

Analyzing Passing Metrics in Ice Hockey using Puck and Player Tracking Data

David Radke, Jaxin Lu, Jackson Woloschuk, Tin Le[†],
Daniel Radke, Charlie Liu, and Tim Brecht

Cheriton School of Computer Science, University of Waterloo
University of Calgary[†]

Abstract. Traditional ice hockey statistics are inherently biased towards offensive events like goals, assists, and shots. However, successful teams in ice hockey require players with skills that may not be captured using traditional measures of performance. The adoption of puck and player tracking systems in the National Hockey League (NHL) has significantly increased the scope of possible metrics that can be obtained. In this paper, we compute recently proposed passing metrics from 1221 NHL games from the 2021-2022 season. We analyze the distributions of values obtained for each player for each metric to understand the variance between, and within, different positions. We find that forwards tend to complete fewer passes with smaller passing lanes, while defensemen pass to forwards significantly more than their defensive partners. Additionally, because these new metrics do not correlate well with traditional metrics (e.g., assists), we believe that they capture aspects of players' abilities that may not appear on the game sheet.

1 Introduction

The idea of using quantitative evidence to understand player tendencies and performance to inform management and strategic decisions has existed in sports for several decades [9]. In sports classified as “striking games”, such as baseball, analytics has transformed team operations and strategies [4]. This influence has lagged behind in “invasion games” such as football (soccer), basketball, handball, and ice hockey due to limitations in data collection and the complexities of the sport. Traditional (publicly available) statistics captured in ice hockey revolve around easily measurable offensive events (i.e., goals or shots) leading to the performance of offensive players being disproportionately captured. Successful teams in ice hockey, like all invasion games, require players with diverse abilities that existing offensively biased metrics do not capture, such as passing. This limited information makes constructing teams using quantitative evidence more difficult. The recent implementation of the puck and player tracking (PPT) systems in the National Hockey League (NHL) has led to several new metrics to quantify player behavior [12, 13]. In this paper, we utilize a larger dataset to study how passing metrics can be utilized to understand the variance in

behavior among players and players at different positions (metrics with larger variance may provide more opportunities to find under-valued players).

The main motivation behind the development of passing metrics in ice hockey was to capture other player contributions that might not show up on a game sheet [12]. Understanding how players compare to each other within the distribution of passing metrics provides valuable context for team building and management. We perform a deeper analysis into recently proposed ice hockey metrics from NHL puck and player tracking (PPT) data to show how passing metrics can be used to identify diverse behaviors among individuals. The contributions of this work are:

- We perform significant amounts of data cleaning to calculate passing metrics using PPT data from 1221 games in the 2021-2022 NHL season.
- We analyze the distributions of various passing metrics for forwards and defensemen. This provides insights into how much better highly-ranked players are when compared with other players.
- We find that after normalizing for ice time, forwards tend to complete fewer passes than defensemen and have smaller passing lanes, whereas defensemen complete significantly more passes to forwards and overtake more opponents.
- We show that the number of players overtaken with completed passes and the size of the passing lanes for completed passes do not correlate well with traditional offensive-oriented statistics like assists. We believe this demonstrates that some of our metrics capture aspects of players’ abilities that might not show up on the game sheet.

2 Related Work

Understanding how multiple players/agents work together most effectively is a significant area of research in organizational psychology and AI [1, 15]. A general finding is that group diversity, role specialization, and cohesiveness is important for group performance [7, 1, 23, 14]. Similar results have been found in football and sports analytics. Those analytics have focused on the performance of groups of players together [8, 11, 3, 10]. We use football to refer to *association* football (also known as soccer), not American football.

The implementation of passing metrics in football allows the analysis of a player’s decision-making and passing ability [20, 19], ability to overtake players with passes [21], impact on scoring probability [5], and ability to act under pressure [2]. In a low-scoring game like football, these models provide insight into players’ behaviors independently from offense and enables team building with diverse skills. Similar advancements in ice hockey have analyzed passing lane probabilities [17], as well as passing scenarios and pressure [12, 13].

Despite the development of models that use PPT data in ice hockey, no previous work has analyzed passing models to help understand general distributions, trends, or differences among players in the NHL. In this paper, we calculate

various passing metrics from recent work [12, 13] for 1221 games of the 2021-2022 NHL season. We analyze the distributions of each metric among players at different positions and within each position. Furthermore, we cross-reference related metrics to gain insight into how individual players behave with respect to multiple metrics.

3 Puck and Player Tracking Data

In the NHL, hockey is played on an ice surface that is 200 feet long and 85 feet wide. Tracking data is collected by SportsMEDIA Technology [18] (a partner of the NHL). They then derive event-level data (including completed passes) from the location tracking data. These event labels contain information about the time of the event and the identities and locations of the players involved. This paper focuses specifically on completed pass events. We have been granted early access by the NHL to the first full season for which the NHL used the PPT system (2021-2022). The PPT data and our resulting metrics are considered unofficial by the NHL, as the models used for creating event labels continue to be validated and improved. Additionally, the process of making statistical data official requires approval in the collective bargaining agreement, an ongoing process that has not been completed at this time. As a result, we do not provide information about individual player’s metrics. Also note that this data may differ from other datasets that contain complete and/or incomplete passes (e.g., a hand labeled dataset). We have processed 1221 of the 1312 regular season NHL games.¹

Location data is collected through tracking technology that is embedded into pucks and inserted into the sweaters of each player (on the back of the sweater, slightly right of the center of the shoulders). Location information contains x, y, and z-coordinates to record locations in 3-dimensional space. The x and y locations are relative to center ice (which is 0, 0). It is our understanding that when tested, the margin of error for the x and y coordinates is about 3 inches (the diameter of the puck) and very often as little as 1 inch. This is accurate enough for our purposes, as a puck traveling at speeds between 30 and 100 MPH would travel between 8 inches and 28 inches, respectively between readings. Further, our metrics are not overly sensitive to small changes in the puck’s location. The z coordinates (not used in this paper) are relative to the surface of the ice. Location data is recorded 60 times per-second for the puck and 12 times per-second for each player on the ice, resulting in a total of about 734,400 location readings of interest in a 60-minute game. Additionally, location data is obtained once per-second for players that are determined to be off of the ice. We interpolate all puck and player locations to 100 readings per-second to more easily identify the positions of all on-ice entities at precisely the same time.

¹ Some games could not be processed due to issues with the data sets and/or some special cases that our code hasn’t yet handled.

4 Background of Metrics

We briefly discuss the passing models used to derive the metrics in this paper. We refer the reader to the original paper [12] for further passing model details and to [13] for extensions and improvements to the original passing lane model. To ensure comparisons of different players are fair and are not simply a measure of ice time, we normalize our metrics by time on ice and/or games played, where appropriate (e.g., as is done in Section 6).

4.1 Passing Lane Model

The passing lane model we use in this paper is originally proposed in [12] and enhanced in [13]. The model uses the spatial locations of players in PPT data to estimate the available space between a passer p and any receiver r .

Figure 1 (adapted from [12]) shows the passing lane shape for a direct pass from p to r with three opponents. For each passing event, the model constructs a teardrop-like passing lane shape around the passer p and extending beyond the location of the receiver r (shaded regions). The size of the passing lane is determined by the nearest opponent to the pass and assigns a non-negative real-numbered value γ to be the openness of the pass. Figure 1 shows three passing lanes with respect to each of the three opponents. The γ value of this pass is $\gamma = 0.6$, since o_1 restricts the passing lane the most. We use the enhanced version from [13], where the *expected* locations of the receiver r and all opposing players based on current velocities are used to determine the passing lane. The enhancement also considers indirect passes off the boards. We developed a new constant-time algorithm to directly calculate γ instead of the previous binary search method. Refer to our previous paper [13] for more details.

4.2 Pass Overtaking Model

Previous work proposed and implemented models to understand progressing the puck beyond opponent players with passes [12]. At a high level, this model is represented as a zero-sum game, where a passer p gains a positive value for overtaking opponents with passes, and each opponent overtaken o receives a corresponding negative value. Formally, for a completed pass from p to r , if $\delta(x, y)$ is Euclidean distance between location x and y and NET is the center of the entrance to o 's net, o is considered overtaken if $\delta(p, \text{NET}) > \delta(o, \text{NET})$ and $\delta(o, \text{NET}) > \delta(r, \text{NET})$.² Because defensemen have greater opportunities to overtake more players, the model uses the fraction of players that are possible to be overtaken as the allocated credit. For example, if there are 3 players between p and the net (not counting the goalie) and the pass overtakes 2 opponents, the pass overtake value is 0.67. The passer p receives a positive value of +0.67 and each of the two overtaken players receives a negative value of -0.33 while the remaining non-overtaken player is unchanged.

² See Section 8 for variations we plan to consider in future work.




Fig. 1: Passing Lane.

Fig. 2: Adapted from [12]. The passing lane model for direct passes. The passing lane (shaded regions) surrounds the passer p and receiver r . The size and shape of this lane scales to the nearest opponent o (we show three examples of passing lanes with respect to three opponents). We use an expanded version that incorporates expected movement and indirect passes [13].

These values are aggregated into various metrics, including OVT (overtake total), BTT (beaten total), and PPM (passing plus-minus), calculated as $PPM = OVT - BTT$. We also calculate OVA, the average fraction of players overtaken with each pass ($\frac{OVT}{\text{num_passes}}$). Because there can be significant differences in the number of games played by different players, we use average values per game where appropriate. This ensures a fair comparison when examining and comparing different players.

5 Data Cleaning

When beginning our analysis we found several anomalies that needed to be corrected. Specifically, when using the timestamps associated with a fair number of completed passes, the puck was located at a relatively large distance from the passing player (e.g., significantly outside the reach of the player). To mitigate this issue, we performed a pass timestamp correction phase to better identify and adjust the time at which the event occurred. Adjusting these timestamps is also important to correctly identify the locations of all players on the ice at the time of the event. This is critical to obtain accurate passing metrics. All results in this paper are computed after adjusting the timestamps, which has significantly improved our metrics.

Our adjustment process begins by finding the timestamp for an event t in the PPT data. At a high level, our approach is to find a more accurate timestamp t' where the puck is sufficiently close to the passing player (i.e., within reach of the player). We determined a threshold of $\delta(p, \text{puck}) \leq 4$ feet to be a reasonable value, based on discussions with people at the NHL and personal measurements.

Metric	Description
avgOVT_20	The sum of the fraction of opponents overtaken by a player’s passes. We scale to 20 minutes of ice time and average per-game.
avgBTT_20	The sum of the fraction a player was overtaken by opponents’ passes. We scale to 20 minutes of ice time and average per-game.
avgOVA	The average fraction of opponents overtaken by a pass in a game. We average this value per-game.
avgPAA	Average γ (passing lane) value for completed passes. We average this value per-game.

Table 1: Summary of passing metrics discussed in this paper. Additive metrics (totals; end with “T”) are averaged over players’ games played (“avg”) and scaled to 20 minutes per-game if necessary (“_20”).

Any passes that could not be corrected using this technique are omitted from our dataset. This was only about 2.6% of the total number of completed passes. There are several possible ways to improve the accuracy of this approach including examining changes in the direction and speed of the puck. However, determining the accuracy of various techniques requires knowing ground truth, as a result this is a topic for future research.

6 Distribution Analysis of Passing Metrics

In the original work where we proposed these passing models we only had access to smaller PPT datasets so we did not conduct a detailed analysis for players [12]. In this paper, we analyze 1221 games and examine whether or not there are differences in passing metrics between forwards and defensemen and study the differences among individual players within the same position. We provide a summary of the metrics we analyze in Table 1. Our dataset includes 1000 players. To ensure that we have a sufficient sample size for various metrics we exclude players that did not play in at least 10 games and average at least 10 minutes of ice time per game. This reduced our dataset to 750 players (478 forwards and 272 defensemen). Because our work in this paper focuses on passing, we do not include goaltenders in any of our metrics or player counts.

To allow for fair comparisons among players that receive different amounts of ice time (since some metrics correlate with ice time) we normalize metrics (where appropriate) to 20 minutes per game. For each of the metrics in Table 1 we average over a player’s games. Thus, a metric such as OVT, the total fraction of opponents a player overtakes with their passes, will be represented as *avgOVT_20*: averaged over a player’s games (“avg”) and scaled to 20 minutes of ice time per-game (“_20”), where appropriate.

6.1 Distributions of Metrics Based on Position

We perform the Welch t-test [22] to analyze how the distributions of metrics vary between forwards and defensemen. When rejecting the null hypothesis for the mean values of a metric being equal between forwards and defensemen at a p -value of 0.05, we find that the mean between forwards and defense for both traditional statistics (e.g., goals, assists, points, shots, and shots blocked) and the new passing metrics in Table 1 are sufficiently different for every metric except for hits (which we do not consider in this paper). As a result, we analyze forwards and defensemen separately and use cumulative distribution functions (CDFs) to analyze the variance of distributions at each position.

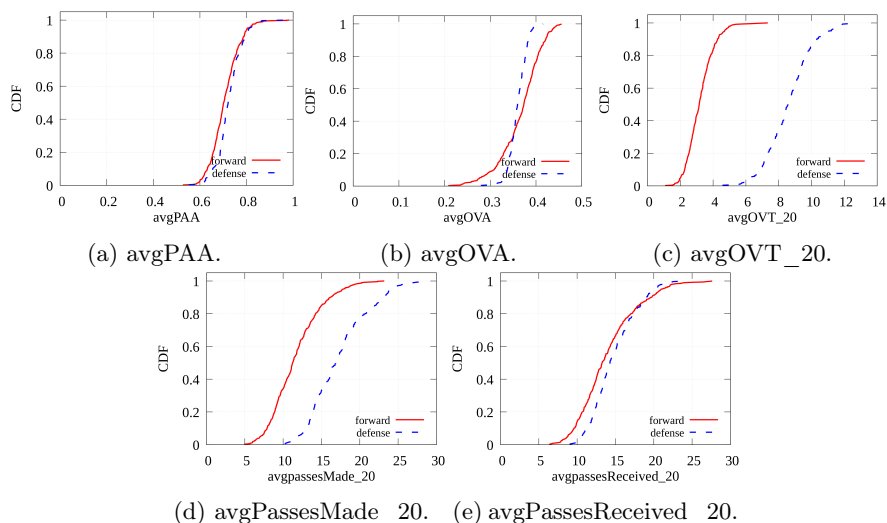


Fig. 3: CDFs plots for passing metrics separated by position, including (a) avgPAA: the average γ value for completed passes (lower indicates smaller passing lanes). (b) avgOVA: the average fraction of opponents overtaken by a pass (larger is better). (c) avgOVT_20: the total fraction of opponents overtaken by a pass (larger is better). Metrics for each player are averaged over the number of games they have played (“avg”) and when appropriate scaled to 20 minutes of playing time per-game (“_20”). (d) The average passes made by players per 20 minutes. (e) The average passes received by players per 20 minutes.

Passing Metrics Figure 3a shows the CDF for avgPAA, the per-game average γ value (passing lane size) for completed passes. Our results show that forwards and defensemen have distributions with similar shapes; however, the median defensemen tends to complete passes with slightly larger passing lanes. The for-

wards with the lowest avgPAA complete passes with about 47% smaller passing lanes than the forwards with the highest avgPAA.

Figure 3b shows the CDF for avgOVA, the average fraction of opponents a player overtakes per-pass, per-game. Higher values of avgOVA suggest the player overtakes a higher fraction of opponents with each pass (i.e., a stretch pass beating four of five players gives $\frac{4}{5} = 0.8$, while beating only the last defender gives $\frac{1}{1} = 1$). The 30th percentile values of each position are similar (about 0.35). Defensemen have lower variance in avgOVA with the range from the 20th percentile to 80th percentile being from 0.34 to 0.37 per-pass (34% to 37%) of the possible players per-pass. Comparatively, forwards have over double the variance than defensemen in avgOVA and the forwards with the highest avgOVA have over double the overtake value per-pass compared to the lowest forwards (0.45 compared to 0.21). The 20th percentile of forwards overtake an average of 33% of the possible players per-pass and the 80th percentile forwards overtake an average of about 40% of the possible players per-pass. The larger variance among forwards is likely caused by forwards typically having fewer opponents to overtake (2 or 3) compared to defensemen (4 or 5). We note that lowest percentile forwards are players that tend to make fewer than five passes per 20 minutes. We acknowledge that there may exist some players within our dataset that circumvent the intent of our filter and if they have a low number of passes, that could skew the distributions of some metrics. Future work could consider filtering techniques to remove players with too few passes.

Figure 3c shows the CDF for the avgOVT_20, the per-game average of the *total* fraction of opponents overtaken with passes normalized for 20 minutes of ice time. Higher values of avgOVT_20 suggest the player overtook a large fraction of opponents with their passes throughout a game. Our results show that the median defenseman achieves 2.8 times higher avgOVT_20 than the median forward (comparing 3 for forwards to 8.5 for defensemen). This difference of 5.5 avgOVT_20 increases to about 6 at the 80th percentile of forwards and defensemen (comparing 3.9 for forwards to 9.5 for defensemen). This change in the differences means the top defensemen for avgOVT_20 overtake more opponents compared to other defensemen than the top forwards compared to the rest of the forward population. Since the median avgOVA value for forwards is 4% higher than the median defensemen (see Figure 3b), we can conclude that higher values of avgOVT_20 for defensemen indicate that they complete more passes than forwards. This is confirmed in Figure 3d which shows the distribution of completed passes per 20 minutes. The shapes of the distributions among each position are almost identical, but defensemen tend to complete about five more passes than forwards at every percentile. The median forward completes about 12 passes per 20 minutes, whereas the median defenseman completes about 17. The forwards that complete the most passes complete up to 23 passes per 20 minutes (92% more than the forward median) and the defensemen that complete the most passes complete up to 27 passes (59% more than the defense median).

On average, despite defensemen completing roughly five more passes each than forwards, both positions tend to receive about the same number of passes

(Figure 3e). Comparing Figures 3d and 3e allows us to draw an interesting conclusion: defensemen complete passes to forwards significantly more often than to their defensive partner.

To understand this insight, consider that at even strength, the players that a forward can pass to are the two other forwards and the two defensemen. Assuming each of the other players is equally likely to be chosen (which may not be true), the probability of passing to a forward or defensemen is equal at 0.5. However, for defensemen there are three forwards and one defensemen to choose from. Again assuming the probability of passing to each of the other four players is equally likely (i.e., 0.25), the probability of passing to a forward is 0.75 and their defensive partner is 0.25. For the average pass reception curves for defensemen and forwards to be similar (as seen in Figure 3e) defensemen must complete passes to forwards three times more often. Considering that defensemen typically complete about five *more* passes than forwards (Figures 3d), defensemen must pass to forwards even more. Since these passes are likely up-ice, the higher frequency of passes from defensemen to forwards must be the main reason for high values of avgOVT among defensemen (Figures 3c).

Takeaways: We find that forwards make passes with slightly smaller passing lanes than defensemen. The variance among forwards for overtaking opponents with a pass (avgOVA) is significantly larger than with defensemen; however, the median forwards are only 4% higher than the median defensemen in avgOVA. Despite slightly lower median avgOVA, defensemen accumulate significantly higher *totals* for overtaking opponents (avgOVT_20) and complete about 5 more passes each game compared to forwards. Using Figures 3d and 3e, we find that defensemen pass to forwards significantly more than to their defensive partners.

6.2 Analyzing Player Differences

In this section we analyze individual players across a variety of metrics to gain insights into differences among players. One of the main passing metrics derived in our previous work [12] and discussed in Section 4 is passing plus-minus (PPM), defined as $PPM = OVT - BTT$. PPM gives insight into if a player overtakes more opponents than they are overtaken themselves; however, the metric removes additional context that may be important when understanding player behaviors. For example, a player that rarely overtakes opponents while also never being overtaken could have the same PPM value as a player that overtakes many opponents but often gets overtaken.

Figure 4a compares the two components of PPM to analyze the distribution of players along these two dimensions. The x -axis shows avgOVT_20, the total fraction of opponents that a player overtakes with their passes (per-game average, higher is better) and the y -axis shows avgBTT_20, their total fraction of being overtaken by opponents (per-game average, lower is better). Red triangles represent forwards and blue triangles represent defensemen. Players in the lower right corner overtake more opponents while not being overtaken by many opposing team passes. Analyzing where players are in these distributions may be

important when constructing forward lines or defensive pairings as a coach, or a roster as a manager.




Fig. 4: (a) The two components of passing plus-minus (PPM; presented in Section 4.2). The avgOVT_20 metric (overtake total; x -axis), the total (per-game average) fraction of opponents a player overtakes with their passes and avgBTT_20 (beaten total; y -axis), the total (per-game average) fraction that players are overtaken by opponents. (b) The average γ value (passing lane) for a player’s completed passes (avgPAA; x -axis) compared to players’ average completed passes per 20 minutes.

Figure 4a shows that there is diversity (or variation) among forwards with respect to both avgOVT_20 or avgBTT_20. Table 2 shows the avgOVT_20 and avgBTT_20 values with 95% confidence intervals for the forwards and defensemen with the highest, median, and lowest values for each metric. None of the confidence intervals for the three forwards intersect for either metric; thus, we can confirm that there exist forwards with differences that are statistically significant. In comparison, Figure 4a shows that defensemen mostly vary along the dimension of how they overtake opponents with passes (avgOVT_20). Table 2 confirms that the confidence intervals for defensemen do not intersect for avgOVT_20 but do intersect for avgBTT_20. Therefore, we conclude that defensemen mostly distinguish themselves from their peers by overtaking more opponents with their passes (avgOVT_20).

Figure 4b compares the per-game average value of γ (passing lane size) for a player’s completed passes (avgPAA_20; x -axis) and the average number of passes made by that player (avgPassesMade_20; y -axis). For both forwards and defensemen, players that complete the most passes (higher on the y -axis) tend not to have the lowest or highest values of avgPAA (x -axis) compared to the other players within their position. This implies that the players who complete a large number of passes do so in situations that are not anomalous (i.e., they are not mostly passing in easier situations).

Takeaways: There exists diversity among forwards with respect to overtaking opponents with passes and being overtaken by opponent passes. Defense-

Player	avgOVT_20	95% CI	avgBTT_20	95% CI
Fwd. Highest	7.28	± 1.84	4.66	± 1.09
Fwd. Median	3.16	± 0.42	3.00	± 0.19
Fwd. Lowest	1.05	± 0.43	2.05	± 0.17
Def. Highest	12.46	± 1.91	1.58	± 0.73
Def. Median	8.57	± 0.60	0.83	± 0.24
Def. Lowest	4.55	± 0.52	0.51	± 0.10

Table 2: Analyzing the mean and 95% confidence intervals for the highest, median, and lowest values for forwards and defensemen for avgOVT_20 and avgBTT_20. Our results show diversity among forwards with respect to both metrics while the highest defensemen tend to mostly separate themselves from their peers with respect to avgOVT_20.

men mainly separate themselves from their peers by overtaking more opponents, while there is less distinction with how defensemen are overtaken by opponents. At both positions, players that complete the most passes tend to do so with an average passing lane size instead of completing a disproportionate amount of easier passes with bigger passing lanes.

7 Comparative Analysis

Inspired by the work on “Meta-Analytics” (to examine stability, discrimination and independence of metrics) proposed by Franks *et al.* [6], we present a simple analysis of some of our metrics to show that avgOVT_20 and avgPAA do not correlate well with assists (i.e., to provide some indication of independence from a traditional offensive oriented statistic). We also compare the avgOVA, avgOVT_20, and avgPAA metrics obtained from the first 50% of the games with the same metrics computed across the last 50% of the games we have processed (to examine the stability of those two metrics). We divide games using the unique value assigned to each game (game id) which are typically ordered by scheduled date. Note that because a small number of games were postponed due to COVID-19, the split may not be precisely by the date games were played. In the future we plan to conduct an in depth analysis of all of our metrics (and other existing statistics) using the “Meta-Analytics” framework.

7.1 Comparison with Traditional Statistics

Figure 5a compares assists_82 (normalized to 82 games with 20 minutes per-game) and avgOVT_20 (the sum of the fraction of opponents overtaken by a player’s passes normalized to 20 minutes). Advancing the puck and overtaking opponents is a valuable aspect in invasion games like ice hockey [16]. Figure 5a shows there exists many players at both positions who overtake a significant number of opponents with completed passes who do not record a large number

of assists. These players with high avgOVT_20 values may not always show up on a game sheet; however, they may be playing important roles on their team.

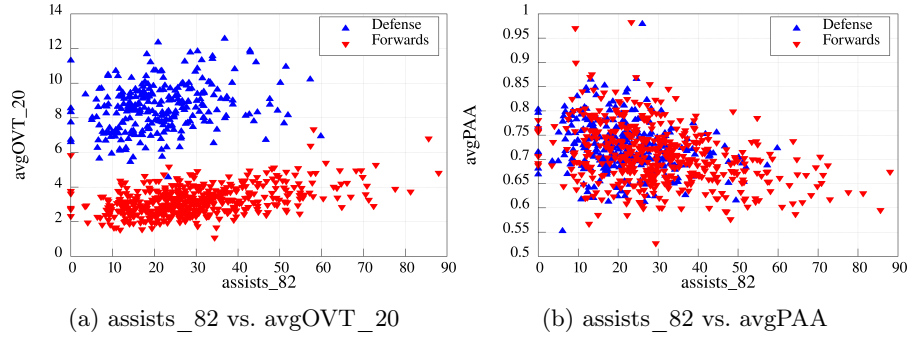


Fig. 5: (a) Assists projected to 82 games and 20 minutes per-game (assists_82) versus the total fraction of opponents overtaken by passes (avgOVT_20; per-game average; higher is better). (b) Projected assists (assists_82) versus the average value of γ for completed passes (avgPAA; per-game average; lower means smaller lanes).

Figure 5b compares assists_82 with avgPAA for players' completed passes in a game (avgPAA is the average γ value, or passing lane size; lower indicates smaller lanes). Note that there are no players with both high avgPAA and high assists_82 (i.e., no players in the top right of Figure 5b). However, many players with the highest assists_82 values have relatively low avgPAA (between 0.59 and 0.70). This may suggest a connection between recording many assists and being able to complete passes with smaller lanes. In future work we plan to examine this question more closely by separating, studying and comparing passing lanes for completed passes that result in assists. Again, we believe that considering traditional offensively-oriented statistics for a player could reduce one's ability to see other potentially important skills.

7.2 Evaluating Stability

Figure 6 compares metrics computed over the first 610 games with the same metric computed over the last 611 games. If the metrics obtained for each player during the first half of the games were able to perfectly predict the metric computed over the second half of the games, all data points would fall exactly on the diagonal line. These graphs indicate that the avgOVA and avgPAA metrics are well correlated across the two halves of the season (their correlation coefficients, r , are 0.87 and 0.89, respectively). The avgOVT_20 metric is strongly correlated with $r = 0.99$. For comparison we found (details and graphs have been excluded for brevity) that the correlation coefficient for players' points is $r = 0.80$ and

for goals is $r = 0.72$. This indicates that our new metrics are more stable (i.e., future values may be more predictable) than points and goals.




Fig. 6: Comparing different metrics from the first half of the games with the same metric computed over the second half of the games. Points on the diagonal line are perfectly correlated.

8 Discussion

While we perform an extensive analysis of several metrics and their distributions across players, our work has several limitations. One limitation is that we do not consider different factors such as coaching style (or team systems), manpower (e.g., even strength or not), goal differential, time of the game, and play location that may provide further insights. Future work may consider analyzing these scenarios separately.

Another limitation is the aggregation of metrics while including players with few samples. We filter our dataset by excluding players that don't receive a minimum average amount of ice time per game or have not played a minimum number of games. However, among the unfiltered players, some players recorded relatively few completed passes. Future work could apply additional filters (e.g., filtering players by a minimum number of samples).

Similar to the limitations with previous work [12, 13], we are only able to analyze completed passes. In the future we hope to discern or obtain information about unsuccessful passes. Additionally, our model for overtaking opponents does not consider potentially valuable passes such as those from close to (or behind) the net to the slot area, or east-to-west passes on odd-man-rushes, as overtaking opponents. Future work may adapt our model to include these types of passes.

9 Conclusions

Traditional ice hockey statistics disproportionately capture the offensive perspective of players. Understanding other characteristics of players' behaviors is important for constructing forward lines, defensive pairings, or entire teams. In this paper, we analyze several recently proposed passing metrics using PPT data from 1221 games of the NHL 2021-2022 season. We find that forwards tend to complete passes with slightly smaller passing lanes compared to defensemen; however, defensemen complete more passes and overtake more opponents. Examining players by comparing their scores on the basis of two metrics reveals the diversity of behavior among players with regards to pass overtaking and being overtaken by passes. Finally, because these new metrics do not correlate well with traditional metrics, we believe they capture aspects of players' abilities that may not appear on a traditional game sheet. This analysis may be of significant interest to coaches and managers as they attempt to construct successful teams.

Acknowledgments

We thank Rogers Communications for providing funding for this project. In addition, this research is partially funded by the Natural Sciences and Engineering Research Council of Canada (NSERC), two University of Waterloo Math Undergraduate Research Awards and a University of Waterloo Undergraduate Research Fellowship. We thank the members of the National Hockey League's Information Technology Team and the Stats and Information Team for their participation in fruitful discussions and their insights related to this work. From AWS, we thank Candi Jeronimo for providing us with credits and Varad Ram for technical assistance that enabled us to process the games using AWS. We thank Neel Dayal and Jonah Eisen from Rogers Communications for making this project possible. We thank Robby Brecht for his feedback on an earlier draft of this paper. Finally, we thank the anonymous reviewers for their constructive and helpful feedback.

References

1. Andrejczuk, E., Rodriguez-Aguilar, J.A., Sierra, C.: A concise review on multi-agent teams: contributions and research opportunities. *Multi-Agent Systems and Agreement Technologies* (2016)
2. Andrienko, G., Andrienko, N., Budziak, G., Dykes, J., Fuchs, G., Landesberger, T.V., Weber, H.: Visual analysis of pressure in football. *Data Mining and Knowledge Discovery* **31**, 1793–1839 (2017)
3. Bransen, L., Van Haaren, J.: Player chemistry: Striving for a perfectly balanced soccer team. *arXiv preprint arXiv:2003.01712* (2020)
4. Elitzur, R.: Data analytics effects in major league baseball. *Omega* **90**, 102001 (2020)

5. Fernández, J., Bornn, L., Cervone, D.: Decomposing the immeasurable sport: A deep learning expected possession value framework for soccer. In: 13th MIT Sloan Sports Analytics Conference (2019)
6. Franks, A., D’Amour, A., Cervone, D., Bornn, L.: Meta-Analytics: Tools for understanding the statistical properties of sports metrics. *Journal of Quantitative Analysis in Sports* **12**(4), 151–165 (2016)
7. Hong, L., Page, S.E.: Groups of diverse problem solvers can outperform groups of high-ability problem solvers. *Proceedings of the National Academy of Sciences* **101**(46), 16385–16389 (2004)
8. Ingersoll, K., Malesky, E., Saiegh, S.M.: Heterogeneity and team performance: Evaluating the effect of cultural diversity in the world’s top soccer league. *Journal of Sports Analytics* **3**(2), 67–92 (2017)
9. Lewis, M.: *Moneyball: The art of winning an unfair game*. WW Norton & Company (2004)
10. Ljung, D., Carlsson, N., Lambrix, P.: Player pairs valuation in ice hockey. In: *MLSA@PKDD/ECML* (2018)
11. Nsolo, E., Lambrix, P., Carlsson, N.: Player valuation in european football. In: *Machine Learning and Data Mining for Sports Analytics: 5th International Workshop, MLSA 2018, Co-located with ECML/PKDD 2018, Dublin, Ireland, September 10, 2018, Proceedings 5*. pp. 42–54. Springer (2019)
12. Radke, D.T., Radke, D.L., Brecht, T., Pawelczyk, A.: Passing and pressure metrics in ice hockey. *Workshop of AI for Sports Analytics* (2021)
13. Radke, D., Brecht, T., Radke, D.: Identifying completed pass types and improving passing lane models. In: *Linköping Hockey Analytics Conference*. pp. 71–86 (2022)
14. Radke, D., Larson, K., Brecht, T.: Exploring the benefits of teams in multiagent learning. In: *IJCAI* (2022)
15. Radke, D., Larson, K., Brecht, T.: The importance of credo in multiagent learning. *ALA Workshop at AAMAS* (2022)
16. Radke, D., Orchard, A.: Presenting multiagent challenges in team sports analytics. In: *AAMAS* (2023)
17. Ritchie, R., Harell, A., Shreeves, P.: Pass evaluation in women’s olympic ice hockey. In: *Proceedings of the 5th International ACM Workshop on Multimedia Content Analysis in Sports*. pp. 65–73 (2022)
18. SMT: SportsMEDIA technology. <https://www.smt.com> (2021), accessed: 2023-03-19
19. Spearman, W.: Beyond expected goals. In: *Proceedings of the 12th MIT sloan sports analytics conference*. pp. 1–17 (2018)
20. Spearman, W., Basye, A., Dick, G., Hotovy, R., Pop, P.: Physics-based modeling of pass probabilities in soccer. In: *Proceeding of the 11th MIT Sloan Sports Analytics Conference* (2017)
21. Steiner, S., Rauh, S., Rumo, M., Sonderegger, K., Seiler, R.: Outplaying opponents—a differential perspective on passes using position data. *German Journal of Exercise and Sport Research* **49**, 140–149 (2019)
22. West, R.M.: Best practice in statistics: Use the welch t-test when testing the difference between two groups. *Annals of Clinical Biochemistry* **58**(4), 267–269 (2021)
23. Zaccaro, S.J., Dubrow, S., Torres, E.M., Campbell, L.N.: Multiteam systems: An integrated review and comparison of different forms. *Annual Review of Organizational Psychology and Organizational Behavior* **7**(1), 479–503 (2020)