Valuing Actions and Ranking Hockey Players With Machine Learning (Extended Abstract)

Oliver Schulte
School of Computing Science
Simon Fraser University
Burnaby, Canada

Linkökping Hockey Analytics Conference
May 2022

Abstract

A fundamental goal of sports analytics is to rank player performance. A common approach is to assign a value to each player action and rank a player by their aggregate action value. A recent AI-based approach is to measure the value of a player's action by how much it increases their team's chance of success, that is, their team's chance of scoring the next goal. This requires a model that outputs a success probability estimate, given a match context and an action. This talk describes machine learning techniques for building success probability models from data. The techniques range from easy-to-implement probabilistic classifiers to advanced reinforcement learning methods. The results of success probability models are illustrated with action values and player rankings for the National Hockey League.

1 Introduction: Success Probabilities in Sports Analytics

During a match, each action by a sports team is directed towards maximizing the chance of future success. Therefore the *probability of (future) success* is a key statistical quantity for evaluating the strength of a team, the impact of an action, and the contributions of a player. This note accompanies a talk that shows how success probabilities can be used to value actions and ranking players, and describes techniques for estimating them from data.

2 Defining Success Probabilities

For the purposes of this note, success is flexibly defined as a binary event that a team seeks to bring about. An analyst is free to define success in different ways depending on the question they are investigating. Examples of success concepts that have appeared in the literature include the following (using hockey examples where possible).

- Winning the Match [Pettigrew, 2015].
- Scoring a goal within a short time interval (e.g. 1 minutes) [Schuckers and Curro, 2013].
- Scoring a goal within the next 5 actions [Decroos et al., 2019].
- Scoring the next goal in the match [Routley and Schulte, 2015].
- Drawing the next penalty [Routley and Schulte, 2015]. This is a failure event that would be interesting to a coach who is concerned to minimize the number of penalties incurred by their team.

The notation

$$P(S_i|\mathbf{X}_t)$$

denotes the probability that i achieves (future) success given the current match context $\mathbf{X}_t$. We discuss in Section 5.1 below how a context vector can be computed from play-by-play data.

A success probability is a dynamic quantity; a success probability ticker shows an estimated probability for each time in a match [Liu and Schulte, 2018]; see Figure 1.

3 From Success Probabilities to Action Values

From success probabilities we can assign a value to actions called the **impact** of an action occurring at time $t+1$ [Liu and Schulte, 2018].

$$impact_i(t+1) \equiv P(S_i|\mathbf{X}_{t+1}) - P(S_i|\mathbf{X}_t)$$

Thus the impact of an action is the difference in success probabilities before and after the action occurred. Figure 3 shows boxplots for impact values. Note that impact values can vary widely for the same action, depending on context.

4 From Action Values to Player Ranking

We can compute a player performance metric from impact values in a straightforward way: for each player, and for each of their actions, we can compute the impact of the action. The **goal impact metric** (GIM) is simply the total impact of the player's actions.

Our metric can be used to *identify undervalued players*. For instance, Johnny Gaudreau and Mark Scheifele drew salaries below what their GIM rank would suggest. Later they received a $\$5M+$ contract for the 2016-17 season.

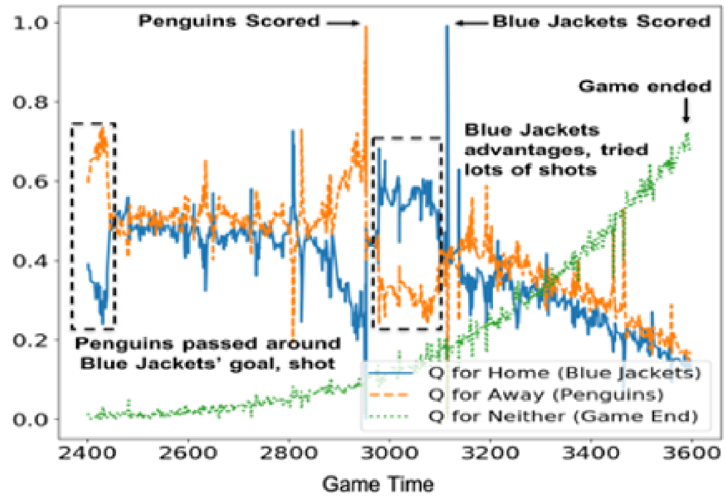

Figure 1: A success probability ticker for a match between the Penguins and the Blue Jacket. The y-axis shows the estimated probability of scoring the next goal.

While we do not have ground truth for evaluating player rankings, the goal impact player rankings have been validated indirectly in several ways.

1. GIM correlates well with standard success metrics (e.g., Points) [Liu and Schulte, 2018].
2. GIM converges close to half-way through the season. This means that the beginning of the season can be used to evaluate player strength (predict the player's final ranking).
3. The GIM values per player correlate well across different seasons [Routley, 2015, Pettigrew, 2015], which is evidence that they measure a stable quality of players.

The impact metric approach has also been validated in other sports, such as soccer [Decroos et al., 2019, Liu et al., 2020a, Fernández et al., 2021] and basketball [Cervone et al., 2014].

5 Learning Success Probability Models

Given the usefulness of success probability models, a major concern of machine learning for sports analytics is to develop machines for building such models from data. In the following I will discuss methods for building success probability models from event data, also known as play-by-play data. Estimating success probabilities from tracking data is less studied because tracking data is less

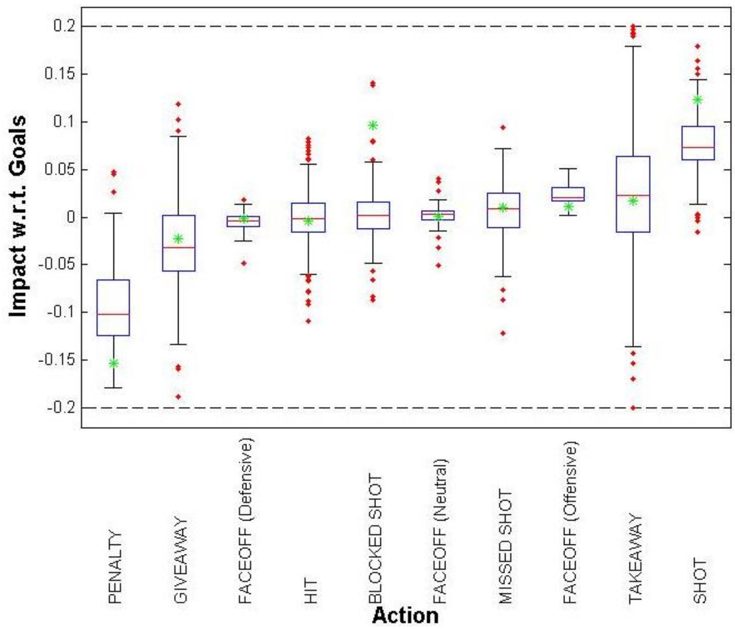

Figure 2: Impact on the probability of Scoring the Next Goal. Higher numbers are *better* for the team that performs the action. Action Impact Values vary with context. The central mark is the median, the edges of the box are the 25th and 75th percentiles. The whiskers are at the default value, approximately 2.7 s.d. Based on the model of [Routley and Schulte, 2015].

commonly available (but see [Dick and Brefeld, 2019, Fernández et al., 2021]). Table 2 illustrates play-by-play data.

5.1 Classifier Approach

A straightforward approach is to annotate each event at time t with a binary target $Y_{i,t} \in \{0, 1\}$ that denotes whether team i acting at time t achieved future success after time t. For example in the game of Figure 1, the Penguins scored around time $t = 3,900$ sec. So for all previous times $2,400 < t' < 3,900$, we would have $Y_{Penguins,t'} = 1$ and $Y_{Flyers,t'} = 0$. Then estimating success probabilities can be modelled as predicting a binary label given the information $\mathbf{X}_t$ available at time t.

It is straightforward to include in the context vector $\mathbf{X}_t$ values for time indexed features score differential, manpower differential, time remaining, location etc. [Routley and Schulte, 2015, Liu et al., 2018]. The main difficulty is how to include the match history prior to time t. A simple approach is to fix a window size k and then to append to $\mathbf{X}_t$ the time-indexed features for the previous times $t-1, \ldots, t-k$. See [Decroos et al., 2019] for a model of this approach applied

Table 1: 2015-2016 Top-20 Player Impact Scores. Based on the model of [Liu and Schulte, 2018].

Name	Impact	Assists	Goals	Points	+/-	Salary
Taylor Hall	96.40	39	26	65	-4	$6,000,000
Joe Pavelski	94.56	40	38	78	25	$6,000,000
Johnny Gaudreau	94.51	48	30	78	4	$925,000
Anze Kopitar	94.10	49	25	74	34	$7,700,000
Erik Karlsson	92.41	66	16	82	-2	$7,000,000
Patrice Bergeron	92.06	36	32	68	12	$8,750,000
Mark Scheifele	90.67	32	29	61	16	$832,500
Sidney Crosby	90.21	49	36	85	19	$12,000,000
Claude Giroux	89.64	45	22	67	-8	$9,000,000
Dustin Byfuglien	89.46	34	19	53	4	$6,000,000
Jamie Benn	88.38	48	41	89	7	$5,750,000
Patrick Kane	87.81	60	46	106	17	$13,800,000
Mark Stone	86.42	38	23	61	-4	$2,250,000
Blake Wheeler	85.83	52	26	78	8	$5,800,000
Tyler Toffoli	83.25	27	31	58	35	$2,600,000
Charlie Coyle	81.50	21	21	42	1	$1,900,000
Tyson Barrie	81.46	36	13	49	-16	$3,200,000
Jonathan Toews	80.92	30	28	58	16	$13,800,000
Sean Monahan	80.92	36	27	63	-6	$925,000
Vladimir Tarasenko	80.68	34	40	74	7	$8,000,000

to soccer (they used $k = 3$ as the window size). After extracting the window information as context, the data will be a list of $\langle \mathbf{X}_t, Y_{i,t} \rangle$ pairs, which is the standard format for any classifier package available in systems like R, Weka, scikit-learn.

An alternative to using a fixed window size is to apply a recurrent neural network, which can take as input a sequence without the need for preprocessing.

5.2 Reinforcement Learning

Reinforcement learning (RL) is the branch of machine learning that studies learning to act [Sutton and McCallum, 2007]. Estimating success probabilities from sequential data is one of the basic well-studied problems in RL. In RL, a mapping from match states to success probabilities is known as a **value function** and estimating a value function is called the **prediction problem**. The classifier approach described in the previous subsection (implicitly) treats all match states as independent, and hence ignores the correlations between success probabilities due to the temporal dynamics of ice hockey. In contrast, reinforcement learning seeks to exploit the temporal dynamics to efficiently learn success probabilities.

If we discretize the spatial rink coordinates, we can model hockey dynamics

Table 2: Sample Play-By-Play Data in Tabular Format.

gameId	playerId	period	teamId	xCoord	yCoord	Manpower	Action Type
849	402	1	15	-9.5	1.5	even	lpr
849	402	1	15	-24.5	-17	even	carry
849	417	1	16	-75.5	-21.5	even	check
849	402	1	15	-79	-19.5	even	puckprot.
849	413	1	16	-92	-32.5	even	lpr
849	413	1	16	-92	-32.5	even	pass
849	389	1	15	-70	42	even	block
849	389	1	15	-70	42	even	lpr
849	389	1	15	-70	42	even	pass
849	425	1	16	-91	34	even	block
849	395	1	15	-97	23.5	even	reception

in a framework known as a **discrete Markov decision process** [Routley and Schulte, 2015, Schulte et al., 2017a,b]. The key parameters in a Markov decision process are *state transition probabilities* that describe what is likely to happen next in a hockey game. Given an estimate of state transition problems, the dynamic programming algorithm can be used to compute success probabilities for any match state.

While discretization can simplify learning and in many cases increases the interpretability of success probabilities, it also loses information and introduces unnatural discontinuities in a success probability model. Reinforcement learning provides so-called model-free methods for learning success probabilities that do not require discrete state transition probabilities. Combining model-free methods with neural networks provides a method for learning success probabilities that can take as input continuous spatio-temporal data "as is" without the need for discretization or fixing a window size. Model-free deep RL has been developed in several recent approaches for sports dynamics [Liu et al., 2018, 2020b,a]. Figure 3 summarizes the options for learning success probabilities discussed.

6 Conclusion

Estimating success probabilities is a basic statistical problem in hockey analytics. A good success probability model can be leveraged to solve important analytics problems such as quantifying the value of an action and the contributions of a player. Machine learning models can include rich match contexts to provide useful success probabilities. Probabilistic classifiers based on a sliding window are relatively straightforward to implement and can serve as a strong baseline for evaluating the usefulness of success probabilities in an application. Reinforcement learning is especially suitable for handing complex dynamic domains like ice hockey and provides a powerful set of tools for increasing the complexity and accuracy of a hockey model.

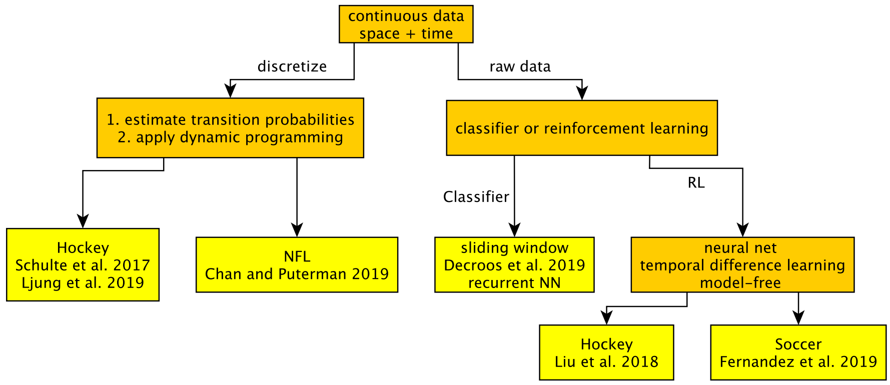

Figure 3: Approaches for Learning Success Probabilities

References

Stephen Pettigrew. Assessing the offensive productivity of nhl players using in-game win probabilities. In *9th Annual MIT Sloan Sports Analytics Conference*, 2015.

Michael Schuckers and James Curro. Total hockey rating (thor): A comprehensive statistical rating of national hockey league forwards and defensemen based upon all on-ice events. In *MIT sloan Sports Analytics Conference*, 2013.

Tom Decroos, Lotte Bransen, Jan Van Haaren, and Jesse Davis. Actions speak louder than goals: Valuing player actions in soccer. In *Proceedings of the 25th International Conference on Knowledge Discovery & Data Mining (KDD-19)*, pages 1851–1861, 2019.

Kurt Routley and Oliver Schulte. A markov game model for valuing player actions in ice hockey. In *Uncertainty in Artificial Intelligence (UAI)*, pages 782–791, 2015.

Guiliang Liu and Oliver Schulte. Deep reinforcement learning in ice hockey for context-aware player evaluation. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 3442–3448. International Joint Conferences on Artificial Intelligence Organization, 7 2018. doi: 10.24963/ijcai.2018/478. URL https://doi.org/10.24963/ijcai.2018/478.

Kurt Routley. A markov game model for valuing player actions in ice hockey. Master's thesis, Simon Fraser University, April 2015.

Guiliang Liu, Yudong Luo, Oliver Schulte, and Tarak Kharrat. Deep soccer analytics: learning an action-value function for evaluating soccer players. *Data Min. Knowl. Discov.*, 34(5):1531–1559, 2020a.

doi: 10.1007/s10618-020-00705-9. URL https://doi.org/10.1007/s10618-020-00705-9.

Javier Fernández, Luke Bornn, and Daniel Cervone. A framework for the fine-grained evaluation of the instantaneous expected value of soccer possessions. *Machine Learning*, 110(6):1389–1427, 2021.

Dan Cervone, Alexander D'Amour, Luke Bornn, and Kirk Goldsberry. Pointwise: Predicting points and valuing decisions in real time with NBA optical tracking data. In *8th Annual MIT Sloan Sports Analytics Conference, February*, volume 28, 2014.

Uwe Dick and Ulf Brefeld. Learning to rate player positioning in soccer. *Big data*, 7(1):71–82, 2019.

Guiliang Liu, Oliver Schulte, Wang Zhu, and Qingcan Li. Toward interpretable deep reinforcement learning with linear model u-trees. In *Proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases (ECML)*, volume abs/1807.05887, pages 1–16, 2018. URL http://arxiv.org/abs/1807.05887.

Charles Sutton and Andrew McCallum. An introduction to conditional random fields for relational learning. In *Introduction to Statistical Relational Learning*, chapter 4, pages 93–127. MIT Press, 2007.

Oliver Schulte, Zeyu Zhao, Mehrsan Javan, and Philippe Desaulniers. Apples-to-apples: Clustering and ranking nhl players using location information and scoring impact. In *MIT Sloan Sports Analytics Conference*, 2017a.

Oliver Schulte, Mahmoud Khademi, Sajjad Gholami, Zeyu Zhao, Mehrsan Javan, and Philippe Desaulniers. A Markov game model for valuing actions, locations, and team performance in ice hockey. *Data Mining and Knowledge Discovery*, pages 1–23, 2017b. ISSN 1573-756X. doi: 10.1007/s10618-017-0496-z. URL http://dx.doi.org/10.1007/s10618-017-0496-z.

Guiliang Liu, Oliver Schulte, Pascal Poupart, Mike Rudd, and Mehrsan Javan. Learning agent representations for ice hockey. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2020, December 6-12 2020*, 2020b.