

Player pairs valuation in ice hockey

Dennis Ljung Niklas Carlsson Patrick Lambrich

Linköping University, Sweden

Abstract. To overcome the shortcomings of simple metrics for evaluating player performance, recent works have introduced more advanced metrics that take into account the context of the players' actions and perform look-ahead. However, as ice hockey is a team sport, knowing about individual ratings is not enough and coaches want to identify players that play particularly well together. In this paper we therefore extend earlier work for evaluating the performance of players to the related problem of evaluating the performance of player pairs. We experiment with data from seven NHL seasons, discuss the top pairs, and present analyses and insights based on both the absolute and relative ice time together.

Keywords: Sports analytics · Data mining · Player valuation.

1 Introduction

In the field of sports analytics, many works focus on evaluating the performance of players. A commonly used method to do this is to attribute values to the different actions that players perform and sum up these values every time a player performs these actions. These summary statistics can be computed over, for instance, games or seasons. In ice hockey, common summary metrics include the number of goals, assists, points (assists + goals) and the plus-minus statistics (+/-), in which 1 is added when the player is on the ice when the player's team scores (during even strength play) and 1 is subtracted when the opposing team scores (during even strength). More advanced measures are, for instance, Corsi and Fenwick¹.

However, these metrics do not capture the context of player actions and the impact they have on the outcome of later actions. To address this shortcoming and to capture the ripple effect of actions (where one action increases/decreases the success of a later action, for example), recent works [11,13,6] have therefore introduced more advanced metrics that take into account the context of the actions and perform look-ahead. The use of look-ahead is particularly valuable in low-scoring sports such as ice hockey.

An important aspect that the above works do not take into account is that ice hockey is a team sport. In particular, individual ratings are not enough to predict outcomes for the team. It is therefore important for coaches to identify players that play particularly well together. In this paper, we therefore extend the above

¹ See, e.g., [https://en.wikipedia.org/wiki/Analytics_\(ice_hockey\)](https://en.wikipedia.org/wiki/Analytics_(ice_hockey)).

recent analysis approach to the related problem of evaluating the performance of **pairs of players**. For our analysis, we extend the work by Routley and Schulte [11] to evaluate the performance of player pairs. More specifically, we use their action-value Q-function to assign values to individual actions performed by players and then sum the value-adding actions associated with player pairs when they are on the ice simultaneously.

In ice hockey there are usually two defenders, three forwards, and a goaltender on the ice simultaneously. However, the set of players changes frequently (e.g., the average shift duration is roughly 45 seconds) and coaches typically select to adjust which players are on the ice at each point in time based on desirable matchups against the other teams and, perhaps most importantly, based on which players play well together. While ice hockey coaches traditionally talk about defense pairings (consisting of two defenders) and forward lines (consisting of three forwards), coaches increasingly work with both defense pairings and forward pairings. Identifying forwards pairs with particularly good "chemistry" is considered simpler and allows some flexibility when there are injuries on a team, for example. By focusing on pairings, we provide a tool that helps identify player pairs that perform particularly good/bad.

For our analysis, we use this tool to identify particularly successful pairings, compare successful pairings to how their coaches assign ice time to these pairings, and to compare the success of different categories of player pairs (e.g., based on position, total ice time together, or the relative fraction of their ice time played together). At a high level we find that coaches' desire to play their top players against the other teams' top players appears to even out the relative impact per minute observed across different player categories.

The remainder of the paper is organized as follows. Section 2 discusses related work, and Section 3 describes the work of [11]. In Section 4 we define our metrics to evaluate pairs of players and in Section 5 we present and discuss the results of our experiments with NHL play-by-play data from the 2007-2008 through 2013-2014 NHL seasons, as provided by [11]. The paper concludes in Section 6.

2 Related work

Regarding player performance in ice hockey, several regression models have been proposed for dealing with the weaknesses of the +/- measure (e.g., [7,8,4]). Other measures have also been introduced, including Corsi, Fenwick, and added goal value [10].

Another measure for player evaluation based on the events that happen when a player is on the ice is proposed in [12]. Event impacts are based on the probability that the event leads to a goal (for or against) in the next 20 seconds.

Other works model the dynamics of an ice hockey game using Markov games where two opposing sides (e.g., the home team and the away team) try to reach states in which they are rewarded (e.g., scoring a goal). In [15] the scoring rate for each team is modeled as a semi-Markov process, with hazard functions for each process that depend on the players on the ice. A Markov win probability

model given the goal and manpower differential state at any point in a hockey game is proposed in [5]. In [11,13,14,6] action-value Q-functions are learned with respect to different targets. (See Section 3 for the model in [11].) Although the approaches use Markov-based approaches, the definitions of states and reward functions are different. The advantages of such approaches (e.g., [14]) are the ability to capture game context (goals have different values in a tie game than in a game where a team is leading with many goals), the ability to look ahead and thereby assigning values to all actions in the game, and the possibility to define a player’s impact through the player’s actions. In this paper we base our work on one of these approaches, i.e., the work in [11].

There is not much work on evaluating player pair performance for ice hockey. In [15] player pairs are rated using their Markov-based model of the game. They selected the 1,000 player pairs with respect to the number of full-strength shifts in five NHL seasons (2007-2008 until 2011-2012) where the players in the pairs are both forwards or both defenders. Lessons learned consisted of knowledge about player pairs where the players performed better than their individual performance, and player pairs where the combination reduced the performance. Further, there is some work in basketball, such as [1,2] where player pairs are ranked according to a +/- measure. Which types of players should be chosen for the top pair of players, is investigated in [3].

3 Background

In this section we explain the model of [11]. In [11] action-value Q-functions are learned with respect to the next goal or the next penalty. In this paper we only use the next goal Q-function.

The state space considers *action events* with three parameters: the action type (Faceoff, Shot, Missed Shot, Blocked Shot, Takeaway, Giveaway, Hit, Goal), the team that performs the action, and the zone (offensive, neutral, defensive).

A *play sequence* is defined as the empty sequence or a sequence of events for which the first event is a start marker, the possible next events are action events, and the possible last event is an end event. If the play sequence ends with an end event, it is a *complete* play sequence. The start/end events are Period Start, Period End, Early Intermission Start, Penalty, Stoppage, Shootout Completed, Game End, Game Off, and Early Intermission End.

Actions and play sequences occur in a context. In [11] a *context state* contains values for 3 context features. Goal Differential is the number of home goals minus the number of away goals. Manpower Differential is the number of home players on the ice minus the number of away players on the ice. Further, the Period of the game is recorded.

A *state* is then a pair which contains a context state and a (not necessarily complete) play sequence.

Actions are performed in specific states. For action a and state $s = \langle c, ps \rangle$, where c is the context state and ps is the play sequence, the resulting state of performing a in state s is denoted by $s * a$ and is defined as $\langle c, ps * a \rangle$, where

$ps * a$ is the play sequence obtained by appending action a to ps . For states with play sequences that are end events, the next state is a state of the form $\langle c', \emptyset \rangle$ where c' is defined by the end event. For instance, a goal will change the goal differential and update the context.

Transition probabilities between different states are based on play-by-play data. The transition probability $TP(s, s')$ for a transition from state s to state s' is defined as $Occ(s, s') / Occ(s)$ where $Occ(s)$ is the number of occurrences of s in the play-by-play data and $Occ(s, s')$ is the number of occurrences of s that are immediately followed by s' in the play-by-play data.

Using a state transition graph with the computed transition probabilities, Q-values for states are learned using a value iteration algorithm. The impact of an action in a certain state $impact(s, a)$ is then defined as $Q_T(s * a) - Q_T(s)$ where T is the team performing the action.

The performance of a player is computed as the sum of the impacts of the actions the player performs (over a game or a season). This is equivalent to comparing the actions taken by a specific player to the actions of an average player.

For our work, we reimplemented the code available from [11] using Python and C++. The reward for goals for is +1 and goals against -1. The resulting impact values for the actions were used as a base for our work on pair performance.

4 Player pair metrics

We base our method for computing performance measures on the impact of actions² as defined by [11]. However, we define the impact of players in different ways. First, we define different sets of actions for players and player pairs (Table 1). We differentiate between actions performed by a player and actions performed (by the player or another player) when a player is on the ice. The player impact is then defined using the actions when the player is on the ice (Table 2). This allows for a measure that includes indirect impact on the game by being on the ice. Even when players do not perform registered actions, they can still influence the game; e.g., by opening up a path for a teammate who may score. Further, we define the direct impact of a player based on the actions the player performs (and this is essentially the impact as defined in [11]).

For player pairs we define the impact using the actions when both players are on the ice. To be able to measure the influence of the players in the pair on each other, we also define the impact of a player without a particular second player, i.e., we use the actions when the first player is on the ice, but not the second.

5 Data-driven analysis

For our analysis, we used NHL play-by-play data from the 2007-2008 through 2013-2014 NHL season, as provided by [11]. To compare against prior works,

² In the remainder we use action as a shorthand for action in a particular state.

Table 1. Basic action sets.

A is the set of all state-action-pairs $\langle s, a \rangle$ where action a is performed in state s
$A_i(p_k)$ is the set of state-action-pairs when player p_k is on the ice
$A_i(p_k, p_l)$ is the set of state-action-pairs when players p_k and p_l both are on the ice
$A_i(p_k, p_l) = A_i(p_k) \cap A_i(p_l)$
$A_p(p_k)$ is the set of state-action-pairs where the action is performed by player p_k
$A_p(p_k) \subseteq A_i(p_k)$

Table 2. Player and player pair impact.

The direct impact of a player is the sum of the impact values of the actions performed by the player: D-impact(p_k) = $\sum_{\langle s, a \rangle \in A_p(p_k)} \text{impact}(s, a)$
The impact of a player is the sum of the impact values of the actions when the player is on the ice: impact(p_k) = $\sum_{\langle s, a \rangle \in A_i(p_k)} \text{impact}(s, a)$
The impact of a player pair is the sum of the impact values of the actions when both players are on the ice: impact(p_k, p_l) = $\sum_{\langle s, a \rangle \in A_i(p_k, p_l)} \text{impact}(s, a)$
The impact of a player without a second player is the sum of the impact values of the actions when the player is on the ice and the second player is not on the ice: impact-without(p_k, p_l) = $\sum_{\langle s, a \rangle \in (A_i(p_k) \setminus A_i(p_k, p_l))} \text{impact}(s, a)$

we primarily focus our analysis on the last two full regular seasons in this set; i.e., the 2011-2012 season and the 2013-2014 season. (The 2012-2013 season was shortened due to a lockout.)

5.1 Top pairings

We first present the top pairings according to the impact metrics, as calculated over the entire seasons, for three categories of pairings: forward pairs, defense pairs, and mixed pairs (consisting of a forward and a defender). Tables 3 and 4 summarize these results. Here, we include the player position (defender (D) or the three forward positions: rightwing (R), leftwing (L), and center (C)) and their high-level stats over the entire season (goals (G), assists (A), and plus-minus (+/-)), together with the team, the pairs' total impact (rounded), and the pairs' joint time on ice (TOI), measured in seconds.

Looking first at the 2011-2012 result, we note that many of the names on this list placed high in the scoring race (e.g., Stamkos 2nd, Spezza 4th, Kovalchuk 5th) or were responsible for a large fraction of their teams scoring (e.g., Pavelski/Thornton and O'Reilly/Landeskog). They also all were among the most relied pairings and hence accumulated among the most ice time among all forward pairs. For example, Kovalchuk/Parise, Pavelski/Thornton, and O'Reilly/Landeskog all placed in the top-5 in joint ice time among forward pairs. For the defense pairs and mixed pairs, the correlation was even greater, with four out of five in the corresponding top-five TOI sets.

Table 3. Top pairs 2011-2012 according to total impact.

	Player 1					Player 2					Pair stats		
	Name	Pos	G	A	+/-	Name	Pos	G	A	+/-	Team	Impact	TOI
Forwards	Ilya Kovalchuk	R	37	46	-9	Zach Parise	L	31	38	-5	NJD	121.17	40,163
	Ryan O'Reilly	C	18	37	-1	Gabriel Landeskog	L	22	30	+20	COL	115.74	39,021
	Joe Pavelski	C	31	30	+18	Joe Thornton	C	18	59	+17	SJS	112.65	39,353
	Steven Stamkos	C	60	37	+7	Martin St. Louis	R	25	49	-3	TBL	111.77	35,941
	Milan Michalek	L	35	25	+4	Jason Spezza	C	34	50	+11	OTT	111.73	36,689
Defenders	Dan Girardi	D	5	24	+13	Ryan McDonagh	D	7	25	+25	NYR	155.28	55,911
	Filip Kuba	D	6	26	+26	Erik Karlsson	D	19	59	+16	OTT	134.74	47,985
	Francois Beauchemin	D	8	14	-14	Cam Fowler	D	5	24	-28	ANA	125.54	45,795
	Josh Gorges	D	2	14	+14	P.K. Subban	D	7	29	+9	MTL	125.16	44,390
	Carl Gunnarsson	D	4	15	-9	Dion Phaneuf	D	12	32	-10	TOR	123.06	36,181
Mixed	Jason Spezza	C	34	50	+11	Erik Karlsson	D	19	59	+16	OTT	110.58	35,990
	Joe Pavelski	C	31	30	+18	Dan Boyle	D	9	39	+10	SJS	106.04	35,612
	Joe Thornton	C	18	59	+17	Dan Boyle	D	9	39	+10	SJS	102.96	35,160
	Tomas Fleischmann	L	27	34	-7	Brian Campbell	D	4	49	-9	FLA	98.08	31,804
	Stephen Weiss	C	20	27	+5	Brian Campbell	D	4	49	-9	FLA	96.79	32,995

Table 4. Top pairs 2013-2014 according to total impact.

	Player 1					Player 2					Pair stats		
	Name	Pos	G	A	+/-	Name	Pos	G	A	+/-	Team	Impact	TOI
Forwards	James van Riemsdyk	L	30	31	-9	Phil Kessel	R	27	43	-5	TOR	166.42	51,910
	Alex Ovechkin	L	21	38	-35	Nicklas Backstrom	C	18	61	-20	WSH	113.85	40,815
	James van Riemsdyk	L	30	31	-9	Tyler Bozak	C	19	30	+2	TOR	111.39	32,567
	Phil Kessel	R	27	43	-5	Tyler Bozak	C	19	30	+2	TOR	109.57	34,648
	Chris Kunitz	L	35	33	+25	Sidney Crosby	C	36	68	+18	PIT	107.04	36,296
Defenders	Duncan Keith	D	6	55	+22	Brent Seabrook	D	7	34	23	CHI	143.19	46,762
	Dan Girardi	D	5	19	+6	Ryan McDonagh	D	14	29	+11	NYR	129.03	50,102
	Carl Gunnarsson	D	3	14	+12	Dion Phaneuf	D	8	23	31	TOR	124.39	41,219
	Justin Faulk	D	5	27	-9	Andrej Sekera	D	11	33	+4	CAR	121.08	42,966
	Jan Hejda	D	6	11	+8	Erik Johnson	D	9	30	+5	COL	116.57	40,941
Mixed	Phil Kessel	R	27	43	-5	Dion Phaneuf	D	8	23	31	TOR	106.39	32,608
	James van Riemsdyk	L	30	31	-9	Dion Phaneuf	D	8	23	31	TOR	100.82	30,847
	Jason Spezza	C	23	43	-26	Erik Karlsson	D	20	54	-15	OTT	92.91	31,898
	Sean Couturier	C	13	26	+1	Braydon Coburn	D	5	12	-6	PHI	90.43	26,893
	Anze Kopitar	C	29	41	+34	Drew Doughty	D	10	27	+17	LAK	88.74	33,209




Fig. 1. CDFs and CCDFs of the shared TOI for different NHL seasons.

The 2013-2014 results are similar in that the top-5 list of forward pairs include three of the top-ten names in the scoring race (e.g., Crosby 1st, Kessel 6th, and Ovechkin 8th) and responsible for a large portion of the points on their respective teams, including the line consisting of the three Toronto (TOR) players van Riemsdyk, Kessel, and Bozak. Despite this, none of these three players are still with Toronto, as Kessel was traded to Pittsburgh (PIT) in 2015 and van Riemsdyk and Bozak signed with Philadelphia (PHI) and St. Louis (STL), respectively, the first week of July 2018, illustrating how quickly a team can change direction and build. The other pairings on this list have since combined for three Stanley Cups (PIT in 2015-2016 and 2016-2017 and WSH in 2017-2018), with each of these players contributing to the championships. In fact, Kessel (TOR above) was part of both Pittsburgh (PIT) championship teams. Interestingly, on a related note, the top defense pairings above won the Stanley Cup the 2013-2014 season, as part of the Chicago (CHI) championship team. Keith also won the Norris trophy, as the top defensemen in the league. Other recent Norris trophy winners show up in the top-5 list of the mixed category, including Karlsson (2011-2012 and 2014-2015) and Doughty (2015-2016).

5.2 TOI-based analysis

Coaches typically carefully select player combinations and match these against the opponents' lineups so to maximize the chance of success. This means that good players and player combinations typically get more ice time, but also that they may be matched up against tougher competition. We next look at the relationship between TOI and the impact per minute played together for each pair with at least one minute played together during at least one game of the season.

To help interpret the observed relationships, we first present Figure 1. Here, we plot the cumulative distribution functions (CDFs) and complimentary CDFs (CCDFs) for the joint TOI across all player pairs meeting our threshold criteria, for each of the seasons 2007-2008 through 2013-2014, with the y-axes plotted on linear and logarithmic scales, respectively. First, note that the s-shaped CDFs (plotted on lin-log scale) display close to straight-line behavior on lin-log scale, suggesting that the distribution only has slightly heavier tail than the exponential distribution (for which we would expect straight line behavior) [9].





Fig. 2. Impact per minute played together, as a function of the joint TOI.

Second, we note that with exception of the 2012-2013 season (which only had 48 games per team, rather than the typical 82 games per team, due to a player strike/league lockout), the distributions are relatively overlapping, suggesting that this result may be invariant across seasons.

Now, let us get back to the relationship between impact per minute and TOI. Figure 2(a) plots distribution statistics for the impact per minute observed across all player pairs with a joint TOI falling into the time interval $[2^i, 2^{i+1})$, during the 2011-2012 and 2013-2014 seasons, where we vary i from 0 to 9. Here, the body of the bars shows the 75%-iles, the whisker lines the 90%-iles, and the markers (\times) the medians. We also include the (almost overlapping) overall medians across all pairs (0.192 and 0.190, respectively). As expected, the variation decreases the more time that pairs play together. However, it is interesting to note that the medians are highest for the pairs that play 16-256 minutes together, and actually decrease for some of the top TOI pairings. This may partially be due to some pairings (especially on weaker teams) having to be relied upon more than perhaps is healthy for their performance. However, it is also an indication that coaches rely on some of these pairings to play “tough minutes”, against the other teams’ top players. Yet, the relatively flat median values and overall higher values for pairings that play significant minutes together suggest that coaches overall do a good job distributing the load and/or that more short-term pairings (due to overlapping shifts or bad line changes, where some players are caught out on the ice tired, for example) may perform worse.

When breaking down the analysis based on the player positions of the players in each pair, we have only observed small variations across the traditional combination types: forward pairs and defense pairs. For example, Figure 2(b) shows relatively similar normalized impact (per minute) for defense pairs and forward pairs, with similar joint TOI. Compared with these categories, the forward-defense pairs have higher impact per pair and contribute with the more pairings (e.g., 7,586 pairs during 2013/14, compared to 4,743 forward and 1,316 defense pairs for the same season). While part of the high-scoring pairings can be attributed to good forwards (e.g., Spezza and Kopitar) being matched with good puck moving defensemen (e.g., Karlsson and Doughty), further analysis of what makes good forward-defense pairings leaves room for interesting future work.



(a) Impact per minute played (b) Relative impact together vs not

Fig. 3. Impact metrics as a function of fraction of time played together.

5.3 Relative ice time together

We note that coaches also observe players during practice and off the ice, where part of the chemistry between two players may be developed. It is therefore interesting to analyze if players that spend most of a game together in fact produce better during the time they play together or when they play with other players. As a first-cut analysis to look into this question, we plot the impact per minute as function of the time the players play together during the games that they played at least one minute together (Figure 3(a)) as well as the ratio between the impact per minute when playing together during those games and when not playing together during those same games (Figure 3(b)). Here, we use an exponentially moving weighted average (EWMA) with $\alpha = 0.02$ to smooth out the curves. Despite using significant smoothing, the variations are significant compared to the relative trends, making it difficult to identify clear patterns. However, in general, it is interesting to see that the players that spend the largest fraction together often have lower relative impact when playing together. At first this may appear counterintuitive. However, for players that also have significant playing time in those games (e.g., subset of the players with at least 300 minutes together), this may be due to matchups against the other teams' top lines. In other cases, this may be an effect of fourth line players taking their opportunities when on the ice with top line players, with whom they spend less of their total TOI with.

6 Conclusion

In this paper we extend a recent analysis approach for evaluating the performance of players [11] to the related problem of evaluating the performance of pairs of players. In particular, we defined measures for player pairs' impact and analyzed NHL play-by-play data from the 2007-2008 through 2013-2014 NHL seasons using these new metrics.

Our analysis helps identify pairings that have particularly good "chemistry", that performed well across the season (e.g., top pairings), or that the coaches for other reasons rely more heavily on (e.g., that may have played long and tough minutes against other teams' top players). Some of the lessons learned are that

for the top pairings, according to the impact metrics, many of the names on this list placed high in the scoring race or were responsible for a large fraction of their teams scoring. Further, forward-defense pairs have higher impact per pair, and the players that spend the largest fraction together often have lower relative impact when playing together. Using the data, we also hypothesize that coaches desire to play their top players against the other teams' top players appears to even out the relative impact per minute observed across different player categories.

Regarding future work, one direction is to work with different reward functions in the Q-learning algorithm to investigate impact of player actions for different desirable outcomes (e.g., shots on goals, powerplays). We also intend to investigate alternative pair impact definitions. For instance, the current definition credits the pair for the actions when they are on the ice (indirect impact), while it would be interesting to compare to direct impact as well (e.g., the pair receives credit only for direct impact actions by one of the players in the pair).

References

1. 82games: Which players play well together? - 2003-04 season, <http://www.82games.com/ppairs0304.htm>
2. 82games: Which players play well together? - 2004-05 season, <http://www.82games.com/ppairs0405.htm>
3. Ayer, R.: Big 2's and Big 3's: Analyzing how a team's best players complement each other. In: MIT Sloan Sports Analytics Conference (2012)
4. Gramacy, R.B., Jensen, S.T., Taddy, M.: Estimating player contribution in hockey with regularized logistic regression. *Journal of Quantitative Analysis in Sports* **9**, 97–111 (2013), doi: 10.1515/jqas-2012-0001
5. Kaplan, E.H., Mongeon, K., Ryan, J.T.: A Markov Model for Hockey: Manpower Differential and Win Probability Added. *INFOR Information Systems and Operational Research* **52**(2), 39–50 (2014), doi: 10.3138/infor.52.2.39
6. Liu, G., Schulte, O.: Deep reinforcement learning in ice hockey for context-aware player evaluation. In: Lang, J. (ed.) *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*. pp. 3442–3448 (2018), doi: 10.24963/ijcai.2018/478
7. Macdonald, B.: A Regression-Based Adjusted Plus-Minus Statistic for NHL Players. *Journal of Quantitative Analysis in Sports* **7**(3) (2011), doi: 10.2202/1559-0410.1284
8. Macdonald, B.: An Improved Adjusted Plus-Minus Statistic for NHL Players. In: MIT Sloan Sports Analytics Conference (2011)
9. Mahanti, A., Carlsson, N., Mahanti, A., Arlitt, M., Williamson, C.: A tale of the tails: Power-laws in internet measurements. *IEEE Network* **27**(1), 59–64 (Jan/Feb 2013), doi: 10.1109/MNET.2013.6423193
10. Pettigrew, S.: Assessing the offensive productivity of NHL players using in-game win probabilities. In: MIT Sloan Sports Analytics Conference (2015)
11. Routley, K., Schulte, O.: A Markov Game Model for Valuing Player Actions in Ice Hockey. In: Meila, M., Heskes, T. (eds.) *Uncertainty in Artificial Intelligence*. pp. 782–791 (2015)

12. Schuckers, M., Curro, J.: 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)
13. Schulte, O., Khademi, M., Gholami, S., Zhao, Z., Javan, M., Desaulniers, P.: A Markov Game model for valuing actions, locations, and team performance in ice hockey. *Data Mining and Knowledge Discovery* **31**(6), 1735–1757 (2017), doi: 10.1007/s10618-017-0496-z
14. Schulte, O., Zhao, Z., Javan, M., Desaulniers, P.: Apples-to-apples: Clustering and Ranking NHL Players Using Location Information and Scoring Impact. In: MIT Sloan Sports Analytics Conference (2017)
15. Thomas, A., Ventura, S.L., Jensen, S., Ma, S.: Competing Process Hazard Function Models for Player Ratings in Ice Hockey. *The Annals of Applied Statistics* **7**(3), 1497–1524 (2013)