Published in the Artificial Intelligence for Sports Analytics (AISA) Workshop at IJCAI '21
Montreal, QC, Canada
August 21-23, 2021

# Passing and Pressure Metrics in Ice Hockey

**David Radke**[1] , **Daniel Radke**[2] , **Tim Brecht**[1] and **Alex Pawelczyk**[1]
[1]David R. Cheriton School of Computer Science, University of Waterloo
[2]Department of Computer Science, Universität des Saarlandes
dtradke@uwaterloo.ca, s8daradk@stud.uni-saarland.de, {brecht, alex.pawelczyk}@uwaterloo.ca

## Abstract

Advancements in player tracking technology and analytics have revolutionized how professional sports are managed, played, and even consumed. However, these tracking systems have mainly failed to translate to ice rinks, leaving ice hockey to collect relatively coarse and offensively-biased data for player statistics. As a result, current artificial intelligence (AI) models for player valuation and optimal group formation for ice hockey are limited to comparing with offensively-biased data, reinforcing these biases in models. The National Hockey League (NHL) used a new tracking system in the 2020 Stanley Cup Playoffs, and we design a suite of novel analytics that add new insights into the performance and behaviour of players, groups of players, and teams. We calculate metrics for passing lanes and player movement, passing effectiveness, and pressure which have not previously been possible to compute in hockey. We expect our analytics to support future and more accurate performance and coalition models, and to be of direct interest to the ice hockey and AI communities.

## 1 Introduction

Data analytics in professional sports has revolutionized how and what data is collected and used, and is expected to reach revenues of $4.5 billion by 2024 [BusinessWire, 2018]. In the National Hockey League (NHL), a team can roster 18 skating players and two goalies per-game, with five skaters and one goalie in the game at a time. Players interchange during play with fatigue and according to the situation, meaning players are often involved in different scenarios all over the ice. However, current hockey statistics consider mainly offensive events like goals, assists, and shots-on-goal (SOG), which occur relatively infrequently. The 2019-2020 NHL season saw an average of 3.02 goals scored per-team per-game [NHL, 2020]; thus, traditional statistics would suggest a maximum of nine players (45% of the roster) receive a point for a goal or assist (assuming the maximum of two assists for every goal). In addition, a single player accumulating multiple points on multiple goals means the number of possible players that receive statistical updates begins to decrease. SOG are more common than goals, but the 2019-2020 average of only 1.7 SOG per-player per-game also fails to accurately represent performance throughout a game, especially a player's non-offensive play. In some cases, not including the recorded time-on-ice, there may be no indication from current statistics that a player actually played in a game despite potentially bringing value to their team in other ways such as defense or play making. This may lead to misaligned incentives between players trying to improve their statistics and their team relying on them to play a less offensive role.

Success in ice hockey relies on possession of the puck and passing between teammates. Despite hundreds of passes every game, no current ice hockey statistic records players' passing effectiveness. As a result, existing artificial intelligence (AI) models for player valuation and optimal pairing, or coalition formation have no choice but to compare their results or use models that rely on existing offensively-biased statistics based around goals, assists, and SOG. These statistics which overlook valuable position players who may not generate as much offense [Ljung *et al.*, 2018; Luo *et al.*, 2020]. Utilizing biased data is a known problem in AI and has dramatic impact when deployed. New hockey analytics that capture other critical aspects of the game are necessary for a deeper understanding of player value and improved AI models.

Insights from modern analytics across other sports, many of which depend on high-resolution player tracking systems [Amin, 2018; Baysal and Sahin, 2016; Manafifard *et al.*, 2017], are used to improve player performance, roster management, and various AI applications [Lewis, 2003; Lindström *et al.*, 2020]. Unfortunately, the tracking methods used in other sports have failed to transfer to ice hockey due to technical challenges caused by the fast pace, small puck, white-coloured ice, and other hardware challenges [Walkters *et al.*, 2020; Vats *et al.*, 2020; Douglas and Kennedy, 2019]. This has led to hockey metrics which infer performance based on the aggregation of sparse events instead of tracking actual behaviour [Hammond, 2011]. Recent technological advancements have led to the implementation and testing of a new player and puck tracking system during the 2020 NHL Stanley Cup Playoffs for the first time in real gameplay. While tracking itself is a technological accomplishment, the data does not provide general managers, coaches, players, agents, fans, or AI models with meaningful information.

We propose a suite of novel player analytics in hockey designed to extract and aggregate insightful information from the raw tracking data. To the best of our knowledge, we are the first to propose new ice hockey analytics using the recently collected tracking data. Our goal is to represent situational player strengths, weaknesses, and trends throughout the ice surface to offer new perspectives into player performance and value. We make the following contributions:

- We develop new hockey analytics designed to improve our understanding of passing, player movement, trends, and how players respond under pressure.
- We propose a new passing lane metric and a metric to determine the degree to which potential pass receivers are open (or available). These metrics can can be adapted to other sports that involve passing.
- We utilize real tracking data from the 2020 Stanley Cup Playoffs to compute our new metrics and discuss initial insights that our analytics provide.

The aim of our work is two fold. First, we urge the AI and hockey communities to recognize the need for higher resolution hockey analytics that provide insight into how players perform apart from sparse offensive events. Second, we hope our work provides new initial benchmarks for player valuation and coalition formation AI models in hockey, helping to provide a more balanced view of player performance. Establishing new metrics are necessary for AI to mature in ice hockey; thus, we view our contributions as the first steps in the new frontier of ice hockey analytics supported by tracking data and towards a large corpus of future work involving game theory [Yan *et al.*, 2020], multi-team systems [Zaccaro *et al.*, 2020], and cooperative AI [Dafoe *et al.*, 2021].

## 2 Background

Ice hockey in the NHL is played on an ice surface that is 200 feet long and 85 feet wide (imperial units are used in the NHL). A game consists of two teams competing for three 20-minute periods. Overtime rules vary between the regular season and playoffs. A maximum of six players of any combination of defense, forwards, and one goalie are allowed on the ice at any time for each team. Penalties remove a player from the ice surface for two or five minutes depending on the severity of the infraction, so that the penalized team temporarily has fewer players in the game. Both teams playing with the same number of players on the ice is referred to as "even strength" and without penalties a team is at "full strength".

The objective of hockey is to score a goal by putting the puck into the opposing team's net, referred to as the *attacking net*. The team with the most goals at the end of the game wins. Throughout the paper, we refer to a player in possession of or passing the puck as $p$; we use $r$ to denote a teammate receiving the pass; and $o$ refers to any player on the opposing team.

## 3 Related Work

### 3.1 Player Tracking in Soccer

Over the past decade, advancements in player tracking technology for professional soccer have enabled researchers to develop advanced analytics for team and individual performance. Several studies leverage player and ball tracking data to analyze passing, quantifying a pass's disruption to defensive formations [Goes *et al.*, 2019] and the number of outplayed opponents [Steiner *et al.*, 2019]. Other work, with sufficient tracking data, has used deep learning to evaluate the behaviour of players in soccer by predicting the performance of "league average" players and teams in simulated scenarios [Meerhoff *et al.*, 2019; Lindström *et al.*, 2020]. [Kempe *et al.*, 2018] evaluate tracking data to show that scoring events and player performance typically highlighted by human annotators can be determined and assessed automatically using ridge regression. Similarly in [Fernández, 2019], tracking data, deep learning-based models and stochastic processes are used to calculate the likelihood that a soccer possession ends in a goal and assign value to passes.

### 3.2 Analytical Insights in Ice Hockey

Existing research in ice hockey player valuation and analytics has mainly utilized event data, typically using the SPORTLOGiQ NHL dataset which records the location, involved players, and time of events, such as shots, hits, and passes [Liu *et al.*, 2018; Silva *et al.*, 2018; Yu *et al.*, 2019]. These have been utilized to extract representations of players' abilities through clustering and Markov Decision Processes [Schulte and Zhao, 2017], and deep learning [Liu *et al.*, 2020; Mehrasa *et al.*, 2018; Guo *et al.*, 2020]. These models typically benchmark their results with the offensively-biased currently recorded statistics due to a lack of recorded ground truth statistics based around other aspects of the game.

Current *advanced statistics* used by the NHL, such as Corsi and Fenwick, have been shown to be good performance indicators for teams [Macdonald, 2012]. Corsi is a plus/minus rating of shot attempts (*shots for* minus *shots against*) during even strength play. Shot attempts include blocked shots, missed shots, and SOG, which have traditionally been tracked by human annotators [Hammond, 2011]. Fenwick is similar to Corsi, but omits blocked shots since this could be a player's positive skill. Positive Corsi and Fenwick scores imply the player's team produced more offense than their opponent during the even strength time the player was on the ice.

## 4 Dataset

We utilize a proprietary sample dataset made available through an exclusive contract with Rogers Communications [Rogers, 2021] and Sportsnet [Sportsnet, 2021]. Tracking data is collected by SportsMEDIA Technology [SMT, 2021] (a partner of the NHL) with infrared puck and player tracking systems in every NHL arena. Our dataset consists of player and puck tracking data from games five and six of the 2020 Stanley Cup Finals between the Tampa Bay Lightning and the Dallas Stars. Sensors are located inside the puck and on the right shoulder of every player, sampled at rates of 60 times per second (60 Hertz) for the puck and 12 Hertz for players. The location coordinate system is based on the dimensions of a hockey rink with center-ice being the origin $(0, 0)$. The x-axis has a range of $-100 \leq x \leq 100$, and the y-axis has a range of $-42.5 \leq y \leq 42.5$. Our dataset includes

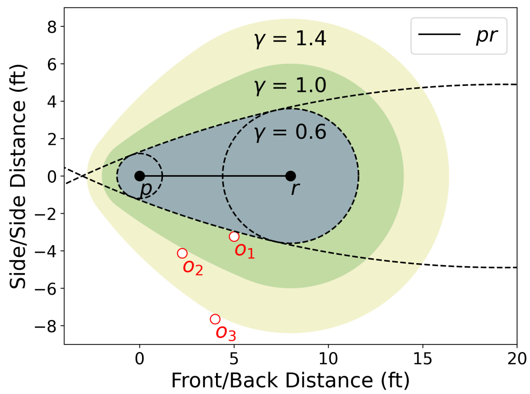

Figure 1: Our method to calculate a passing lane. PA is the value of $\gamma$ corresponding with the largest passing lane not containing any $o$. In this case, the pass availability PA = 0.6.

pre-labeled events such as hits, shots, passes, faceoffs, and puck possession gains and losses. We perform data cleaning and pre-processing for player and puck locations rounded to 0.1 second intervals to align with provided event labels. Note that our analytics can be implemented using data from any tracking system that provides frequent enough data for the locations of all players and the puck (or ball).

## 5 Our Analytics

Ice hockey is a unique sport due to its speed and dynamic game flow, with specific plays and decisions often taking only fractions of a second. Player and puck tracking makes high resolution analysis into specific situations possible. In this section, we detail our new player analytics with tracking data. A portion of our analytics are averaged over 60-minutes, as a game in our dataset extended to double overtime, lasting almost 90 minutes. Due to space constraints we often report the mean of each metric, however one could also analyze the distribution of values or changes over time to obtain more detailed information. Our implementation is based in ice hockey, however we expect our contributions to translate to other multi-player possession sports.

### 5.1 Passing Lanes

When a team is at full strength, a passer has four passing lanes to their teammates, not including the goalie. We develop a novel method for analyzing passing lanes and record two metrics: the passing lane availability average of $p$'s completed passes PAA; and the average "openness", or passing lane availability, to a player who is a potential receiver $r$ for their teammate OPA. PAA provides insight into players' decisions, passing skill, and risk, while OPA shows their ability to become open for a pass from a teammate.

Calculating a passing lane value has several challenges. Specifically, a proper method should: **1)** always assign a passing lane a real number to allow aggregation; **2)** incorporate the area surrounding $p$ and $r$ (i.e., consider the proximity of all opponents); **3)** be asymmetric with respect to $p$ and $r$ as passes are directed events, and **4)** scale with respect to the pass length, as longer passes take more time to reach $r$, giving opponents more time to potentially disrupt the pass.

One attempt to analyze passing lanes in soccer uses the difference in angle between the direct passing lane, the line $pr$ from $p$ to $r$, and the most threatening opponent $o$ within the distance of the pass [Steiner *et al.*, 2019]. While this method addresses 3) and 4), it fails to satisfy 1) and 2). The circle-based $\beta$-skeleton [Kirkpatrick and Radke, 1985] from graph theory calculates a maximal empty region between two points, addressing 1), 2), and 4), but point 3) remains unsatisfied. Therefore, we propose a novel method using Euclidean geometry to simultaneously satisfy all four points above.

Our algorithm uses a non-negative real-number parameter $\gamma$ to define the passing lane region from $p$ to $r$ $(PL(p, r, \gamma))$, shown as the coloured regions in Figure 1. Specifically, $\gamma$ determines the shape and size of $PL(p, r, \gamma)$. PA, the pass availability of a potential pass, is the value of $\gamma$ with the largest $PL(p, r, \gamma)$ that does not contain any $o$. For example, Figure 1 shows three passing lanes from $p$ to $r$ with $\gamma \in \{0.6, 1.0, 1.4\}$. Since no opponents are inside the passing lane region when $\gamma \leq 0.6$, PA = 0.6. PAA is the average value of PA for a player's completed passes.

Our passing lane is asymmetrical, meaning $PL(p, r, \gamma) \neq PL(r, p, \gamma)$. The boundary of our passing lane is constructed by two asymmetric circles $C_p$ and $C_r$, centered at $p$ and $r$ respectively, and the area between the circles contained by two arcs tangent to these circles, seen in Figure 1. This asymmetry accounts for the directional nature of a pass, $p \rightarrow r$, and the notion that an opponent close to $r$ has more time to react.

The radii of circles $C_p$ and $C_r$, $R_p$ and $R_r$ respectively, are calculated using Equations 1 and 2, where $d = \delta(p, r)$, the euclidean distance from $p$ to $r$, and $t = 0.25$ is fixed to ensure passing lane asymmetry, scaling the growth of $R_r$ 3 times faster than $R_p$ as $\gamma$ increases.

$$R_p = \gamma d t \qquad (1) \qquad\qquad R_r = \gamma d(1 - t) \qquad (2)$$

The arcs come from two larger circles tangent to both $C_p$ and $C_r$ with radii $R_a = \frac{cd}{\gamma}$, where $c$ is used to scale the numerator to ensure the centre points of the arcs are real and not imaginary. Specifically, $c$ must be $\geq 3$ due to our choice of $t = 0.25$, so we use $c = 4$. Since $R_r$ grows faster than $R_p$ as $\gamma$ increases, when $\gamma \geq 2$, $C_r$ encompasses $C_p$ and the passing lane is simply $C_r$.

The parameter $\gamma$ monotonically affects the size and shape of $PL(p, r, \gamma)$, so that $PL(p, r, \gamma_1) \subseteq PL(p, r, \gamma_2)$ for $0 \leq \gamma_1 \leq \gamma_2$. The naïve method of calculating $\gamma$ for any $p$, $r$, and $o$ is to initialize $\gamma = 0$ so that $PL(p, r, \gamma) = pr$, and subsequently increment $\gamma$ by a small value until the lane contains $o$. The most recent value of $\gamma$ is then determined as the PA value when a pass is made (i.e., the largest $\gamma$ not containing $o$). For efficiency, we use a binary search with a stopping resolution of 0.01 to calculate $\gamma$. Since there are always opponents, we guarantee $\gamma < \infty$ for each pass and PA $< \infty$. We summarize our passing lane metrics in Table 1.

### 5.2 Passing Performance

A completed pass is when the puck moves from one player, the passer $p$, to another player on the same team, the receiver $r$. A pass that advances the puck beyond an opponent $o$, overtaking them, can leave $o$ in a poor defensive position. Despite

| Sym | Description |
|---|---|
| PAA | Pass availability (value of $\gamma$) for a player's completed passes (average). |
| OPA | Openness (value of $\gamma$) as a receiver for each timestep a teammate has the puck (average). |

Table 1: Passing lane analytics for a single player.

hundreds of completed passes every game, there currently exists no recorded ice hockey metric to understand a player's passing with respect to opponent positions. Therefore, we introduce metrics to analyze passing in ice hockey inspired by overtaking opponents in soccer [Steiner *et al.*, 2019] and zero-sum games [Dafoe *et al.*, 2020].

We determine a player's average number of successful passes over 60 minutes (PASA) and for each pass compute the ratio of possible opponents overtaken. The sum of those values is the overtaken ratio total (OVT) and the mean of those values is the overtaken ratio average (OVA). On the defensive side, we similarly evenly distribute that ratio of opponents overtaken across all players that have been beaten with the pass. This sum of these values is denoted as BTT (beaten total). Note that OVT and BTT across all players is a zero-sum game. Inspired by the current plus/minus statistic for goals, for each player we calculate the difference between overtaking players and being overtaken as "passing plus/minus" PPM = OVT - BTT. Additionally, we normalize PPM by the number of times a player overtakes, or is overtaken by an opponent (count) for "normalized PPM", NPPM = $\frac{\texttt{PPM}}{\texttt{count}}$. As is the case with the scoring plus/minus statistic, we only consider even-strength play to avoid biases for players given heavily offensive or defensive roles during penalties. Finally, we record a player's average number of turnovers, an immediate change of possession between teams, over 60 minutes (TOA).

Previous work in soccer records the number of opponents overtaken by each pass and computes a total from all passes [Steiner *et al.*, 2019]. A pass overtaking an opponent results in $o$ being in poor defensive position, where $r$ receiving the pass is now closer to $o$'s net. Formally, if NET is the entrance to $o$'s net, $o$ is considered overtaken if $\delta(p, NET) > \delta(o, NET)$ and $\delta(o, NET) > \delta(r, NET)$. Directly translating this metric to ice hockey results in heavily defensemen-biased metrics since formations are more compact and defensemen typically have more opportunity to overtake opponents in their offensive direction than forwards. Therefore, we scale the resulting values by calculating the ratio of players overtaken with a pass that were possible to overtake. For example, if there are 3 players between $p$ and the net (not counting the goalie) and the pass overtakes 2 opponents, the computed value is 0.67. In this example, OVT for $p$ would increase by 0.67, and BTT for the two opponents overtaken would increase by $\frac{0.67}{2} = 0.33$ each. Our intuition is the closer the value is to 1, the fewer players remaining for $r$ to beat. We summarize our passing analytics in Table 2.

| Sym | Description |
|---|---|
| PASA | Successful passes made (average). |
| OVA | Overtaken opponents with passes (average). |
| OVT | Overtaken opponents with passes (total). |
| BTT | Beaten by opponent's passes (total). |
| PPM | OVT - BTT |
| NPPM | $\frac{\texttt{OVT}-\texttt{BTT}}{\texttt{count}}$, PPM is normalized by the number of passes when OVT or BTT are incremented. |
| TOA | Turnovers made (average). |

Table 2: Passing analytics for a single player over 60 minutes.

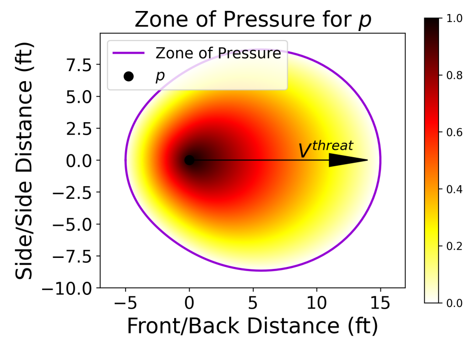


Figure 2: Zone of pressure (ZoP) around $p$, direction of $V^{threat}$, and pressure $p$ experiences from $o \in (0, 1)$ regarding the location of $o$. Darker corresponds with more pressure.

## 5.3 Metrics Indicating Pressure

Exerting pressure on the puck possessor $p$ is a tactic used by one or more opponents $o$ to try and force $p$ to make a mistake. Calculating pressure can help analyze how players respond to different scenarios, if they are putting themselves in difficult situations and if they are reacting quickly enough. To understand pressure when players are in offensive situations we calculate the average amount of pressure on the player when they move the puck PrMA and shoot the puck PrSHA. On the defensive side we compute the average pressure each player exerts on shooting opponents PrSOA.

To calculate pressure we slightly modify the Zone of Pressure (ZoP) used in soccer [Andrienko *et al.*, 2017]. First, we decrease the ZoP size since team formations and the playing surface are smaller in ice hockey. Next, we scale the dispersion of pressure in response to the smaller ZoP and to account for the use of hockey sticks which increasing the reach of $o$.

As in [Andrienko *et al.*, 2017], let the direction in front of $p$ be represented by vector $V^{threat}$ (Figure 2). Instead of $V^{threat}$ pointing up-ice, we define $V^{threat}$ to point towards the center of the opening to the attacking net. The pressure boundary is calculated using the same method as in [Andrienko *et al.*, 2017], resulting in an oval-like structure surrounding $p$ with greater distance along $V^{threat}$.

The limits of the ZoP directly in front and behind $p$ are denoted as $d_{front}$ and $d_{back}$. We modify $d_{back}$ to be 5 feet,

| Sym | Description |
|---|---|
| PrMA | Pressure when moving the puck (average). |
| PrSHA | Pressure for each SOG by $p$ (average). |
| PrSOA | Pressure excerted on a shooting opponent while inside the ZoP of $o$ (average). |

Table 3: Pressure analytics for a single player.

| | Player | PAA | OPA | PASA | OVA | OVT | BTT | PPM | NPPM | TOA | PrMA | PrSHA | PrSOA |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | Bogosian* | 1.19 | 0.44 | 15.00 | 0.55 | 4.38 | 0.98 | 3.40 | 0.28 | 1.00 | 0.28 | 0.47 | 0.57 |
| 2 | Cernak | 0.97 | 0.45 | 21.25 | 0.42 | 4.56 | 0.35 | 4.21 | 0.27 | 2.00 | 0.31 | 0.00 | 0.38 |
| 3 | Sergachev | 0.78 | 0.62 | 18.04 | 0.43 | 4.53 | 0.79 | 3.74 | 0.25 | 0.40 | 0.29 | 0.11 | 0.48 |
| 36 | Maroon | 1.42 | 0.36 | 4.01 | 0.67 | 0.27 | 2.27 | -2.00 | -0.14 | 1.60 | 0.63 | 0.00 | 0.56 |
| 37 | Volkov* | 0.63 | 0.44 | 3.00 | 0.75 | 0.75 | 2.60 | -1.85 | -0.14 | 0.00 | 0.55 | 0.00 | 0.48 |
| 38 | Benn | 0.93 | 0.44 | 9.62 | 0.40 | 0.64 | 6.39 | -5.75 | -0.15 | 4.01 | 0.39 | 0.76 | 0.43 |

Table 4: Analytic results, top three and bottom three players (excluding goalies) ordered by normalized passing plus/minus (NPPM) in decreasing order. DAL = green, TBL = blue. (* = played one game)

equivalent to about the length of a hockey stick so that $o$ can reach $p$ when they are directly behind $p$. Keeping the same scale as in [Andrienko *et al.*, 2017], we make $d_{front}$ three times larger (15 feet). We also simplify and alter their formula for pressure from $o$ onto $p$ to scale linearly, so that the pressure $o$ exerts on $p$ inside of the ZoP, $o(p) = 1 - \frac{\delta(p,o)}{\delta(p,z)}$, where $z$ is the point on the ZoP boundary in the direction of $o$. This dispersion is less dramatic than in soccer because hockey sticks increase the reach of $o$. The total pressure experienced by $p$ with $O$ pressers is calculated by $\texttt{Pr(p)} = \sum_o o(p)$. Figure 2 shows pressure in relation to $p$, with darker red corresponding to higher pressure depending on the location of $o$. We summarize pressure analytics in Table 3.

## 6 Evaluation

We compute our analytics using the dataset from two 2020 Stanley Cup Playoff games between the Tampa Bay Lightning (TBL) and the Dallas Stars (DAL). Due to space limitations, Table 4 shows only the top and bottom three players across both games according to decreasing normalized passing plus/minus (NPPM). The rows are coloured blue for TBL and green for DAL players.

The top three players in Table 4 are all defensemen for TBL, and the bottom three are all forwards. All players display variation in PAA and OPA, with a slight openness advantage for the defensemen, likely due to their positioning. Bogosian (TBL), Cernak (TBL), and Sergachev (TBL) all frequently overtake opponents (high OVT) and do not get overtaken much (low BTT), leading to relatively high NPPM. While Volkov (TBL) and Maroon (TBL) have the highest OVA, their frequency of passes is low and they are overtaken more than other players, leading to low NPPM. Benn (DAL) averages almost 10 passes over 60 minutes, however he does not overtake opponents as often as some other players, seen through OVA and OVT, while he is also more frequently overtaken with BTT = 6.39. Pressure measures are most useful when compared relative to other players, or cross-referenced with other metrics as described next.

### 6.1 Cross-Referencing Metrics

We cross-reference a subset of our analytics to show how metric pairings could potentially lead to further player insights. Figure 3a compares the average pressure on puck movement (PrMA) with the average turnovers (TOA) and Figure 3b compares the pass availability average (PAA) with the average turnovers (TOA). We highlight a handful of players which

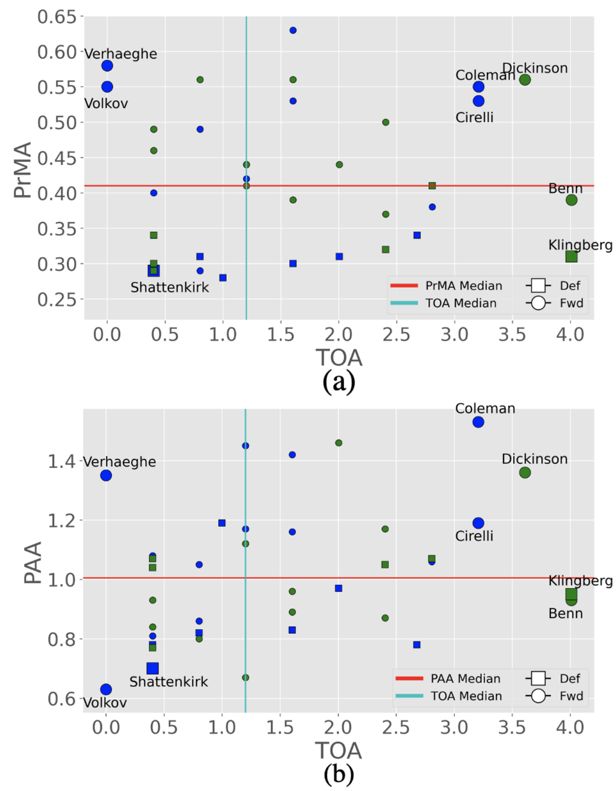


Figure 3: Cross-referencing TOA with (a) PrMA and (b) PAA.

demonstrate potentially interesting observations. In Figure 3a, Coleman (TBL), Cirelli (TBL), and Dickinson (DAL) have both high PrMA and TOA, suggesting they move the puck when under a lot of pressure and turn the puck over relatively frequently. In contrast, Verhaeghe (TBL) and Volkov (TBL) move the puck under similar pressure, but did not turn the puck over. Shattenkirk (TBL), Klingberg (DAL), and Benn (DAL) move the puck under less pressure than the median, however Shattenkirk averages fewer turnovers.

We note that although PAA shows completed passes, it also captures general trends of risk level for a player's decisions (lower values means smaller passing lanes and riskier passes). In Figure 3b, we see that Coleman (TBL), Cirelli (TBL), Dickinson (DAL), and Verhaeghe (TBL) have higher PAA than the dataset median, suggesting they find and/or use more open passing lanes. Comparing these metrics with average turnovers (TOA), we see that Verhaeghe has a much lower turnover rate than Coleman, Cirelli, and Dickinson. Players

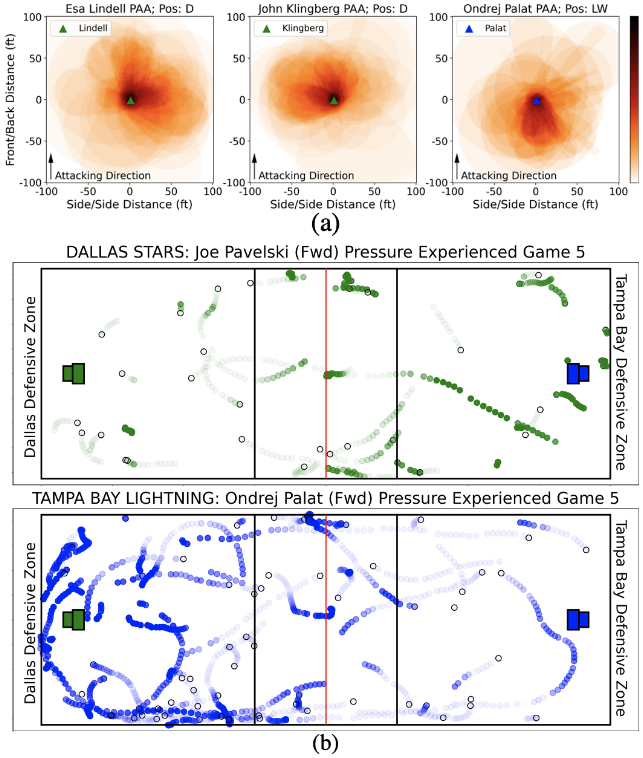


Figure 4: (a) PAA heatmaps, (b) Pressure maps.

below the dataset median for PAA are Volkov (TBL), Shattenkirk (TBL), Klingberg (DAL) and Benn (DAL), meaning smaller (possibly riskier) passing lanes. Note the differences in TOA for these players. This type of analysis could help adjust a player's game or anticipate opponents' actions.

### 6.2 Spatial Trends

**PAA Heatmaps:** We have constructed pass availability average (PAA) heatmaps for each player by overlaying the passing lane $PL(p, r, \gamma)$ oriented to their attacking direction for completed passes. In these heatmaps darker areas represent more passes being made in that direction due to more overlapping passing lanes and the variation in colours shows passing behaviour across positions and players. Figure 4a shows three heatmaps for two defensemen, Lindell (DAL) and Klingberg (DAL) who are often paired together and one forward, Palat (TBL). Each chart is coloured according to its own scale. We notice significantly different trends in passing across defense and forwards in our dataset. For example, Lindell and Klingberg tend to pass between each other and up-ice, whereas Palat makes most of his passes backwards. A larger dataset is needed for more general trends, however initial studies of passing lanes and direction help identity key differences in passing style due to position and player. Identifying trends becomes more important for coalition models pairing defensemen, three forwards, or a pairs of defense with forwards.

**Pressure Maps:** Another method of visualizing our metrics is to display them directly on a 2-dimensional rink, as shown in Figure 4b. The spatial aspect of understanding pressure is important when designing representative models of simulated player behaviour, which has been done in soccer [Lindström *et al.*, 2020]. Modeling how a player spatially handles pressure helps understand their playing style beyond when they move or shoot the puck. Figure 4b shows each possession by Joe Pavelski (DAL; top) and Ondrej Palat (TBL; bottom) throughout game five. Note that although teams change ends at the start of each period, we have adjusted these diagrams so they are relative to the direction of the opponents net. Each dot represents a timestep of possession beginning with hollow black circles. The shade in subsequent dots represents the experienced pressure normalized between 0 (transparent; low pressure) and 1 (dark; high pressure). This view suggests significant differences in playing style and handling pressure which could inform future simulations and models. Specifically, Pavelski tends to possess the puck for shorter durations on average than Palat, which we confirmed with an average possession time of 0.8 s compared to 1.13 s. Palat also tends to skate towards pressure when possessing the puck, whereas Pavelski doesn't carry the puck as long. These visualizations provide some preliminary analysis of how our analytics might be used to identify different playing styles. Further analysis into distributions, timelines, and play-by-play situations could further expose behavioural differences.

## 7 Discussion and Future Work

Although our analytics provide us with the possibility of identifying trends in player behaviour, it is important to emphasize that our dataset only includes two games due to data availability limitations. Therefore, we emphasize that these initial example insights may not be representative of a player's true longer-term behaviour. Furthermore, some of our metrics result in slightly defensive-bias results since forwards tended not to overtake defense often. While our methods allow for forwards to overcome this bias with larger rewards in OVT, we believe a complete representation of performance can be achieved when combined together with the current offensively-biased data. Additionally, our metrics are most effective when comparing players of the same position as their in-game situations would be more similar.

We expect the broader impact of our work to change the way AI and ice hockey work together, supporting an abundance of work around game theory, optimal coalition formation, and learning deep representations that are of broader interest to the AI community. Specifically, our metrics provide a new perspective of performance and we hope future models of performance can learn and validate which actions are valuable apart from sparsely recorded offense. Additionally, analytics can change the way players, coaches, general managers, and fans interact and understand the game, potentially helping to identify undervalued or overvalued players. By visualizing passing lanes and pressure for each timestep a player has the puck, new temporal and situational analysis of players' trends can be studied.

In conclusion, our contributions help capture attributes about every player in the game, not just those who generate offense. Future work involves improving AI in ice hockey with new player evaluation and roster management models and trying to quantify the elusive notion of momentum.

## Acknowledgments

We thank the Natural Sciences and Engineering Research Council of Canada (NSERC) for financial support. We thank The NHL, Rogers, Sportsnet, and the University of Waterloo for organizing the 2020 Sportsnet Hockey Hackathon: Powered by Rogers 5G where the first and fourth authors were given access to the tracking data. We thank Braden Hill, Rachel Li, Marhew Maradin, Navpreet Kaur, Ryan Trainor, Hanwen Zhang, and Paula He for their teamwork during the initial hackathon. The authors also broadly thank all coaches and support staff for their guidance throughout the authors' combined 75+ years of amateur, collegiate, and professional hockey playing experience.

## References

[Amin, 2018] Jiten B. Amin. Soccer player tracking system. *International Journal for Research in Applied Science and Engineering Technology*, 6:3455–3461, 2018.

[Andrienko *et al.*, 2017] G. Andrienko, N. Andrienko, G. Budziak, J. Dykes, G. Fuchs, T. V. Landesberger, and Hendrik Weber. Visual analysis of pressure in football. *Data Mining and Knowledge Discovery*, 31:1793–1839, 2017.

[Baysal and Sahin, 2016] S. Baysal and P. D. Sahin. Sentioscope: A soccer player tracking system using model field particles. *IEEE Transactions on Circuits and Systems for Video Technology*, 26:1350–1362, 2016.

[BusinessWire, 2018] BusinessWire. Global $4.5 billion sports analytics market forecasts up to 2024. https://www.businesswire.com/news/home/20181205005823/en/Global-4.5-Billion-Sports-Analytics-Market-Forecasts, 2018. Accessed: 2021-04-08.

[Dafoe *et al.*, 2020] A. Dafoe, Edward Hughes, Yoram Bachrach, T. Collins, Kevin R. McKee, Joel Z. Leibo, K. Larson, and T. Graepel. Open problems in cooperative AI. *ArXiv*, abs/2012.08630, 2020.

[Dafoe *et al.*, 2021] Allan Dafoe, Yoram Bachrach, Gillian Hadfield, Eric Horvitz, Kate Larson, and Thore Graepel. Cooperative AI: machines must learn to find common ground. *Nature*, 593:33–36, 2021.

[Douglas and Kennedy, 2019] A. Douglas and C. Kennedy. Tracking in-match movement demands using local positioning system in world-class men's ice hockey. *Journal of strength and conditioning research*, 2019.

[Fernández, 2019] J. Fernández. Decomposing the immeasurable sport: A deep learning expected possession value framework for soccer. 2019.

[Goes *et al.*, 2019] F. Goes, Matthias Kempe, L. Meerhoff, and K. Lemmink. Not every pass can be an assist: A data-driven model to measure pass effectiveness in professional soccer matches. *Big data*, 7 1:57–70, 2019.

[Guo *et al.*, 2020] Tianxiao Guo, Kuan Tao, Qingrui Hu, and Yanfei Shen. Detection of ice hockey players and teams via a two-phase cascaded cnn model. *IEEE Access*, 8:195062–195073, 2020.

[Hammond, 2011] Richard Hammond. Understanding advanced stats, part one: Corsi and Fenwick. https://www.matchsticksandgasoline.com/2011/7/29/2290643/understanding-advanced-stats-part-one-corsi-fenwick, 2011. Accessed: 2021-02-09.

[Kempe *et al.*, 2018] Matthias Kempe, Floris R. Goes, and Koen A. P. M. Lemmink. Smart data scouting in professional soccer: Evaluating passing performance based on position tracking data. 2018.

[Kirkpatrick and Radke, 1985] D. Kirkpatrick and J. D. Radke. A framework for computational morphology. *Machine Intelligence and Pattern Recognition*, 2:217–248, 1985.

[Lewis, 2003] M.M. Lewis. *Moneyball: The Art of Winning an Unfair Game*. Norton paperback. W.W. Norton, 2003.

[Lindström *et al.*, 2020] Per Lindström, Ludwig Jacobsson, Niklas Carlsson, and Patrick Lambrix. Predicting player trajectories in shot situations in soccer. In Ulf Brefeld, Jesse Davis, Jan Van Haaren, and Albrecht Zimmermann, editors, *Machine Learning and Data Mining for Sports Analytics*, pages 62–75, Cham, 2020. Springer International Publishing.

[Liu *et al.*, 2018] Guiliang Liu, Wang Zhu, and O. Schulte. Interpreting deep sports analytics: Valuing actions and players in the NHL. In *MLSA@PKDD/ECML*, 2018.

[Liu *et al.*, 2020] Guiliang Liu, O. Schulte, P. Poupart, Mike Rudd, and M. Javan. Learning agent representations for ice hockey. In *NeurIPS*, 2020.

[Ljung *et al.*, 2018] Dennis Ljung, N. Carlsson, and P. Lambrix. Player pairs valuation in ice hockey. In *MLSA@PKDD/ECML*, 2018.

[Luo *et al.*, 2020] Y. Luo, O. Schulte, and P. Poupart. Inverse reinforcement learning for team sports: Valuing actions and players. 2020.

[Macdonald, 2012] B. Macdonald. Adjusted plus-minus for NHL players using ridge regression with goals, shots, fenwick, and corsi. *Journal of Quantitative Analysis in Sports*, 8, 2012.

[Manafifard *et al.*, 2017] M. Manafifard, H. Ebadi, and H. Moghaddam. A survey on player tracking in soccer videos. *Comput. Vis. Image Underst.*, 159:19–46, 2017.

[Meerhoff *et al.*, 2019] L. Meerhoff, F. Goes, A. D. Leeuw, and A. Knobbe. Exploring successful team tactics in soccer tracking data. In *PKDD/ECML Workshops*, 2019.

[Mehrasa *et al.*, 2018] Nazanin Mehrasa, Yatao Zhong, F. Tung, L. Bornn, and G. Mori. Deep learning of player trajectory representations for team activity analysis. 2018.

[NHL, 2020] NHL. NHL league averages. https://www.hockey-reference.com/leagues/stats.html, 2020. Accessed: 2021-04-17.

[Rogers, 2021] Rogers. Rogers communications. https://www.rogers.com, 2021. Accessed: 2021-04-30.

[Schulte and Zhao, 2017] O. Schulte and Z. Zhao. Apples-to-apples: Clustering and ranking NHL players using location information and scoring impact. 2017.

[Silva *et al.*, 2018] R. M. Silva, J. Davis, and T. B. Swartz. The evaluation of pace of play in hockey. *Journal of Sports Analytics*, 4(2):145–151, 2018.

[SMT, 2021] SMT. SportsMEDIA technology. https://www.smt.com, 2021. Accessed: 2021-04-30.

[Sportsnet, 2021] Sportsnet. Sportsnet. https://www.sportsnet.ca, 2021. Accessed: 2021-04-30.

[Steiner *et al.*, 2019] Silvan Steiner, S. Rauh, M. Rumo, Karin Sonderegger, and R. Seiler. Outplaying opponents—a differential perspective on passes using position data. *German Journal of Exercise and Sport Research*, 49:140–149, 2019.

[Vats *et al.*, 2020] Kanav Vats, M. Fani, Pascale B. Walters, David A Clausi, and J. Zelek. Event detection in coarsely annotated sports videos via parallel multi receptive field 1d convolutions. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 3856–3865, 2020.

[Walkters *et al.*, 2020] Pascale Walkters, Mehrnaz Fani, David Clausi, and Alexander Wong. Bendernet and ringernet: Highly efficient line segmentation deep neural network architectures for ice rink localization. *Journal of Computational Vision and Imaging Systems*, 2020.

[Yan *et al.*, 2020] Tom Yan, Christian Kroer, and Alexander Peysakhovich. Evaluating and rewarding teamwork using cooperative game abstractions. *ArXiv*, abs/2006.09538, 2020.

[Yu *et al.*, 2019] David Yu, Christopher Boucher, Luke Bornn, and Mehrsan Javan. Playing fast not loose: Evaluationg team-level pace of play in ice hockey using spatio-temporal possession data. 2019.

[Zaccaro *et al.*, 2020] S. Zaccaro, Samantha Dubrow, E. Torres, and L. Campbell. Multiteam systems: An integrated review and comparison of different forms. 2020.