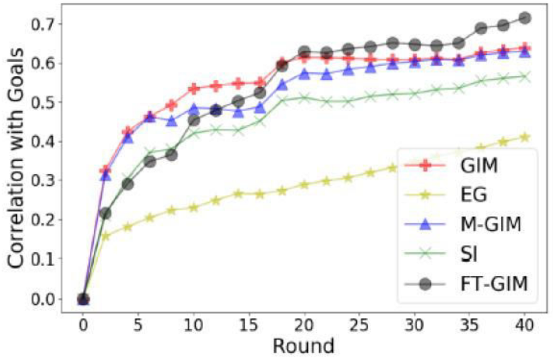

Conclusion

Key takeaways :

- **Q function from Sarsa Temporal Difference (TD) Learning:**
  1. Neural function approximator fits well with the high dimensional Spatial-temporal data.
  2. TD method provides a promising evaluation to action-state.
- **Action Impact :**
  1. Cancel influence from previous action and focus only on current actions.
  2. Effectively evaluate the influence of agent's action
- **Domain knowledge:**
  1. Home/away team behaves differently.
  2. Players in different soccer league should be evaluated separately.
  3. Action impact correlates well with standard success measures (e.g., goal, shot, etc.,)