Edited by:
Patrick Lambrix
Mikael Vernblom
Niklas Carlsson
Tim Brecht

Linköping Hockey Analytics Conference
LINHAC 2025

Organizing Committee

Patrick Lambrix, Linköping University, Sweden (chair)
Tim Brecht, University of Waterloo, Canada (co-chair)
Niklas Carlsson, Linköping University, Sweden (co-chair)
Mikael Vernblom, Linköping Hockey Club, Sweden (co-chair)

Program Committee for the Research Track and Community Notes

Tim Brecht, University of Waterloo, Canada (chair)
Niklas Carlsson, Linköping University, Sweden (chair)
Martin Arlitt, University of Calgary, Canada
David A Clausi, University of Waterloo, Canada
Jack Davis, University of Waterloo, Canada
Ulf Johansson, Jönköping University, Sweden
Patrick Lambrix, Linköping University, Sweden
Erik Lignell, Frölunda HC, Sweden
Joshua Pohlkamp-Hartt, Boston Bruins, USA
David Thomas Radke, Chicago Blackhawks, USA
Mikael Svarén, Dalarna University, Sweden
Andrew C Thomas, Detroit Tigers, USA
Erik Wilderoth, Färjestad BK, Sweden

In cooperation with

E.H.C. Alliance of European Hockey Clubs

In cooperation with

Sportlogiq

Sponsor

The Swedish Research Council for Sport Science

Silver industry collaborators

BEARMIND

KINEXON SPORTS

Bronze industry collaborators

Preface

LINHAC 2025 took place June 2-4, 2025, and was organized by Linköping University and Linköping Hockey Club. LINHAC brought together professionals and academics with an interest in hockey analytics. It featured the latest research in hockey analytics in academia and companies, discussions with analysts and coaches, industry sessions with the latest hockey analytics products, and an analytics competition for students.

In addition to the research track and community notes, the program included talks by Andreas Hänni (49ing) and Olof Simander (Linköping University) on large language model interfaces to hockey databases, by Andreas Hänni (49ing) on analytics for officiating in the Swiss League, by Adam Almqvist Andersson and Johan Andersson (Swedish Ice Hockey Association) on examples of concrete hockey analytics, by Erik Wilderoth (Färjestad BK) with the title 'No room for stagnation: We need to get better', by Mikael Svarén (Dalarna University and Swedish Sports Confederation), on using analytics for performance optimization from an applied and academical perspective, by Robert Bandemer (d-fine) and Karl Schwarzenbrunner (German Ice Hockey Federation) on how data enriches the ice hockey game, by Curtis Harvey and Samuel Howorth (SMT) on skater and goalie workloads using player and puck tracking, and by Robin Blidstrand (Blidgency AB) on how to create stories and memories in sports and TV production.

Further, there were panel discussions moderated by Mike Helber. A first panel was made up of analysts from different European teams (Miika Arponen (Ässät Pori), Martin Lundholm (Skellefteå AIK), Karl Malmquist (Linköping Hockey Club), Erik Lignell (HC Fribourg-Gottéron SA)). The second panel discussed the state of the art and future of hockey analytics from the industry perspective (Tom Bertrand (Bearmind), Michael Elmer (KINEXON), Leo Girod (Sportcontract), Andreas Hänni (49ing), Freddy Sjögren (Freddie Sjögren Consulting AB), Morgan Zeba (Spiideo)). The third panel discussed sports analytics in use, where we looked broader than hockey (Tom Bertrand (Bearmind), Ola Lidmark Eriksson (Playmaker AI), Albin N Maelum (Stretch on Sense), Devin Pleuler (Maple Leaf Sports & Entertainment)). The participants of the fourth panel consisted of NHL analysts (AJ Bernstein (San Jose Sharks), Miranda McMillan (Montreal Canadiens), Josh Pohlkamp-Hartt (Boston Bruins), David Radke (Chicago Blackhawks), Matthew Hamann (Nashville Predators)). The final panel discussed the use of analytics in the media (Robin Blidstrand (Blidgency AB), Andreas Hänni (49ing), Mike Kelly (Sportlogiq, NHL Network)).

Our industry collaborators presented their products: KINEXON Sports, Bearmind, 49ing, Spiideo and Stretch on Sense.

This year's LINHAC had an invited session on football analytics. Football analytics has reached a higher level of maturity than hockey analytics and we invited representatives from academia, teams and companies. Jesse Davis (KU Leuven) talked about an effort to define a common data format for football match data. Pieter Robberechts (KU Leuven) from the same research group presented their work on evaluating sports analytics models: challenges, approaches and lessons learned. John Wall (Columbus Crew) discussed his use of analytics to create sustainable success with succession planning and purpose during his time as assistant coach with the Jamaican national team 2022-2024. Further, there was a talk by Ola Lidmark Eriksson on the Playmaker AI system.

Finally, there was a student competition where the task was to provide insights based on sequences of events in a hockey game. Data was provided by the SHL and Sportlogiq. A jury consisting of Tim Brecht (University of Waterloo), Hassaan Inayatali (Chicago Blackhawks), Patrick Lambrix (Linköping University), David Radke (Chicago Blackhawks) and Mikael Vernblom (Linköping Hockey Club) selected the project *Identifying and Analyzing SHL Ice Hockey Match Styles Based on Event Data Aggregation and K-Means Clustering* by Yanjie Lyu, Qingxuan Cui, Huaide Liu, Han Xia, and Yi Yang as the winner of the competition.

LINHAC is the only conference of its kind in Europe, and to our knowledge, it is the only hockey analytics conference that covers all aspects related to hockey analytics. This book includes the research track papers and community notes as well as contributions from industry, the student competition papers, and insights from contributors to LINHAC regarding their experience with hockey analytics and thoughts about its future. The research track papers and community notes are also published at Linköping Electronic Conference Proceedings as ECP 214, `https://doi.org/10.3384/ecp214`. The research track program committee with chairs Tim Brecht and Niklas Carlsson, selected the research paper *Ice Hockey Action Recognition via Contextual Priors* by Kseniia Buzko, Amir Nazemi, David A. Clausi and Yuhao Chen as the winner of the best research paper award.

We thank our moderator Mike Helber, our conference service Way, and the members of our local organization committee Mina Abd Nikooie Pour, Ying Li, Priyansh Gupta, Adriana Concha, Jenny Rydén, Veronica Gunnarsson, Lene Rosell, Anders Cronstierna, and Daniel Jemander, for their excellent support.

Last, but not least, we thank our collaborators the Alliance of European Hockey Clubs and Sportlogiq, our sponsor the Swedish Research Council for Sport Science, and our silver (Bearmind, KINEXON Sports), and bronze (49ing, Spiideo, Stretch on Sense) industry collaborators.

September 2025

Patrick Lambrix (chair),
Tim Brecht (co-chair),
Niklas Carlsson (co-chair),
Mikael Vernblom (co-chair)

Contents

Research papers

Ice Hockey Action Recognition via Contextual Priors

Kseniia Buzko[1 [0009−0009−9930−8497]], Amir Nazemi[1 [0000−0002−8405−473X]], David A. Clausi[1 [0000−0002−6383−0875]], and Yuhao Chen[1 [0000−0001−6094−0545]]

Department of Systems Design Engineering, University of Waterloo, Waterloo, ON N2L 3G1, Canada
{kbuzko,amir.nazemi,dclausi,yuhao.chen1@}@uwaterloo.ca

Abstract. Skeleton-based action recognition models, which are developed for generic human-pose data, struggle with ice-hockey broadcasts player action recognition, where the players appear smaller, move abruptly, and wield sticks that are invisible to standard skeleton models. To address these issues, we propose CP-Hockey, a context-aware pipeline that incorporates two domain-specific priors. First, a temporal player's bounding-box normalization stabilizes player scale across the player tracklet, raising top-1 accuracy from 31 % to 57 % on a six-class NHL dataset. Second, we design hockey-specific skeletons that include stick end-points and optional detailed head landmarks. A 15-keypoint body-plus-stick model improves the accuracy to 64 %, while our full 20-keypoint configuration reaches 65 %. Experimental results with STGCN++ and 2s-AGCN show that both contextual priors are necessary: scale normalization reduces spatial jitter, and stick keypoints disambiguate visually similar movements such as stickwork versus striking a puck with a stick. CP-Hockey establishes a strong baseline for fine-grained ice-hockey analytics and provides a blueprint for adapting skeleton pipelines to other equipment-centric sports.

Keywords: Ice Hockey Action Recognition · Skeleton-based Action Recognition · Contextual Priors

1 Introduction

Player action recognition is a fundamental task in sports analytics, enabling the automatic identification and classification of particular player movements during the game, such as skating, striking a puck, or maintaining position. Understanding these actions allows deeper insights into player strategies and game analysis. Although sports such as basketball [2, 5] and soccer [3, 18] have seen extensive development in player action recognition methodologies due to their popularity, ice hockey remains comparatively underexplored. Traditional approaches like the Action Recognition Hourglass Network (ARHN) [1] rely mainly on static pose input, ignoring temporal and ice hockey contextual information. By 'contextual information' of broadcast ice hockey videos, we point to structured domain-specific cues such as camera-driven scale changes and stick pose that go beyond

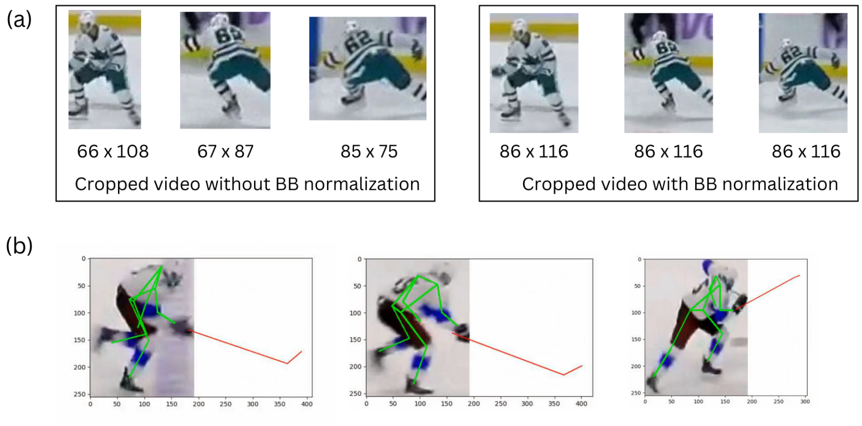

Fig. 1. Illustration of the two contextual priors exploited by CP-Hockey. (a) Temporal bounding-box (BB) normalization suppresses scale jitter by resizing every frame in the player tracklet to the maximum BB dimensions. (b) A 15-keypoint ice hockey skeleton (green body joints, red stick joints) with the stick pose which is vital for player action recognition and not included in the generic 17-keypoint COCO skeleton.

the joint coordinates; the two contextual priors that we leverage in this work are visualized in Figure 1.

Player action recognition in broadcast ice hockey footage is challenging for several reasons. The players bounding boxes are small and fast-moving (Figure 2); rapid camera pans create scale jitter; bulky protective gear and frequent occlusions confound generic pose estimators.

Most pose-based solutions for action recognition ignore the temporal contextual features of the data domain [22, 7] and rely on generic 17-keypoint human skeletons (such as keypoint COCO model [14]). However, static cropping and resizing of a player's bounding box on a frame-by-frame basis disregards the temporal context, often resulting in irregular player scales and spatial jitter across consecutive frames. This jitteriness of the players' poses decreases the performance of downstream models, as the player's pose rapidly changes in size or position due to camera effects rather than actual movement of the player's poses. Although generic skeletal configurations are widely used in human pose estimation [14], they are not tailored for the ice hockey domain. Many of the keypoints in these models correspond to body joints that can be easily occluded by hockey gear like helmets or pads. In addition, they do not consider key objects such as the hockey stick in their solution. From the perspective of contextual priors in ice hockey, we argue that the hockey stick serves as a vital cue to recognize actions in the sport. It can differentiate various actions, such as stickwork, where the stick primarily makes contact with the ice, and striking the puck, where the stick is in a winding-up motion.

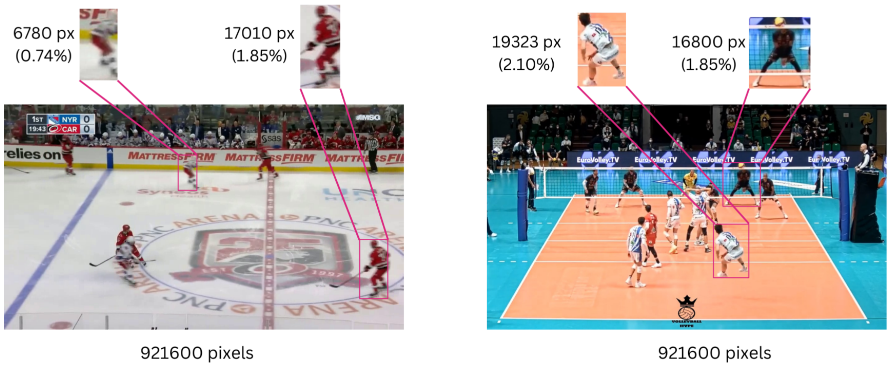

Fig. 2. Comparison of bounding-box pixel sizes between ice hockey (left) and volleyball (right) broadcast footage [4]. In addition to the rapid movement of players and the camera, the bounding boxes of players are relatively smaller in ice hockey broadcast videos, which makes player action recognition more complicated.

In this work, we propose CP-Hockey, a novel ice hockey action recognition pipeline that leverages contextual priors to address the mentioned challenges. Our CP-Hockey pipeline benefits from contextual priors at two points (Figure 1). First, a context-sensitive temporal normalization stabilizes the bounding box of each player's tracklet throughout the action clip, suppressing camera-driven scale variability. Second, our solution benefits from skeleton with stick end-points, giving the action recognition model explicit access to stick pose and motion. Through extensive experiments on an NHL hockey video dataset, we demonstrate that integrating such contextual priors markedly improves player action recognition performance in ice hockey broadcast videos.

This paper makes the following unique and novel contributions to player action recognition in ice hockey broadcast videos:

- We implement a temporal bounding box normalization method that reduces spatial jitter of a player's tracklet in broadcast footage and improves the skeleton-based player action recognition in ice hockey broadcast videos.
- We implement a novel skeletal configuration specifically tailored to the ice hockey domain, which integrates stick keypoints to capture important stick-related interactions, an aspect that to the best of our knowledge, has not been previously explored.
- We perform extensive evaluations comparing graph convolutional networks (GCNs) such as 2s-AGCN [9] and STGCN++ [6], generating empirical evidence advocating for the use of STGCN++ with contextual priors in recognizing ice hockey-specific actions.

2 Background Research

2.1 Pose Estimation

Human pose estimation models identify key anatomical points, such as joints and limbs, from images or videos, providing a structured representation of human posture. Pose estimation is particularly relevant in sports analytics, as it enhances the ability to capture precise body configurations necessary for effectively classifying specific movements. In ice hockey, accurate pose estimation can significantly improve action recognition by providing robust temporal skeletal representations of players' complex movements.

However, standard human pose estimation methods, typically trained on general datasets such as COCO [14] or MPII [15], face unique difficulties when applied to ice hockey broadcasts. Players wear bulky protective equipment that alters the shape of the body, leading to frequent misinterpretation of limb positions. In addition, uniform colors, such as white jerseys, visually blend with ice or boards, complicating limb identification. Pose estimation in ice hockey is further challenged by rapid and agile player movements, such as quick skating transitions or sudden turns, which often result in severe motion blur and unusual body poses [1, 10].

Recent work has begun to address these challenges associated with ice hockey-specific poses. Balaji et al. [11] introduced a multimodal approach that utilizes language cues to manage occluded keypoints, focusing on the player's body and using textual prompts for expected stick positions. This method significantly improved pose accuracy in an ice-hockey dataset by leveraging domain context, such as equipment knowledge. We incorporate pose estimation [10, 11] into our CP-Hockey pipeline, as accurate skeletal representations are essential for reliable ice hockey action recognition, especially for subtle actions like puck striking and stick handling.

2.2 Skeleton-Based Action Recognition

Building on pose estimation, researchers have explored skeleton-based action recognition for sports, where player joint sequences are input to an action classifier. Graph convolutional networks (GCNs) have become a dominant paradigm for this problem [22], as they naturally model the human skeleton as a graph, where each vertex represents a body joint and edges are split into temporal and spatial connections. The edges that join vertices inside a frame are referred to as spatial edges, while the edges that join the same vertex across consecutive frames are referred to as temporal edges. This representation significantly outperformed convolutional neural networks (CNNs) and recurrent neural networks (RNNs) based approaches [16, 17] that employed the skeleton modality without taking joint dependence into account.

ST-GCN [8] demonstrated the effectiveness of spatio-temporal keypoint graphs for action recognition, outperforming purely appearance-driven approaches. Subsequent improvements, such as 2s-AGCN [9] and STGCN++ [6], introduced adaptive graph structures and multi-stream inputs, such as joint and bone data, for improved performance. These models have proven particularly effective in sports analytics [23], where structured skeletal data remains reliable despite variations in clothing, lighting, and background clutter.

Despite advancements, action recognition in ice hockey remains underexplored compared to sports with a larger following, such as basketball or soccer. Prior research has focused mainly on coarse tasks such as player tracking [19] and puck tracking [20], often relying on traditional vision pipelines or CNN-based trackers. For particular action recognition, such as forward skating, stickwork, and rapid deceleration, remains underexplored.

2.3 Contextual and Object-Aware Action Recognition

Early skeleton pipelines inferred actions almost exclusively from joint coordinates, treating each frame or short clip in isolation. Recent work shows that contextual priors like temporal neighborhoods, interacting objects, and high-level semantics, can greatly sharpen recognition, especially when skeleton trajectories alone are ambiguous. For example, Cioppa et al. design CALF [12], a context-aware loss that spots soccer events by supervising not the single annotated frame but a short temporal window around it. While Wen et al. add dynamically detected object centers to the ST-VGCN graph [7], demonstrating that even coarse object localization can disambiguate actions sharing near-identical limb motions.

3 Dataset

To develop and evaluate our CP-Hockey pipeline, we collected video clips from National Hockey League (NHL) broadcasts captured at 30 frames per second (fps). These clips were sourced from multiple NHL games, covering 29 different teams, which provides a wide range of arenas, lighting conditions, and team uniforms for a diverse and challenging dataset. The raw footage was segmented into individual shots, where each clip lasts from a few seconds to over a minute.

Six ice-hockey-specific classes were selected for their frequency, tactical importance, and visual distinctness. Table 1 summarizes the class definitions and their distribution. In total, the dataset contains 1,547 annotated action instances, each spanning 2 seconds (60 frames), where the start and end frames are defined as ± 30 frames around the anchor action frame.

4 Methodology

4.1 Pipeline Overview

The CP-Hockey pipeline for ice-hockey action recognition comprises four stages that together inject two complementary contextual priors: a temporal prior that stabilizes player scale and position, and an object prior that models the hockey stick explicitly. Figure 3 gives a high-level view of the proposed pipeline.

Table 1. Action classes with descriptions and frequency.

Action Class	**Description**	**Occurrences**
Rapid Deceleration	Player abruptly reduces skating speed to change direction or avoid contact	52
Backwards Skating	Player skates in reverse, typically to maintain defensive positioning	116
Maintaining Position	Player remains largely stationary while reading the play or screening the goalie	164
Strike Puck with Stick	Player's stick makes direct contact with the puck for a pass, shot, or block	337
Forwards Skating	Player accelerates toward the offensive zone while controlling balance and speed	421
Stickwork	Player performs intricate stickhandling to retain or regain puck control	457
Total		**1547**

Each annotated action segment is first preprocessed by automatically extracting a player tracklet with the VIP-HTD tracker [19]. Bounding boxes are then temporally normalized to enforce scale consistency (Section 4.3, Contextual Prior #1). Next, an ice-hockey-specific pose estimator [10, 11] converts every normalized crop to a skeleton, yielding 2-D key-point sequences with six alternative configurations (Section 4.4, Contextual Prior #2). Finally, a GCN-based model classifies the action (Section 4.5).

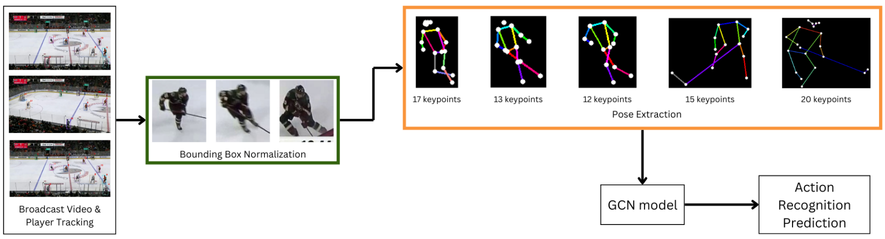

Fig. 3. Overview of the CP-Hockey pipeline. Contextual Prior #1 (green) stabilizes bounding-box scale; Prior #2 (orange) augments the skeleton with stick key-points. The enriched graph is fed to GCN-based model for action recognition.

4.2 Preprocessing

To automate bounding box extraction and ensure temporal consistency across frames, we leverage a tracking approach [19], specifically tailored for player detection in ice hockey broadcast footage. After extracting bounding boxes, frame-wise bounding box cropping and resizing to a normalized spatial representation is performed. These cropped bounding boxes are the inputs for pose extraction, which uses ice hockey-optimized models presented in [10, 11].

4.3 Bounding Box Normalization

A significant challenge in using broadcast video for pose estimation is the variability in player size and position across frames. A naive approach might resize each frame independently, stretching the player's skeleton and distorting the body shape, which negatively impacts downstream performance. To mitigate this, we propose a robust bounding box normalization strategy. Specifically, we compute the maximum width and height of the bounding box for each annotated action across all frames within the action segment. We then re-extract the bounding box from the original video, expanding each frame's bounding box to consistently match these maximum dimensions, ensuring the player remains centered.

To illustrate the effectiveness of this approach, Figure 4 compares keypoint trajectories for a sample data sequence without and with bounding box normalization. On the left, the raw joint trajectories without normalization are scattered, showing significant variability due to changes. In contrast, the trajectories on the right, after applying bounding box normalization, are significantly more stable and coherent.

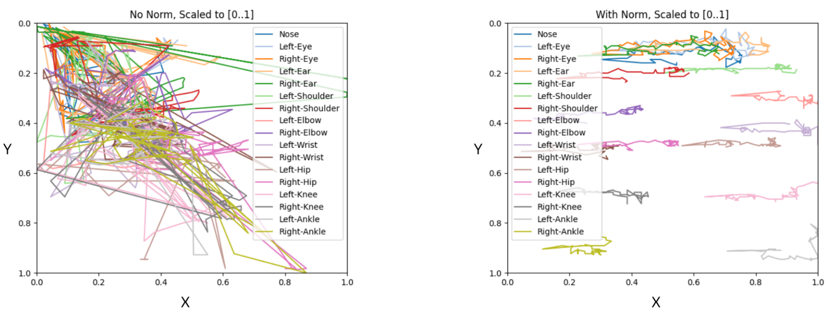

Fig. 4. Impact of bounding box normalization on joint trajectories for one data sample. Left: Without bounding box normalization, the joint trajectories are scattered and inconsistent due to variability in player scale. Right: After applying bounding box normalization, trajectories appear significantly more stable and coherent, clearly reducing positional variance and making movement patterns easier to interpret.

4.4 Pose Keypoint Extraction

We apply ice hockey-optimized 2D pose estimators [10, 11] to generate skeletal keypoints from players in normalized bounding boxes belonging to a sequence. We experimented with various keypoint configurations (Figure 5):

1. **17 Keypoints (COCO [14]):** Standard full-body representation.
2. **13 Keypoints:** Merged head keypoints.
3. **12 Keypoints:** Head keypoints omitted to focus solely on torso and limbs.
4. **15 Keypoints:** 12 body keypoints plus 3 ice hockey stick keypoints (butt end, heel, toe).
5. **20 Keypoints:** Comprehensive setup with 5 head, 12 body, and 3 stick keypoints.
6. **3 Stick Keypoints Only:** Focused solely on stick interaction.

Fig. 5. Five skeletal configurations used in this study, from left to right: 17-keypoint COCO, 13-keypoint merged head, 12-keypoint head-omitted, 15-keypoint with ice hockey stick endpoints, 20-keypoint comprehensive configuration.

4.5 GCN for Action Recognition

We employ the 2s-AGCN [9] and STGCN++ [6] architectures to classify skeleton-based action sequences. Training and evaluation leverage the mmAction2 toolbox [21]. Our approach exclusively utilizes the joint modality (keypoint coordinates) to capture spatial-temporal dynamics. This choice simplifies model inputs while effectively capturing both static postures and dynamic movements necessary for accurate ice hockey action recognition.

5 Experiments

5.1 Experimental Setup

To ensure a robust model evaluation, we partitioned the dataset described in Section 3 into training and validation sets using a 70%-30% stratified split, maintaining an approximately proportional representation of each action class. We evaluate action recognition performance using standard metrics.

We adopt the mmAction2 toolbox [21] for training and evaluation. The training process leverages the GCN-based architectures (described in Section 4.5) with cross-entropy loss. Experiments were conducted on an NVIDIA RTX 3090 Ti GPU. The training setup for our GCN experiments is as follows: a learning rate of 0.1, weight decay of 5×10^{-4}, batch size of 16, and training for a total of 20 epochs using stochastic gradient descent (SGD) as the optimizer.

5.2 Experimental Results

We evaluated our ice hockey action recognition framework using exclusively the joint modality, focusing clearly on different skeletal configurations and their impact on performance. Results across these configurations are presented in Table 2.

Table 2. Comparison of action recognition accuracy using joint modality across different skeletal configurations and models.

Model	Keypoint Configuration	Accuracy (%)
STGCN++	17 Keypoints (No Bounding Box Norm)	31
	12 Keypoints	40
	13 Keypoints	44
	17 Keypoints	57
	15 Keypoints (12 body + 3 stick)	<u>64</u>
	20 Keypoints (5 head + 12 body + 3 stick)	**65**
	3 Stick Keypoints Only	52
2s-AGCN	17 Keypoints	51
	15 Keypoints (12 body + 3 stick)	56

The 20-keypoint skeleton with STGCN++ attains the highest 65 % accuracy. Using the same model, bounding-box normalization plus the standard 17-keypoint human skeleton yields 57 %. Adding only three stick joints to that baseline raises accuracy by 7 pp to 64 %, and further appending five fine-grained head landmarks adds another 1 pp, reaching 65 %. Conversely, dropping head landmarks altogether (13 or 12 keypoints) lowers accuracy to 44 % and 40 %, respectively. These step-wise gains and losses show that both detailed head cues and explicit stick endpoints are essential for recognizing hockey-specific actions. Additionally, results from the 3 stick keypoints alone configuration (52%) motivate the significant role of stick interactions. STGCN++ consistently outperformed 2s-AGCN, validating our selection of STGCN++ as the core model.

5.3 Qualitative Results

Figure 6 illustrates qualitative predictions from our models on six representative action clips. Each image captures the midpoint frame of a two-second action sequence along with ground-truth labels and predictions from five experimental setups: (1) without bounding box normalization, (2) 2s-AGCN with 17 keypoints, (3) STGCN++ with 15 keypoints, (4) STGCN++ with 17 keypoints, and (5) STGCN++ with our optimized 20-keypoint configuration.

The model without bounding box normalization often misclassifies subtle actions, mistaking stick interactions for general stickwork due to spatial inconsistencies. While the 2s-AGCN model with 17 keypoints struggles to differentiate stick actions from skating movements, our 15-keypoint and 17-keypoint STGCN++ configurations significantly improve accuracy. However, the 20-keypoint

Ground Truth	Strike puck with stick	Forwards skating	Backwards skating	Maintaining position	Rapid deceleration	Stickwork
Our Prediction (w/o bb norm)	Stickwork	Stickwork	Stickwork	Strike puck with stick	Stickwork	Stickwork
Our Prediction (2s-AGCN 17 kps)	Stickwork	Forwards skating	Strike puck with stick	Maintaining position	Stickwork	Strike puck with stick
Our Prediction (STGCN++ 15 kps)	Strike puck with stick	Forwards skating	Forwards skating	Maintaining position	Stickwork	Stickwork
Our Prediction (STGCN++ 17 kps)	Strike puck with stick	Forwards skating	Stickwork	Strike puck with stick	Stickwork	Stickwork
Our Prediction (STGCN++ 20 kps)	Strike puck with stick	Forwards skating	Backwards skating	Maintaining position	Stickwork	Stickwork

Fig. 6. Qualitative comparison of ice hockey action recognition predictions across different model setups on six action clips from our dataset. Ground-truth labels are indicated in green (correct predictions), while incorrect predictions are shown in red. The optimized STGCN++ model with 20 keypoints demonstrates improved accuracy, effectively differentiating subtle actions compared to models without bounding box normalization or fewer keypoints.

STGCN++ model excels, effectively distinguishing visually similar actions such as forward skating versus backward skating, and reliably differentiating "strike puck with stick" from general stickwork. These findings highlight the importance of incorporating detailed head and stick keypoints, along with effective bounding box normalization, to enhance ice hockey action recognition accuracy.

6 Conclusion

In this paper, we introduced CP-Hockey, a pipeline that injects two complementary contextual priors: a temporal prior that stabilizes player scale and an object prior that models the hockey stick into the skeleton-based action-recognition stack for broadcast ice-hockey footage. By counteracting camera-induced scale jitter and making stick pose explicit, CP-Hockey tackles the key challenges of tiny, fast-moving players, heavy occlusion from protective gear, and subtle stick-centric motions.

Exploiting the temporal prior alone, our bounding-box normalization lifts top-1 accuracy from 31 % to 57 % with the general 17-keypoint skeleton. Adding the object prior (three stick end-points) further raises accuracy to 64 %. Our most expressive 20-keypoint configuration, which also restores detailed head landmarks, tops out at 65 %. These gains confirm that action recognition benefits not merely from more keypoints but from the right context-aware keypoints, aligned with the physics and semantics of ice hockey.

Although CP-Hockey demonstrates strong performance on pre-segmented action clips, extending the method to the action spotting [24] task would be challenging. Future work should focus on developing mechanisms for automatic action spotting in continuous video streams. Integrating multi-view or depth information may also enable more accurate 3D pose estimation. Additionally, incorporating contextual signals such as puck tracking and player interactions could further enhance recognition of complex, team-based plays.

Overall, our results demonstrate that combining robust spatial normalization, sport-specific skeletal representations, and spatio-temporal graph models substantially improves action recognition in visually challenging sports scenarios. Validated on real NHL broadcast footage, CP-Hockey lays a strong foundation for advanced ice hockey analytics and provides a blueprint for extending action recognition methods to other demanding sports domains.

Acknowledgments. This study is funded by Stathletes Inc and NSERC (Natural Sciences and Engineering Research Council) Alliance program.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Fani, M., Neher, H., Clausi, D.A., Wong, A., Zelek, J.: Hockey Action Recognition via Integrated Stacked Hourglass Network. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pp. 85–93 (2017). https://doi.org/10.1109/CVPRW.2017.17
2. Cheng, X., Li, Z.: Research on Basketball Shooting Action Recognition and Optimization System Based on Deep Learning. In: ICMIII, Melbourne, Australia, pp. 546–551 (2024). https://doi.org/10.1109/ICMIII62623.2024.00108
3. Zhu, H., Liang, J., Lin, C., Zhang, J., Hu, J.: A Transformer-based System for Action Spotting in Soccer Videos. In: ACMMM, pp. 103–109 (2022). https://doi.org/10.1145/3552437.3555693
4. Klyukin, M.: volleyball_dataset: An Open Source Volleyball Action Recognition Dataset. Roboflow Universe (2022). Available at: https://universe.roboflow.com/mikhail-klyukin/volleyball_dataset
5. Khobdeh, S.B., Yamaghani, M.R., Sareshkeh, S.K.: Basketball Action Recognition Based on the Combination of YOLO and a Deep Fuzzy LSTM Network. J. Supercomput. **80**, 3528–3553 (2024). https://doi.org/10.1007/s11227-023-05611-7. Available at: https://doi.org/10.1007/s11227-023-05611-7
6. Duan, H., Wang, J., Chen, K., Lin, D.: Pyskl: Towards Good Practices for Skeleton Action Recognition. In: Proceedings of the 30th ACM International Conference on Multimedia, pp. 7351–7354 (2022).
7. Wen, H., Lu, Z., Shen, F., et al.: Improving Skeleton-Based Action Recognition with Interactive Object Information. Int. J. Multimed. Info. Retr.
8. Yu, B., Yin, H., Zhu, Z.: Spatio-Temporal Graph Convolutional Networks: A Deep Learning Framework for Traffic Forecasting. In: Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, pp. 3634–3640 (2018). https://doi.org/10.24963/ijcai.2018/505

9. Shi, L., Zhang, Y., Cheng, J., Lu, H.: Two-Stream Adaptive Graph Convolutional Networks for Skeleton-Based Action Recognition. In: Proceedings of CVPR, pp. 12018–12027 (2019). https://doi.org/10.1109/CVPR.2019.01230
10. Balaji, B., Clausi, D.A.: Towards Agile Human Pose Estimation: A Benchmark Study in Ice Hockey Analytics. JCVI 9(1), 54–57 (2024). https://doi.org/10.15353/jcvis.v9i1.10014
11. Balaji, B., Bright, J., Chen, Y., Rambhatla, S., Zelek, J., Clausi, D.A.: Seeing Beyond the Crop: Using Language Priors for Out-of-Bounding Box Keypoint Prediction. In: NIPS, pp. 102897–102918 (2024).
12. Cioppa, A., Deliège, A., Giancola, S., Ghanem, B., Van Droogenbroeck, M., Gade, R., Moeslund, T.B.: A Context-Aware Loss Function for Action Spotting in Soccer Videos. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR), (2020).
13. Shahroudy, A., Liu, J., Ng, T.-T., Wang, G.: NTU RGB+D: A Large Scale Dataset for 3D Human Activity Analysis. arXiv:1604.02808 (2016)
14. Lin, T.-Y., Maire, M., Belongie, S., Bourdev, L., Girshick, R., Hays, J., Perona, P., Ramanan, D., Zitnick, C.L., Dollár, P.: Microsoft COCO: Common Objects in Context. arXiv:1405.0312 (2015). Available at: https://arxiv.org/abs/1405.0312
15. Andriluka, M., Pishchulin, L., Gehler, P., Schiele, B.: 2D Human Pose Estimation: New Benchmark and State of the Art Analysis. In: 2014 IEEE Conference on Computer Vision and Pattern Recognition, pp. 3686–3693 (2014). https://doi.org/10.1109/CVPR.2014.471
16. Hou, R., Chen, C., Shah, M.: Tube Convolutional Neural Network (T-CNN) for Action Detection in Videos. In: 2017 IEEE International Conference on Computer Vision (ICCV), pp. 5823–5832 (2017). https://doi.org/10.1109/ICCV.2017.620
17. Krishnan, K., Prabhu, N., Babu, R.V.: ARRNET: Action Recognition through Recurrent Neural Networks. In: 2016 International Conference on Signal Processing and Communications (SPCOM), pp. 1–5 (2016). https://doi.org/10.1109/SPCOM.2016.7746614
18. Sen, A., Minhaz Hossain, S.M., Ashraf Russo, M., Deb, K., Jo, K.H.: Fine-Grained Soccer Actions Classification Using Deep Neural Network. In: 2022 15th International Conference on Human System Interaction (HSI), pp. 1–6 (2022). https://doi.org/10.1109/HSI55341.2022.9869480
19. Prakash, H., Chen, Y., Rambhatla, S., Clausi, D.A., Zelek, J.: VIP-HTD: A Public Benchmark for Multi-Player Tracking in Ice Hockey. Journal of Computational Vision and Imaging Systems **9**(1), 22–25 (2024). https://doi.org/10.15353/jcvis.v9i1.10006. Available at: https://openjournals.uwaterloo.ca/index.php/vsl/article/view/5858
20. Li, M., Hu, H., Yan, H.: Ice Hockey Puck Tracking through Broadcast Video. Neurocomputing **551**, 126484 (2023). https://doi.org/10.1016/j.neucom.2023.126484. Available at: https://www.sciencedirect.com/science/article/pii/S0925231223006070
21. MMAction2 Contributors: OpenMMLab's Next Generation Video Understanding Toolbox and Benchmark. Github repository (2020). Available at: https://github.com/open-mmlab/mmaction2
22. Yin, H., Sinnott, R.O., Jayaputera, G.T.: A Survey of Video-Based Human Action Recognition in Team Sports. Artif. Intell. Rev. **57**, 293 (2024). https://doi.org/10.1007/s10462-024-10934-9
23. Wei, J., Yu, B., Zhang, H., Liu, J.: Skeleton Based Graph Convolutional Network Method for Action Recognition in Sports: A Review. In: 2023 8th IEEE International

Conference on Network Intelligence and Digital Content (IC-NIDC), Beijing, China, pp. 66–70 (2023). https://doi.org/10.1109/IC-NIDC59918.2023.10390711
24. Denize, J., Liashuha, M., Rabarisoa, J., Orcesi, A., Héralt, R.: COMEDIAN: Self-Supervised Learning and Knowledge Distillation for Action Spotting Using Transformers. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) Workshops, pp. 530–540 (2024).

A Gaussian Mixture Model Approach for Characterizing Playing Styles of Ice Hockey Players

Priyansh Gupta[1], Mikael Vernblom[2], Niklas Carlsson[1], and Patrick Lambrix[1]

[1] Linköping University, Sweden
[2] Linköping Hockey Club, Sweden

Abstract. Player categorization based on playing style is a highly important task in professional ice hockey, aiding scouting, player development, and strategic decision-making. Traditional methods often rely on simple metrics like goals or assists, which fail to capture the full complexity of a player's style and contributions. Motivated by the increasing availability of detailed event data and advances in machine learning based modeling techniques, this paper explores a richer, data-driven approach to player categorization. We build on recent work in player vector representations and apply Gaussian Mixture Models (GMMs) to cluster forwards and defenders based on event data from five seasons of the Swedish Hockey League (SHL). Our contributions are threefold: (1) we construct detailed player vectors that summarize a wide range of offensive and defensive skills, (2) we apply GMMs to identify soft clusters of players, allowing for nuanced overlapping playing styles, and (3) we analyze the resulting clusters to interpret distinct player profiles and provide concrete examples. Our results offer a more flexible and realistic view of player roles, reflecting the continuous and multi-dimensional nature of playing styles. The approach helps enhance talent evaluation and roster building, and offers an efficient framework for future analyses across leagues and seasons.

1 Introduction

Player categorization based on playing style is an important task in professional ice hockey, supporting scouting, player development, and strategic decision-making. Traditional approaches typically rely on discrete performance metrics, such as goals, assists, or shots, offering only a partial view of a player's overall style and contribution. More recently, increased event data collection and advances in modeling techniques have opened up new possibilities for representing and analyzing player behaviors in more nuanced ways.

In this paper, we build upon recent developments in player vector representations [17] and apply Gaussian Mixture Models (GMMs) to identify clusters of forwards and defenders based on their playing styles. GMMs offer a probabilistic soft clustering approach that is particularly well-suited to model the continuous and overlapping nature of player styles. Unlike hard clustering methods, which assign each player to a single cluster, GMMs allow players to belong to multiple clusters with varying degrees of membership, reflecting the reality that players often exhibit characteristics of multiple styles.

The specific contributions of this paper are threefold. First, we leverage detailed event data from five seasons of the Swedish Hockey League (SHL) to construct player vectors capturing a wide range of offensive and defensive skills. Second, we apply GMMs to these vectors, determining the number of clusters using model selection criteria such as the Bayesian Information Criterion (BIC). Finally, we analyze the resulting clusters to interpret the different playing styles represented among forwards and defenders, and provide examples of players associated with each style.

Organization: Section 2 provides background on finite mixture models, Gaussian mixture models, and model selection methods. Section 3 reviews related work in player evaluation and categorization. Section 4 describes the dataset used in this study. Section 5 outlines our methodology for constructing player vectors and fitting GMMs. Section 6 presents the clustering results and analyzes the identified playing styles, before Section 7 concludes the paper.

2 Background

In this paper, we apply Gaussian Mixture Models (GMMs) to cluster forwards and defenders based on their playing styles. A GMM is a type of Finite Mixture Model (FMM) where each component is a Gaussian distribution. GMMs are particularly well-suited for player data, as different playing styles often overlap and evolve along continuous spectrums, making soft clustering approaches like GMMs more appropriate than hard clustering alternatives.

An FMM models data from a combination of unobserved groups, without knowing in advance which point belongs to which group. With FMMs, each group is associated with its own probability distribution, and the overall dataset is modeled as a weighted sum of these components. Instead of trying to fit just one model to the entire dataset, this allows the FMMs to fit multiple smaller models and combine them [16]. This offers a flexible framework that better captures complex data structures than single-model approaches.

Although various distributions, such as Poisson and Bernoulli, can be used within the FMM framework, the choice of Gaussian distributions in GMMs provides two key advantages: flexibility and interpretability. Gaussian components can model elliptical clusters with different orientations and scales, which is important when different playing styles vary along distinct combinations of performance features. Furthermore, GMMs naturally produce soft assignments of players to clusters, reflecting the intuition that playing styles often exist on a continuum rather than falling into rigid categories.

Formally, an FMM assumes that each observation y (in this study, a player vector representing various skills for the player) comes from one of g different groups (components), each described by its own distribution. These components are mixed using probabilities $\pi_1, \pi_2, ..., \pi_g$, where each π_i is the mixing proportion for the i-th component. These values are all positive and add up to 1. The overall distribution of an FMM can either be a probability density function (PDF), if the data is continuous, or a probability mass function (PMF) in the case of a discrete dataset [2,15]. More specifically, the PDF or PMF of the mixture model is represented as follows:

$$f(y) = \sum_{i=1}^{g} \pi_i f_i(y), \tag{1}$$

where $f_i(y)$ represents the PDF or PMF of the i-th component, π_i represents the mixing proportion of the component, and g is the total number of components [16]. Depending on the type of distribution, the component densities $f_i(y)$ are represented as $f_i(y, \theta_i)$, where θ_i is the vector of unknown parameters for the i-th component density. In the case of a Gaussian distribution, these parameters are the mean and variance $\theta_i = [\mu_i, \sigma_i^2]$, resulting in the following way to represent the PDF or PMF of the mixture model:

$$f(y, \Psi) = \sum_{i=1}^{g} \pi_i f_i(y, \theta_i), \tag{2}$$

where y is the data we want to model, and Ψ is a vector containing all the unknown parameters in the mixture model, such as π_i and θ_i for i-th component. It can be expressed as $\Psi = (\pi_1, ..., \pi_{g-1}, \zeta^T)^T$, where the parameter ζ includes the parameters of the selected distributions for all g components [15]. In the context of this study, y refers to a player vector that characterizes an individual player's style, and the parameters Ψ collectively describe how these playing styles are distributed across the population.

To estimate the values of the parameters of a FFM, i.e., the mixing proportion π and the parameters of each component distribution θ, several approaches exist, but the most commonly used is the Expectation-Maximization (EM) algorithm. EM applies maximum likelihood estimation to fit the FMM. The EM algorithm consists of two main steps: the Expectation (E-step), which calculates the probability that each data point belongs to each component based on the current parameter estimates, and the Maximization (M-step), which updates the parameter estimates using these probabilities to better fit the data. These two steps are repeated iteratively until convergence is achieved [15].

A key challenge in applying FMM is identifying the value of g, i.e., the number of components in the model. Lower values of g may lead to underfitting, while higher values can result in overfitting [15]. To address this, several model selection criteria have been developed, aiming to balance model fit with complexity. Two widely used criteria for model selection are the Akaike Information Criterion (AIC) [1] and the Bayesian Information Criterion (BIC) [27]. Both criteria evaluate models based on the maximized likelihood $\hat{L}$, while introducing a penalty that increases with the number of estimated parameters $|\Psi|$; thus, discouraging overfitting. The AIC provides an estimator of the relative prediction error, calculated as follows,

$$\mathrm{AIC} = 2|\Psi| - 2\ln(\hat{L}), \tag{3}$$

which can be used to compare the quality of multiple models fitted to the same dataset. Lower AIC values indicate models that are expected to predict new data more accurately. Similarly, the BIC is given by

$$\mathrm{BIC} = |\Psi| \ln(n) - 2\ln(\hat{L}), \tag{4}$$

where n is the number of observations. BIC imposes stricter penalty on model complexity compared to AIC, making it more conservative when selecting the number of components, specially for large datasets [28].

In practice, both AIC and BIC are calculated for models with different values of g, and the model with the lowest value of the selected criterion is considered optimal. While AIC favors more complex models, BIC generally performs better in identifying an adequate number of components in FMMs, especially in the context of large datasets [15,28]. In our implementation, we used GaussianMixture and ParameterGrid in the scikit-learn library for Python [18].

3 Related Work

Characterizing and comparing players in ice hockey has been done in different ways. The most common approach is to use performance metrics [8]. These range from the traditional metrics such as goals, assists, and points to Corsi and xG (expected goals) which are all well-known in the hockey discourse. To deal with some disadvantages of the traditional metrics, other advanced data-driven metrics have been proposed, such as extensions for the +/- metric using regularized logistic regression models [14,4]. There is also work on combining metrics, such as in [5] where principal component analysis is used on 18 basic stats. A major critique for traditional metrics has been that context is not taken into account. Therefore, some approaches for player performance metrics take game context into account such as event impacts, e.g., [24,19], and much of the work that models the dynamics of an ice hockey game using Markov games where two opposing sides (i.e., the home team and the away team) try to reach states in which they are rewarded (e.g., scoring a goal) [29,7,20,25,26,11,22,13,9]. We note that the introduction of new metrics may change the way the game is played. For instance, in [6] it was shown that team play transitioned first to taking more shots (high Corsi, shot-based), and then to taking high-quality shots (high expected goals). Player rankings are presented in [23,12,10]. In [23] a generalized additive model was used to predict player performance metrics from player demographics and player performance data, while in [12] a logistic regression model tree was used. In [10] predictive models were generated that can be used to identify and predict players' ranking tier (top 10%, 25% and 50%).

Player categorization is a relatively unexplored field in the context of ice hockey. In earlier work, a player could belong to only one role or category [30,3]. More recent work used soft clustering techniques to categorize players, allowing for a player to belong to different roles with some probability [21,17]. In the latter case, players can be compared based on their membership in different roles.

This work can be seen as a variant of the work of [17]. In that paper, we used player vectors to characterize a player's playing style. The player vectors contain representations of skills that are computed from game event data. Further, we applied fuzzy clustering on the vectors to generate five types of defender playing styles and five types of forward playing styles. For these types, we showed typical skill levels and players with similar styles. The data included complete seasons for the three leagues AHL, SHL, and HockeyAllsvenskan for 2021/22 and 2022/23, as well as data from the 2023/24 season up until Jan. 28th, 2024. In contrast, the present study focuses exclusively on the SHL and uses data from five full seasons (2019/20 to 2023/24). We use the same kind of player vectors, but apply GMMs to perform soft clustering.

4 Data

The dataset used in this research is a proprietary dataset developed by Sportlogiq[3], and consists of event data for all the SHL regular season games for 5 seasons (2019/20 to 2023/24). In total, the dataset consists of 1820 games, 1072 unique players, 16 unique teams, and 6,814,336 events. Among the 1072 unique players, there are 656 forwards, 377 defenders, and 94 goaltenders. We note that in the dataset, 55 players have been marked as playing in more than one position.

5 Method

5.1 Player vectors

We use the same kind of player vectors as introduced in [17]. In this section we recapitulate how these were developed. For defenders and forwards, different skills were identified (Tables 1 and 2). Each skill is represented by a set of features. For each skill, a feature vector was constructed which contains the frequencies of each feature that describes that skill, standardized using MinMaxScaler in the scikit-learn library for Python [18]. Further, non-negative matrix factorization (NMF) was applied to each feature vector using the NMF in the scikit-learn library. After this operation, every skill was represented by one feature and these were concatenated into player vectors. This resulted in player vectors of length 13 for defenders and of length 18 for forwards.

Figs. 1a and 1b use boxplots to show the distributions of the values for the skills for defenders and forwards, respectively, for the dataset containing all seasons. Here, the lower edge of the box represents the lower quartile value (25%) value, the (red) line in the box the median (50%) value, and the upper edge of the box the higher quartile (75%) value. The lower whisker shows the minimum value and the upper whisker the maximum value. Points below the lower whisker or above the upper whisker are outliers.

[3] https://www.sportlogiq.com/hockey/

Table 1: Skills and example actions for defenders [17].

Skills	Actions
Passing	e.g., different types of passes
Skating	e.g., exits, entries, dumps
Shooting	e.g., different types of shots
Defensive Stickwork	e.g., blocked passes, loose puck recoveries
Puck Moving	e.g., some types of passes, dump-in recoveries
Point Producing	e.g., different offensive zone events
Powerplay Playmaking	e.g., powerplay playmaking events
Powerplay Scoring	e.g., powerplay shots and goals
Physical Play	e.g., body checks and defensive plays
Slot Defense	e.g., blocked shots and dump outs
Stay at Home	e.g., different defensive zone events
Penalty Killing	e.g., different penalty killing events related to puck recovery
Penalty Killing Slot Defense	e.g., different penalty killing defensive plays

Table 2: Skills and example actions for forwards [17].

Skills	Actions
Passing	e.g., different types of passes
Skating	e.g., different types of controlled entries
Powerplay Playmaking	e.g., different types of controlled entries and passes in powerplay
Powerplay Slot Engagement	e.g., powerplay actions close to net
Powerplay Scoring	e.g., powerplay shots and goals
Defensive Puck Control	e.g., dump outs and loose puck recoveries
Defensive Zone Play	e.g., different defensive zone actions
Defensive Positioning	e.g., blocked shots and passes
Slot Defense	e.g., rebounds and dump outs
Penalty Killing	e.g., shorthanded defensive plays
Slot Engagement	e.g., offensive actions close to net
Heavy Game	e.g., body checks and defensive plays
Forechecking	e.g., offensive zone loose puck recoveries
Cycling the Puck	e.g., puck protections and receptions
Neutral Zone	e.g., different neutral zone actions
Puck Moving	e.g., some types of passes, entries
Offensive Zone Play	e.g., different offensive zone events
Shooting	e.g., different types of shots

5.2 Gaussian Mixture Model

To decide on the number of clusters for forwards and for defenders, we used the BIC approach. A full-factorial grid search was performed to identify different configurations of the model; resulting in a total of 15,360 different models being evaluated, each with different values of parameters such as number of components, covariance types, maximum number of iterations, and different initialization methods. Table 3 summarizes the example values used for each parameter.

Fig. 2 shows the average of AIC and BIC values based on the number of components applied to the skill vectors for forwards and defenders, respectively. Fig. 2a shows a significant decrease between three and five components for both AIC and BIC, suggesting that up to five components to the model improves

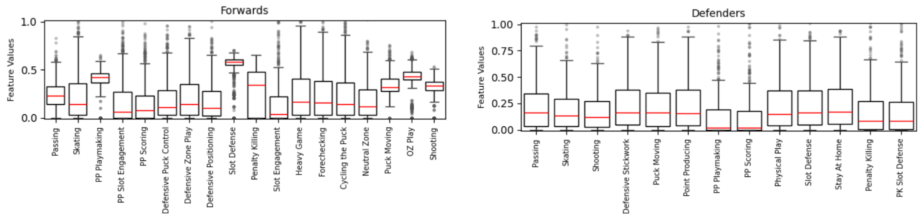

(a) Forwards. (b) Defenders.

Fig. 1: Boxplots of the skill value distributions, as calculated across all seasons.

Table 3: Grid Search Parameters for Gaussian Mixture Models.

Parameter	**Values**
Initialization Method	K-means++, Kmeans, Hybrid hierarchical, Custom Hierarchical
Number of components	3-15
Covariance type	Spherical, Tied, Diagonal, Full
Convergence Threshold	10^{-7}, 10^{-6}, 10^{-5}, 10^{-4}
Regularization covariance	10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}
Max iterations	100, 200, 300, 400, 500

the performance. The lowest average BIC value is observed at ten components. When analyzing each season individually, the average BIC reaches it minimum around five components for both forwards and defenders, as shown in Fig. 2b. Based on this observation, and to remain consistent with earlier work [17], we chose to use five components in this study.

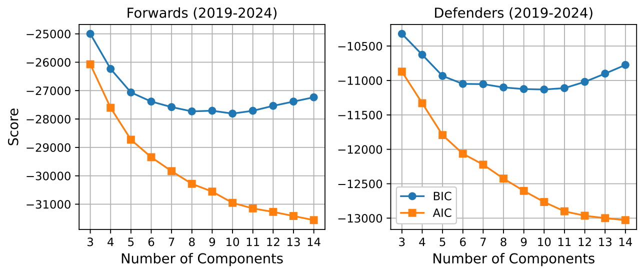

(a) Aggregated over all seasons.

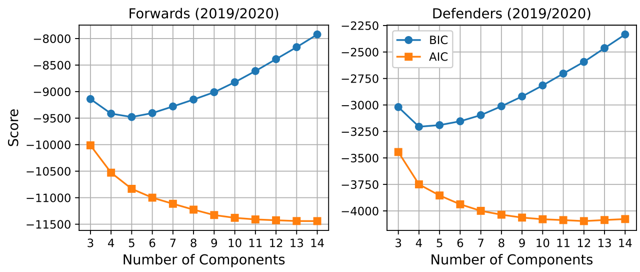

(b) Season 2019/20.

Fig. 2: Average AIC and BIC score by the number of components for forwards (left) and defenders (right).

6 Results

6.1 Forwards

Fig. 3a shows the average skill values of the ten forwards closest to the centroid of each of the five clusters, aggregated across all five seasons. When assigning each player to the cluster for which they have the highest membership, the players are relatively evenly distributed across the five clusters. Specifically, in clusters F0.19-24, F1.19-24, F2.19-24, F3.19-24, and F4.19-24, there are 142, 113, 121, 128, and 127 players, respectively.

The forwards in F0.19-24 have many skills, having strengths in slot defense, puck moving, OZ play, shooting, PP playmaking. Players in F1.19-24 show lower overall skill values but still have strengths in slot defense, PP playmaking, puck moving, and OZ play. The strengths of forwards in F2.19-24 are slot defense, penalty killing, and OZ play, but these players lack in slot engagement, PP slot engagement, and PP scoring. Forwards in F3.19-24 do not excel in any skills, but have strengths in slot defense, puck moving, OZ play, and shooting. Finally, players in F4.19-24 excel in most of the skills, having strengths in both defensive and offensive play.

We also investigated clustering results using data from a single season. As an example, Fig. 3b presents the average skill values for the ten forwards closest to the centroid of each of the five clusters from the 2019/20 season. The cluster skill bar plots for the other seasons show similar patterns. We observed that clusters F0.19-24, F2.19-24, F3.19-24, and F4.19-24 have clear counterpart in the 2019/20 season (F0.19-24 — F3.19/20, F2.19-24 — F4.19/20, F3.19-24 — F1.19/20, F4.19-24 — F2.19/20). Cluster F0.19/20 does not closely match any of the clusters aggregated over all five seasons. The forwards in F0.19/20 have many skills, with strength in defensive play, but lack in PP scoring, PP slot engagement and slot engagement. Example forwards closest to centroids for each cluster are presented in Table 4.

Players develop over time and our approach may give insights in the development of certain players. For instance, Richard Hugg belonged in the 2019/20 season to cluster F3.19/20. In 2020/21 he primarily belonged to a cluster similar to F0.19/20, and in 2021/22, he appeared in a cluster similar to F2.19/20. This progression suggests that over time, from being an offensively-skilled forward, he first improved his defensive play and then developed into a player that excels both defensively and offensively.

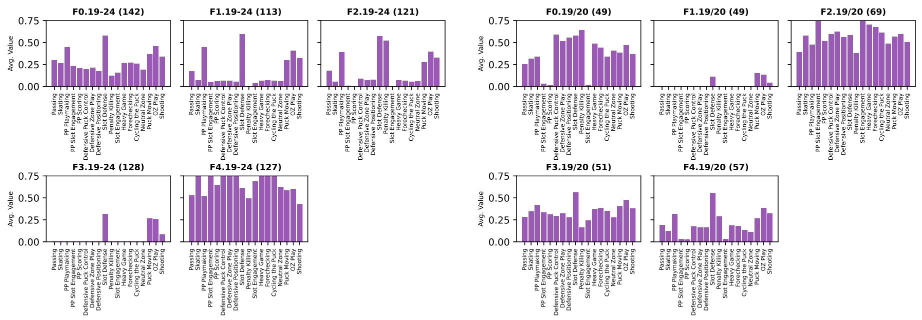

(a) Aggregated over all seasons. (b) Season 2019/20.

Fig. 3: Average skill values of the ten forwards closest to each cluster centroid.

6.2 Defenders

Fig. 4a shows the average skill values of the ten defenders closest to the centroid of each of the five clusters, aggregated over all seasons. Cluster D3.19-24 has the most players regarding highest membership. Defenders in cluster D0.19-24 have low skill levels. The players in this cluster all played few games in the SHL. D1.19-24 defenders are highly skilled and excel in defensive skills such as slot defense, stay at home, penalty killing, PK slot defense, physical play, and defensive stickwork. These players do not play powerplay. D3.19-24 defenders have average skill levels. Defenders in cluster D4.19-24 are highly skilled and in comparison to D1.19-24 defenders, they excel in offensive skills such as passing, shooting, puck moving, point producing, and they also play powerplay. The

Table 4: Forwards closest to the centroids of the clusters. F0.19-24 — F4.19-24 for the 5 seasons aggregated. F0.19/20 — F4.19/20 for season 2019/20.

F0.19-24 (142 players)	**F1.19-24** (113 players)	**F2.19-24** (121 players)	**F3.19-24** (128 players)	**F4.19-24** (127 players)
Marco Kasper	Jacob Micflikier	Mikael Frycklund	Mateusz Szurowski	Simon Ryfors
Dick Axelsson	Peter Holland	Mikkel Boedker	Melvin Fernström	Kalle Östman
Filip Cederqvist	Markus Nenonen	Juuso Ikonen	Johan Lundgren	Linus Fröberg
Markus Modigs	Marcus Paulsson	Joonas Nattinen	William Magnusson	Sebastian Strandberg
Tuomas Kiiskinen	Adam Johnson	Petrus Palmu	Linus Lööf	Andreas Wingerli
F0.19/20 (49 players)	**F1.19/20** (49 players)	**F2.19/20** (69 players)	**F3.19/20** (51 players)	**F4.19/20** (57 players)
Rok Ticar	Tuomas Kiiskinen	Emil Pettersson	Joakim Andersson	Jesper Kandergård
Gustav Possler	Olle Lycksell	Brendan Shinnimin	Johan Johnsson	Melker Eriksson
Marcus Paulsson	Dominik Bokk	Johan Sundström	Adam Pettersson	Alexander Ljungkrantz
Viktor Lodin	Juuso Ikonen	Greg Scott	Axel Wemmenborn	Linus Hedman
Linus Oberg	Johan Ryno	Ted Brithén	John Dahlstrom	Samuel Solem

players in cluster D2.19-24 played few games. In contrast to the players in D0.19-24, these players were junior players and some long-time injured players (e.g., Mattias Bäckman in 2019/20).

When investigating clusters across individual seasons, we observed that clusters similar to D1.19-24, D2.19-24, D3.19-24, and D4.19-24 consistently appeared. However, no clusters similar to D0.19-24 were found in any single season. Fig. 4b shows the average skill values for the ten defenders closest to the centroid of each of the five cluster for the 2019/20 season. In all other seasons, we found equivalent clusters for all except D4.19/20. For D4.19/20, an equivalent cluster was present in every season except 2023/24. In the 2023/24, cluster D2.19/20 appeared to represent an aggregation of two distinct clusters. Example defenders closest to each cluster are presented in Table 5.

Similarly as for forwards, defenders can change cluster during their career. For instance, Rasmus Rissanen belonged mainly to cluster D2.19/20 while in 2023/24 he belonged mainly to the cluster that matches D2.19/20 with high values for all the skills. In this case, his best skills are still in the defensive work, but he has raised the skill level of most of his skills.

6.3 Practical Applications

Beyond descriptive clustering, these playing style profiles have practical applications in player management and roster decisions. For example, clubs can use a player's cluster profile to find comparable players when a replacement is needed. Further, insights into player development and adaptability may be obtained by observing how players' cluster memberships vary across coaching system or seasons, helping to distinguish which skill-based features are intrinsic to the player and which are influenced by team context. Furthermore, by connecting emerging players with established professional archetypes, these clusters can help with scouting in lower-tier leagues or youth programs, i.e., in situations where conventional metrics are typically used for evaluations.

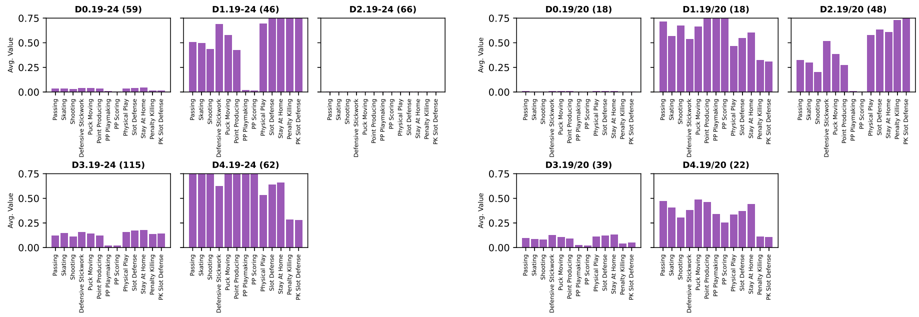

(a) All seasons. (b) Season 2019/20.

Fig. 4: Average skill values of the ten defenders closest to each cluster centroid.

Table 5: Defenders closest to the centroids for the clusters. D0.19-24 - D4.19-24 for the 5 seasons aggregated. D0.19/20 - D4.19/20 for season 2019/20.

D0.19-24 (60 players)	**D1.19-24** (42 players)	**D2.19-24** (67 players)	**D3.19-24** (127 players)	**D4.19-24** (57 players)
Lukas Klok	Anton Mylläri	Albin Thyni Johansson	Daniel Brickley	Matt Caito
Jordan Murray	Oscar Engsund	Nils Strandberg Sarén	Ville Pokka	Oskar Nilsson
Axel Landén	Jonathan Sigalet	Oskar Hassel	Julius Bergman	Lucas Ekeståhl-Jonsson
Theodor Johnsson	Daniel Glad	Jakob Bondesson	Joonas Lyytinen	Kristian Näkyvä
Elias Rosen	Arvid Lundberg	Gustav Berglund	Anton Strålman	Joel Nyström
D0.19/20 (18 players)	**D1.19/20** (41 players)	**D2.19/20** (40 players)	**D3.19/20** (29 players)	**D4.19/20** (19 players)
Gustav Berglund	Jonathon Blum	Oscar Engsund	Filip Johansson	Niklas Hansson
Jakob Bondesson	Nils Lundkvist	Emil Wahlberg	Jonas Junland	Miika Koivisto
Emil Andrae	Jonathan Pudas	Jonathan Sigalet	Lucas Nordsäter	Simon Despres
Albin Thyni Johansson	Ilari Melart	Arvid Lundberg	Patrik Norén	Jesper Sellgren
Christian Lindberg	Erik Gustafsson	Niklas Arell	Julius Bergman	Eric Martinsson

7 Conclusion

In this paper, we presented a Gaussian Mixture Model (GMM) approach for characterizing playing styles among ice hockey defenders and forwards. Our method provides a data-driven framework for identifying distinct player types based on skill profiles, offering new insights into player evaluation and team composition.

In future work, we plan to use data from AHL and HockeyAllsvenskan as well (as in [17]) and investigate whether the playing styles are the same or different in the different leagues. Further, we will use the algorithms from [17] on the data used in this paper and compare the different techniques. Such comparisons are expected to provide insights into the relative strengths and weaknesses of different unsupervised learning techniques for player style characterization.

Overall, our findings contribute to the growing research area on quantitative analysis of player behavior, and we hope they will provide tools and foundation for further research into improved player development, scouting, and strategic decision-making in professional ice hockey.

References

1. Akaike, H.: A new look at the statistical model identification. IEEE Transactions on Automatic Control **19**(6), 716–723 (1974). https://doi.org/10.1109/TAC.1974.1100705
2. Bishop, C.M.: Pattern Recognition and Machine Learning. Springer (2006)
3. Chan, T.C., Cho, J.A., Novati, D.C.: Quantifying the contribution of NHL player types to team performance. Interfaces **42**(2), 131–145 (2012). https://doi.org/10.1287/inte.1110.0612
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**(1), 97–111 (2013). https://doi.org/10.1515/jqas-2012-0001
5. Gu, W., Foster, K., Shang, J., Wei, L.: A game-predicting expert system using big data and machine learning. Expert Systems with Applications **130**, 293–305 (2019). https://doi.org/10.1016/j.eswa.2019.04.025
6. Johansson, U., Wilderoth, E., Sattari, A.: How analytics is changing ice hockey. In: Lambrix, P., Carlsson, N., Vernblom, M. (eds.) Proceedings of the Linköping Hockey Analytics Conference LINHAC 2022 Research Track. Linköping Electronic Conference Proceedings, vol. 191, pp. 49–59 (2022). https://doi.org/10.3384/ecp191006
7. 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). https://doi.org/10.3138/infor.52.2.39
8. Lambrix, P., Carlsson, N.: Data-driven player performance metrics in ice hockey (2025), forthcoming book chapter.
9. Lambrix, P., Carlsson, N., Säfvenberg, R.: Goal-based performance metrics for ice hockey accounting for goal importance. Journal of Sports Analytics **11** (2025). https://doi.org/10.1177/22150218241290987
10. Lehmus Persson, T., Kozlica, H., Carlsson, N., Lambrix, P.: Prediction of tiers in the ranking of ice hockey players. In: Brefeld, U., Davis, J., Van Haaren, J., Zimmermann, A. (eds.) Machine Learning and Data Mining for Sports Analytics. MLSA 2020. Communications in Computer and Information Science, vol. 1324, pp. 89–100 (2020). https://doi.org/10.1007/978-3-030-64912-8_8
11. Liu, G., Schulte, O.: Deep reinforcement learning in ice hockey for context-aware player evaluation. In: Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence. pp. 3442–3448 (2018). https://doi.org/10.24963/ijcai.2018/478
12. Liu, Y., Schulte, O., Li, C.: Model Trees for Identifying Exceptional Players in the NHL and NBA Drafts. In: Brefeld, U., Davis, J., Van Haaren, J., Zimmermann, A. (eds.) Machine Learning and Data Mining for Sports Analytics. MLSA 2018. Lecture Notes in Computer Science, vol. 11330, pp. 93–105 (2019). https://doi.org/10.1007/978-3-030-17274-9_8
13. Ljung, D., Carlsson, N., Lambrix, P.: Player pairs valuation in ice hockey. In: Brefeld, U., Davis, J., Van Haaren, J., Zimmermann, A. (eds.) Machine Learning and Data Mining for Sports Analytics. MLSA 2018. Lecture Notes in Computer Science, vol. 11330, pp. 82–92 (2019). https://doi.org/10.1007/978-3-030-17274-9_7
14. Macdonald, B.: A Regression-Based Adjusted Plus-Minus Statistic for NHL Players. Journal of Quantitative Analysis in Sports **7**(3) (2011). https://doi.org/10.2202/1559-0410.1284

15. McLachlan, G., Peel, D.: Finite Mixture Models. Wiley Series in Probability and Statistics, Wiley (2000)
16. McLachlan, G.J., Lee, S.X., Rathnayake, S.I.: Finite mixture models. Annual Review of Statistics and Its Application **6**, 355–378 (2019). https://doi.org/10.1146/annurev-statistics-030718-104035
17. Olivestam, A., Rosendahl, A., Wilderoth, E., Carlsson, N., Lambrix, P.: Characterizing playing styles for ice hockey players. In: Brecht, T., Carlsson, N., Vernblom, M., Lambrix, P. (eds.) Proceedings of the Linköping Hockey Analytics Conference LINHAC 2023 Research Track. Linköping Electronic Conference Proceedings, vol. 209, p. 39–50 (2024). https://doi.org/10.3384/ecp209004
18. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. Journal of Machine Learning Research **12**, 2825–2830 (2011)
19. Pettigrew, S.: Assessing the offensive productivity of NHL players using in-game win probabilities. In: MIT Sloan Sports Analytics Conference (2015)
20. 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)
21. Säfvenberg, R., Carlsson, N., Lambrix, P.: Identifying player roles in ice hockey. In: Brefeld, U., Davis, J., Van Haaren, J., Zimmermann, A. (eds.) Machine Learning and Data Mining for Sports Analytics. MLSA 2023. Communications in Computer and Information Science, vol. 2035, pp. 131–143 (2024). https://doi.org/10.1007/978-3-031-53833-9_11
22. Sans Fuentes, C., Carlsson, N., Lambrix, P.: Player impact measures for scoring in ice hockey. In: MathSport International 2019 Conference. pp. 307–317 (2019)
23. Schuckers, M.: Draft by Numbers: Using Data and Analytics to Improve National Hockey League (NHL) Player Selection. In: MIT Sloan Sports Analytics Conference (2016)
24. 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)
25. 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). https://doi.org/10.1007/s10618-017-0496-z
26. 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)
27. Schwarz, G.: Estimating the Dimension of a Model. The Annals of Statistics **6**(2), 461 – 464 (1978). https://doi.org/10.1214/aos/1176344136
28. Steele, R.J., Raftery, A.E.: Performance of bayesian model selection criteria for gaussian mixture models. University of Washington, Department of Statistics, Technical Report (570) (2009), https://stat.uw.edu/research/tech-reports/performance-bayesian-model-selection-criteria-gaussian-mixture-models
29. 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). https://doi.org/10.1214/13-AOAS646
30. Vincent, C.B., Eastman, B.: Defining the style of play in the nhl: An application of cluster analysis. Journal of Quantitative Analysis in Sports **5**(1) (2009). https://doi.org/10.2202/1559-0410.1133

Individual Puck Possessions Part I: Frequency, Duration, and Distance Travelled

Evan Iaboni, Sebastian Negulescu, Miles Pitassi, Fauzan Lodhi, and Tim Brecht

Cheriton School of Computer Science, University of Waterloo

Abstract. In this paper we use puck and player tracking data from the 2023-24 NHL season to study individual player possessions (focusing on 5v5 situations). We study metrics such as possession count, average and total possession duration, average and total distance travelled with the puck, and examine relationships between these metrics and traditional measures of success (i.e., goals, assists and points). A key finding is that individual offensive zone possession is strongly correlated with points ($r = 0.70$) and is moderately correlated with goals ($r = 0.64$), assists ($r = 0.54$), and shots on goal ($r = 0.69$). We also observe differences in individual possessions based on position (forwards versus defence), zone of play, and strength and large and statistically significant differences between top ranked players and league averages (across most possession metrics). Finally, we examine the benefits of our individual possession metrics and find that they are highly stable (so they are useful for predictions), able to differentiate players, and provide information not captured by existing metrics.

1 Introduction

In the 2021-22 season, the National Hockey League (NHL) added puck and player tracking (PPT) technologies to all arenas, creating significant opportunities for advanced hockey analytics. The introduction of an individual player possession model into the "DISH" data stream in March 2023 added the ability to measure and understand teams' and players' abilities to control the puck.

In a previous study, we investigated the relationship between puck possession and team success in the NHL. We found that overall team possession had a modest correlation with team success; however, we introduced a new metric, average offensive zone possession time differential (Avg. OZPTD), which showed a strong correlation ($r = 0.77$) with team success, measured by average goal differential. This work advanced our understanding of team-level possession dynamics and highlighted the importance of controlling the puck in high-value areas of the ice. However, it also raised a series of questions about the role of individual possession in driving these team-level outcomes. Is team success simply a function of aggregate possession time, or do individual contributions matter? Which players dominate possession and carry the puck over the greatest distances? How do these metrics vary across positions and game situations, and how do they compare to possession metrics in other sports? These questions drive our present study, marking a natural progression from team-focused analyses to an investigation of individual player performance.

To address these questions, we introduce and compute novel metrics, including average possession duration and distance travelled per possession. Where appropriate, all metrics are normalized to eliminate the impact of differences in ice time. These metrics are analyzed across 5v5 situations unless otherwise noted as power plays, penalty kills, and other non-5v5 situations skew individual possession.

We report on the top players (we exclude goaltenders from our analysis) in several possession metrics categories and discover statistically significant differences between top players and the league average across several metrics. By identifying players who excel at possession related metrics this work offers valuable insights that could be used for talent evaluation, as well as line and roster construction. Most importantly, these metrics can identify players with skills beyond traditional metrics like goals and assists. In this work we make the following contributions:

- We define, compute and study several metrics related to individual possession including: total possession duration, average possession duration, average distance travelled per possession, and possession count. We find that, for most of these metrics, there are statistically significant differences among some of the top 10 players and that the top 10 players differ substantially from the league average.
- We examine correlations between individual possession metrics and other offensive metrics and find there is a strong correlation ($r = 0.70$) between offensive zone possession time and points. When examining the same relationship during power plays we find a much lower correlation ($r = 0.42$).
- We find that individual possession metrics are able to distinguish players from one another, stable throughout the season, and introduce new information not captured by existing metrics.

2 Related Work

Most research in soccer, basketball, and hockey has focused on team possession [1][2][8][15][10]. Relatively little research has examined individual possessions in those sports. We now describe studies that examine individual possession in soccer, basketball, and hockey. Below, we use μ and σ^2 when referencing a statistic to mean that the average is μ and the variance is σ^2. Additionally, we use the notation $m:ss$ to mean m minutes and ss seconds.

Using data from 60 soccer matches from the 2012-13 season of the German Soccer League (Bundesliga), and excluding player's who played less than 45 minutes in a match, Link *et al.* [7] found that the league average was 0:14 $\pm$ 0:13 of possession per 10 minutes of playing time. Note that they use a more restrictive definition than what we introduce for possession in hockey. They found that central midfielders and central defenders had the most instances of ball control with an average of 12.8 possessions per 10 minutes of playing time. They found that, on average, wingers and central forwards had fewer instances of controlled possessions but the differences were not statistically significant. The authors found a strong correlation ($r = 0.85$) between the time a player spent in control of the ball and a player's ball control count.

Next, we examine studies of basketball. To better understand variability, where possible (e.g., if the paper provides sample sizes), we convert standard deviations to 95%

mean confidence intervals (assuming a normal sampling distribution). We use the notation [L, H] to indicate that L is the low and H is the high of a 95% mean confidence interval. Sampaio *et al.* [12] used data from 1,230 games played in the 2013-14 NBA season, and found that players who were on the first, second, or third All-NBA team ($n = 15$) possessed the ball for [5.4 sec, 9.0 sec] per minute played on average. In contrast, non all-star players ($n = 324$), possessed the ball for an average of [4.4 sec, 5.2 sec] per minute played. The authors also found that all-stars averaged [1.94, 2.26] touches of the ball per minute as opposed to non all-stars, who averaged [1.65, 1.75].

A study conducted between 2015 and 2017 using 70 games with point differences of 20 points or less, involving 44 male players from Italy's Lega Basket Serie A and Serie A2 who had played at least 10 minutes in one game prior to the study period, found that, for guards, the mean percentage of playing time in possession of the ball was [9.4%, 14.4%] [3]. This is significantly higher than forwards and centres, who possessed the ball for [2.8%, 4.2%] and [2.1%, 3.7%] of their playing time respectively. Zhang *et al.* [14] used a sample of 699 games with a final score difference of less than or equal to 10 points from the 2015-16 NBA season found no significant difference between the amount of ball touches between players on teams that made the playoffs and players on teams that did not make the playoffs. Although not explained in these papers, we hypothesize that they exclude games with larger score differentials because they do not reflect typical game play.

In a study of the 2002 Olympic Winter men's ice hockey games, the author manually tracked individual possession for the top players, including Mike Modano and Joe Sakic [13]. The study found that these players averaged 1:07 with the puck per game. The study also looked at the top youth players in the USA hockey Tier 1 Youth National Championships and found that the top players averaged 01:06 of possession. The study highlights the extremely small amount of time players handle the puck during a game and emphasizes the benefits of team practices with more puck handling opportunities.

To compare previous findings with our own we normalize soccer results to 20 minutes of game time and find an average of 0:28 $\pm$ 0:26 in control of the ball per 20 minutes. For basketball, using the non-all star group, the same normalization yields [1:31, 1:49] of ball possession time and [33.2, 34.8] touches of the ball (both per 20 minutes of playing time). Our findings show slightly higher amounts of puck possession times for top players than the previous hockey study and that hockey falls somewhere between soccer and basketball, with an average possession duration of [0:39, 0:40] per 20 minutes and an average possession count of [33, 34]. In this paper, we examine the top 10 players across different metrics that capture data regarding individual possession. In addition, we compare forwards and defencemen, and examine the correlations between individual possession metrics and traditional metrics that are based on offensive production.

In Part II of this research [9] on individual puck possessions, we devise a metric to measure players that consistently carry the puck at high speeds (called Bursts20). This is a measure of the average number of times, per game, a player carries the puck for one second or longer where their speed reaches 20 mph or more during the possession. To avoid biases towards players who are given more ice time, we normalize the average to 20 minutes of ice time per game. Using data from the 2023-24 NHL regular season, that work shows significant differences between forwards and defencemen. Forwards average [0.39, 0.48] Bursts20 while defencemen average [0.12, 0.15]. Nathan MacKinnon, the top ranked player, is significantly above average with [2.87, 3.83] bursts per 20 minutes of ice time. Cale Makar, the top ranked defencemen averages [0.50, 0.90] such bursts per 20 minutes. That paper also explores the distribution of individual player possessions during a game within a team. Jain's Fairness Index [6] is used to compute an equity score, which measures the degree to which possessions within a team are equal. The metric produces values between 0 and 1, with 1 being completely equal. The results show that despite there being a significant difference between teams with high equity scores (e.g., FLA with [0.84, 0.85]), medium equity scores (e.g., BOS with [0.80, 0.82]) and low equity scores (e.g., VAN with [0.68, 0.72]). All three of those teams were ranked in the top 4 teams in the league in terms of average 5v5 goal differential, indicating that different styles of individual puck possessions can be successful.

3 Background

3.1 Definition of Individual Puck Possession

We use the definition of individual puck possession deployed by the NHL in the model used to obtain the individual possession data we employ. There are two types of possessions. First, a player is considered to have possession and control of the puck when they have two or more consecutive touches of the puck. The possession starts with the first touch and is confirmed by a second touch. Individual possession ends when the player loses the puck or another player gains possession. The second type of possession includes specific one-touch actions, such as shot attempts, passes, or area plays (e.g., dump-ins), which are considered brief moments of possession [11].

3.2 Dataset Overview

Our study uses proprietary PPT data from the NHL which records x, y, and z coordinates at high frequencies (60 times per second for the puck and 12 times per second for each player on the ice). Players on the bench are updated once per second. In total, approximately 734,400 locations are recorded in a typical 60-minute game. Due to different data collection frequencies and lack of synchronization, all locations are interpolated to uniform one-hundredth of a second timestamps.

In March 2023, an individual player possession model was added to the "DISH" data stream, which features Delayed, Interpolated, Smoothed, and Hundred-Hertz enhancements. That possession model uses the (x,y) coordinates of the puck and players (i.e., their location on the ice). The changes between consecutive readings for all on-ice entities are combined with physics-based models to attempt to determine when a player is and is not in possession of the puck. The dataset we are provided contains information about when an individual possession starts and ends, including the unique identification number for the player. We applying cleaning and filtering techniques to the original data, described in Section 4, to obtain the data used in our analysis. This dataset is considered unofficial by the NHL and may differ from other datasets that track possession information (e.g., a hand-labelled dataset).

There were 1,312 games played in the 2023-24 NHL season. After finding 118 games with significant data issues, as detailed in Section 4.1, 1,194 games remained for our analysis.

4 Dataset Cleaning and Filtering

4.1 Preparing, Cleaning, and Filtering Games

Our dataset preparation, cleaning, and filtering methods are adapted from our previous work [10]. First, we merged the player possession data with a detailed game information file to provide broader game context, such as power play situations, goal differentials, and player locations. This enriched the player possession data with relevant context. Next, we cleaned the dataset by removing possessions that occurred during stoppages or clock resets, and those that were out of sequence, overlapping, or duplicated. Despite these efforts, some errors persisted, leading us to exclude games based on data corruption severity: if the data is compromised for either more than 4% of a games duration or more than 4% of a teams possession time, we exclude the game from our analysis. This resulted in the exclusion of 118 games, leaving 91% of the games for analysis. For more detailed information on the filter criteria, refer to our previous work [10].

4.2 Filtering Individual Players

After filtering down to 1,194 games, we also filtered out players for which there was insufficient data. We chose our filters in a way that ensured players had enough opportunity for our possession and production metrics to be reliably captured. We exclude players with less than 10 games played or less than 10 minutes of 5v5 ice time per game. In the 1,194 game sample, 921 players participated and 250 were excluded, leaving 671 players for our analysis (72.9% of the players who dressed for the included games). With 32 teams in the league and 18 skaters (forwards and defencemen) dressing per team, that equals 576 skaters (32×18). Since our dataset includes 671 players, we believe our analysis captures a representative sample of regularly participating players.

We applied additional filters when conducting power play and short-handed analysis. Players with less than 25 total minutes of ice time in these situations for the season were removed. This left 45.9% of players for power play analysis and 45.4% of players for short-handed analysis.

5 Individual Possession Analysis

We now analyze several metrics related to individual possession and their relationship with traditional player performance metrics such as points, goals, assists, and shots on goal. We found that possession duration is strongly correlated with time on ice ($r = 0.89$), so where appropriate, metrics are normalized to 20 minutes of ice time, providing for a fair comparison across all players (regardless of their ice time).

5.1 Possession Duration and Count

We use total possession duration (or possession time) to denote the average number of minutes that a player possesses the puck per game, normalized to 20 minutes of ice time. Similarly, possession count is the average number of times a player has possession of the puck per game, normalized to 20 minutes of ice time. The average possession duration is a player's total possession duration divided by their possession count (shown in seconds).

Total Possession Duration We start this section by examining the total possession duration. Table 1 shows the top 10 forwards (left) and defencemen (right), ranked by average total possession duration per 20 minutes (Possession Duration in the table). Note that times in the table are denoted as m:ss, to indicate minutes and seconds.

As described in the caption for Table 1, †† indicates that the player is ranked among the top 10 players for average possession duration, total possession duration, and possession count. We note that it should not be surprising that some players appear in all three tables since there is a strong correlation between total possession duration and both average possession duration ($r = 0.87$) and possession count ($r = 0.86$) and there is a moderate correlation between average possession duration and possession count ($r = 0.51$).

Rank	Name	Team	Possession Duration (m:ss)	Name	Team	Possession Duration (m:ss)
1	Mathew Barzal†	NYI	1:21 [1:16, 1:26]	Quinn Hughes††	VAN	1:37 [1:31, 1:43]
2	Jack Hughes††	NJD	1:15 [1:08, 1:19]	Cale Makar††	COL	1:25 [1:19, 1:31]
3	Jack Eichel††	VGK	1:14 [1:09, 1:20]	Jake Sanderson†	OTT	1:18 [1:13, 1:22]
4	Leon Draisaitl††	EDM	1:13 [1:07, 1:18]	Owen Power†	BUF	1:16 [1:11, 1:20]
5	Artemi Panarin†	NYR	1:06 [1:03, 1:09]	Erik Karlsson††	PIT	1:13 [1:08, 1:18]
6	Nikolaj Ehlers	WPG	1:05 [1:02, 1:08]	Scott Perunovich†	STL	1:13 [1:06, 1:19]
7	William Nylander†	TOR	1:05 [1:01, 1:08]	Mike Matheson†	MTL	1:13 [1:08, 1:18]
8	Clayton Keller†	ARI	1:04 [1:01, 1:08]	Cam Fowler†	ANA	1:11 [1:07, 1:18]
9	Connor McDavid†	EDM	1:03 [0:58, 1:08]	Drew Doughty†	LAK	1:11 [1:07, 1:16]
10	Troy Terry†	ANA	1:03 [0:58, 1:07]	Evan Bouchard	EDM	1:10 [1:04, 1:14]
+	Forwards Avg.		0:40 [0:39, 0:40]	Defencemen Avg.		0:49 [0:48, 0:50]

Table 1. Top 10 forwards (left) and defencemen (right) ranked by total possession duration in 5v5 situations. Numbers in square brackets are the low and high of the 95% confidence interval for the by-game mean. † indicates this player appears in another top 10 table in this section, and †† indicates they appear in two other tables (i.e., all three tables).

The player with the highest total possession duration per 20 minutes is Quinn Hughes (VAN) with 1:37 followed by Cale Makar (COL) with 1:25 and Mathew Barzal (NYI) with 1:21. The mean 95% confidence interval for total possession duration is [0:48, 0:50] among defencemen and [0:39, 0:41] among forwards. This suggests that on average, players spend about 3.3% to 4.2% of their ice time with the puck. Moreover, since the confidence intervals do not overlap, there is a statistically significant difference [5] for mean total possession duration between defencemen and forwards.

In Table 2 we show correlations between total possession duration and individual offensive production metrics. Notice that there is a weak correlation with all offensive production metrics when separating forwards and defencemen, however the correlations vanish when combining all positions. It is important to consider forwards and defencemen separately since the nature of their possessions are different. Notably, only 19% of possession time by defencemen is in the offensive zone in contrast to 46% for forwards.

Offensive Production	Forwards r-value	Defencemen r-value	Combined r-value
Points per 20	0.39	0.35	-0.02
Goals per 20	0.24	0.20	-0.14
Assists per 20	0.36	0.31	0.12
SOG per 20	0.38	0.34	-0.02

Table 2. Correlations between 5v5 possession per 20 minutes and player offensive production.

Average Possession Duration Shifting our focus to average possession duration, Table 3 shows the top 10 forwards (left) and defencemen (right), ranked by average possession duration (Avg. Time (sec.) in the table) with 95% confidence intervals (shown in green).

Rank	Name	Team	Avg. Time (sec.)	Name	Team	Avg. Time (sec.)
1	Mathew Barzal†	NYI	2.03 [1.94, 2.11]	Quinn Hughes††	VAN	2.00 [1.93, 2.07]
2	William Nylander†	TOR	1.85 [1.77, 1.93]	Cam Fowler†	ANA	1.89 [1.80, 1.98]
3	Jack Eichel††	VGK	1.82 [1.74, 1.91]	Cale Makar††	COL	1.87 [1.79, 1.95]
4	Troy Terry†	ANA	1.78 [1.70, 1.86]	Scott Perunovich†	STL	1.80 [1.68, 1.93]
5	Evgeny Kuznetsov	CAR	1.78 [1.67, 1.88]	Mike Matheson†	MTL	1.78 [1.70, 1.86]
6	Jack Hughes††	NJD	1.76 [1.69, 1.84]	Owen Power†	BUF	1.78 [1.71, 1.84]
7	Connor Bedard	CHI	1.76 [1.68, 1.84]	Drew Doughty†	LAK	1.74 [1.67, 1.81]
8	Leon Draisaitl††	EDM	1.75 [1.68, 1.82]	John Klingberg	TOR	1.73 [1.55, 1.90]
9	Kent Johnson	CBJ	1.72 [1.59, 1.85]	Bowen Byram	COL	1.72 [1.64, 1.79]
10	Isac Lundestrom	ANA	1.67 [1.54, 1.81]	Erik Karlsson††	PIT	1.71 [1.64, 1.78]
+	Forwards Avg.		1.27 [1.25, 1.28]	Defencemen Avg.		1.32 [1.30, 1.34]

Table 3. Top 10 forwards (left) and defencemen (right) ranked by average possession duration in 5v5 situations. Numbers in square brackets are the low and high of the 95% confidence interval for the by-game mean. † indicates this player appears in another top 10 table in this section, and †† indicates they appear in two other tables (i.e., all three tables).

For the forwards table, with the exception of two players, William Nylander (TOR) and Troy Terry (ANA), all of the top 10 players play centre (according to their listed position on the NHL website). Isolating the centres into their own group, their mean average possession duration is 1.28 seconds (with a 95% confidence interval of [1.25 seconds, 1.31 seconds]). This means that there is no significant difference between the averages of all centres and all forwards (who, from Table 3, have a 95% mean confidence interval of [1.25 seconds, 1.28 seconds). With that being said, there is a small but statistically significant difference between forwards and defencemen (who have a 95% mean confidence interval of [1.30 seconds, 1.34 seconds]).

Table 4 shows the average possession duration by zone and position (in seconds with 95% confidence intervals). For all players (All Positions), the average possession duration is similar in the defensive and neutral zones, but lower in the offensive zone (the differences are statistically significant).

Position	**All Zones**		**Defensive Zone**		**Neutral Zone**		**Offensive Zone**	
All Skaters	1.29	[1.27, 1.30]	1.43	[1.41, 1.45]	1.47	[1.45, 1.49]	1.05	[1.03, 1.06]
Forwards	1.27	[1.25, 1.28]	1.39	[1.36, 1.41]	1.56	[1.53, 1.59]	1.11	[1.09, 1.13]
Defence	1.32	[1.30, 1.34]	1.50	[1.47, 1.53]	1.33	[1.31, 1.35]	0.95	[0.94, 0.97]

Table 4. Avg. Possession Duration (seconds) by Zone for Different Positions in 5v5 Situations.

On average, players hold the puck for 1.29 seconds. In the defensive zone, this average increases to 1.43 seconds and again to 1.47 seconds in the neutral zone. However in the offensive zone, the mean average possession duration decreases to 1.05 seconds. On average, players hold the puck for significantly less time when in the offensive zone. Table 4, also shows that for all zones, defencemen have a slightly higher average possession duration. However, on average per possession, forwards hold the puck longer in the neutral and offensive zones than defencemen.

Possession Count Table 5 shows the top 10 forwards (left) and defencemen (right), ranked by possession count per 20 minutes (Possession Count in the table). Clayton Keller and Jack Hughes top the list for forwards, while Quinn Hughes and Jake Sanderson have the highest possession counts among defencemen. Comparing the two groups, the average possession count for defencemen is higher than the average possession count for forwards and since the 95% confidence intervals do not overlap ([36.6, 37.5] versus [30.9, 31.7]), this difference is statistically significant [5].

5.2 Distance Travelled Per Possession

We now introduce a new metric, average distance travelled per possession. The intent of this metric is to identify puck carriers, which in addition to indicating puck carrying skill, could be useful in constructing lines and defensive pairings. For this metric we only consider possessions longer than one second to capture “puck-carrying” possessions. Table 6 shows the top 10 forwards (left) and defencemen (right), ranked by average distance travelled per possession.

The 95% mean confidence interval for average distance travelled per possession is [37.9 feet, 38.9 feet] for forwards, and [32.5 feet, 33.5 feet] for defencemen. This

Rank	Name	Team	Possession Count	Name	Team	Possession Count
1	Clayton Keller†	ARI	43.4 [41.3, 45.2]	Quinn Hughes††	VAN	48.9 [47.0, 50.6]
2	Jack Hughes††	NJD	43.0 [40.1, 44.4]	Jake Sanderson†	OTT	46.9 [45.1, 48.5]
3	Artemi Panarin†	NYR	42.3 [40.6, 43.6]	Adam Fox	NYR	46.4 [44.8, 48.1]
4	Kevin Fiala	LAK	42.2 [40.3, 44.3]	Cale Makar††	COL	45.6 [44.0, 47.4]
5	Leon Draisaitl††	EDM	41.8 [39.8, 43.6]	Erik Gustafsson	NYR	45.5 [43.3, 47.1]
6	Jordan Kyrou	STL	41.6 [39.5, 43.4]	Zach Werenski	CBJ	44.8 [42.7, 46.7]
7	Nikita Kucherov	TBL	41.5 [39.4, 43.0]	Ryan Johnson	BUF	44.5 [42.0, 47.1]
8	Connor McDavid†	EDM	41.5 [39.4, 43.2]	Rasmus Dahlin	BUF	43.7 [41.9, 45.2]
9	Jack Eichel††	VGK	41.0 [39.3, 42.7]	MacKenzie Weegar	CGY	43.2 [41.2, 44.9]
10	Matt Boldy	MIN	40.9 [39.1, 42.6]	Erik Karlsson††	PIT	43.2 [41.4, 44.7]
+	Forwards Avg.		31.3 [30.9, 31.7]	Defencemen Avg.		37.1 [36.6, 37.5]

Table 5. Top 10 forwards (left) and defencemen (right) ranked by possession count in 5v5 situations. Numbers in square brackets are the low and high of the 95% confidence interval for the by-game mean. † indicates this player appears in another top 10 table in this section, and †† indicates they appear in two other tables (i.e., all three tables).

Rank	Name	Team	Distance Travelled per Possession (ft.)	Name	Team	Distance Travelled per Possession (ft.)
1	Mathew Barzal	NYI	53.8 [51.5, 56.0]	Samuel Girard	COL	43.2 [41.0, 45.4]
2	Denis Gurianov	NSH	53.2 [45.3, 61.2]	Quinn Hughes	VAN	42.9 [41.4, 44.3]
3	William Nylander	TOR	51.4 [49.1, 53.7]	Cam Fowler	ANA	42.0 [40.2, 43.8]
4	Jack Hughes	NJD	51.4 [49.1, 53.6]	John Klingberg	TOR	41.6 [37.7, 45.6]
5	Paul Cotter	VGK	51.3 [48.2, 54.5]	Nikita Zadorov	CGY	41.6 [39.1, 44.1]
6	Jack Eichel	VGK	50.9 [48.5, 53.4]	Cale Makar	COL	41.5 [39.9, 43.1]
7	Noah Gregor	TOR	50.3 [46.7, 53.8]	Wyatt Kaiser	CHI	41.3 [37.9, 44.8]
8	Julien Gauthier	NYI	50.0 [42.8, 57.2]	Mike Matheson	MTL	41.3 [39.6, 43.0]
9	Connor Bedard	CHI	49.8 [47.4, 52.1]	Michael Kesselring	ARI	40.8 [38.5, 43.1]
10	Adam Fantilli	CBJ	49.8 [46.3, 53.2]	Charlie McAvoy	BOS	40.6 [38.8, 42.3]
+	Forwards Avg.		38.4 [37.9, 38.9]	Defencemen Avg.		33.0 [32.5, 33.5]

Table 6. Top 10 players ranked by Distance Travelled per Possession. The numbers in square brackets are the low and high of the mean 95% confidence interval.

is a statistically significant difference since the confidence intervals do not overlap. It indicates that, on average, forwards tend to travel further with the puck per possession.

All the top 10 players in both tables have overlapping confidence intervals, indicating little difference in average distance travelled per possession among them. However, it is noteworthy that the top 10 defencemen have significantly higher average distance travelled than the league mean among defencemen of 33.0 feet. Likewise, the top 10 forwards greatly exceed the league mean among forwards of 38.4 feet. The very top player, Mathew Barzal, averages 15.4 feet more than the league average among forwards.

5.3 Offensive Zone Possession

In Table 7, we list the top 10 players (all forwards) ranked by offensive zone (oZone) possession time per 20 minutes (oZone Pos20 in the table). We are particularly interested in forwards here so we choose to not include separate table for defencemen. We also include columns for offensive production metrics, (normalized per 20 minutes of 5v5 ice time) including points (P20), goals (G20), assists (A20) and shots on goal (SOG per 20). Additionally, we include GP$*$, which refers to the number of games used for that player after filters were applied. Notice that the 95% confidence intervals for the mean overlap for the top 8 players. This indicates that there may not be a statistically significant difference in offensive zone possession time between the top players. Interestingly, several top players in terms of offensive zone possession rank fairly low in term of points (per 20 minutes of 5v5 ice time). So, we can infer that more goes into accruing points than just offensive zone possession.

Rank	**Name**	**Team**	**Pos.**	**GP$*$**	**oZone Pos20 (min.:sec.)**	**P20**	**P20 Rank**	**G20**	**A20**	**SOG per 20**
1	Connor McDavid	EDM	C	68	0:38 [0:35, 0:41]	1.20	1	0.31	0.89	4.4
2	Leon Draisaitl	EDM	C	72	0:37 [0:34, 0:41]	0.89	14	0.25	0.64	3.5
3	Nathan MacKinnon	COL	C	79	0:36 [0:33, 0:38]	1.10	2	0.52	0.58	5.8
4	Jack Eichel	VGK	C	58	0:36 [0:32, 0:39]	0.80	32	0.36	0.44	6.7
5	Mathew Barzal	NYI	C	72	0:35 [0:32, 0:39]	0.66	110	0.23	0.44	3.8
6	Clayton Keller	ARI	R	73	0:34 [0:31, 0:38]	0.66	113	0.33	0.33	4.0
7	Artemi Panarin	NYR	L	77	0:34 [0:31, 0:36]	0.99	5	0.54	0.45	5.0
8	William Nylander	TOR	R	75	0:33 [0:30, 0:35]	0.81	27	0.41	0.40	5.6
9	Luke Evangelista	NSH	R	71	0:31 [0:28, 0:33]	0.59	171	0.30	0.30	3.7
10	Matt Duchene	DAL	C	73	0:31 [0:28, 0:34]	0.67	107	0.24	0.43	3.2
+	League Avg.			58	0:15 [0:14, 0:15]	0.45		0.17	0.28	2.5
+	Forwards Avg.			60	0:18 [0:18, 0:19]	0.56		0.24	0.32	3.1
+	Defencemen Avg.			54	0:09 [0:09, 0:09]	0.28		0.06	0.22	1.6

Table 7. Top 10 Players ranked by 5v5 oZone Pos20 (metrics are normalized to 20 minutes of 5v5 ice time). The numbers in square brackets are the low and high of the 95% confidence interval for the per game mean.

Figure 1 plots the relationship between offensive zone possession and points, split by position. Defencemen are largely clustered in the bottom left of the plot. However, there are defencemen who have offensive zone possession and points numbers competitive with forwards. The top defenceman in terms of offensive zone possession, Quinn Hughes, has 0:29 of offensive zone possession time and 0.54 points (per 20 minutes of 5v5 ice time). Table 8 shows more details examining correlations between offensive zone possession time and some traditional measures of player success (Offensive Production). We see that offensive zone possession has a strong correlation with points when players of all positions are considered ($r = 0.70$). We note that when considering all strengths (i.e., not just 5v5 situations) the correlation is even stronger ($r = 0.78$).

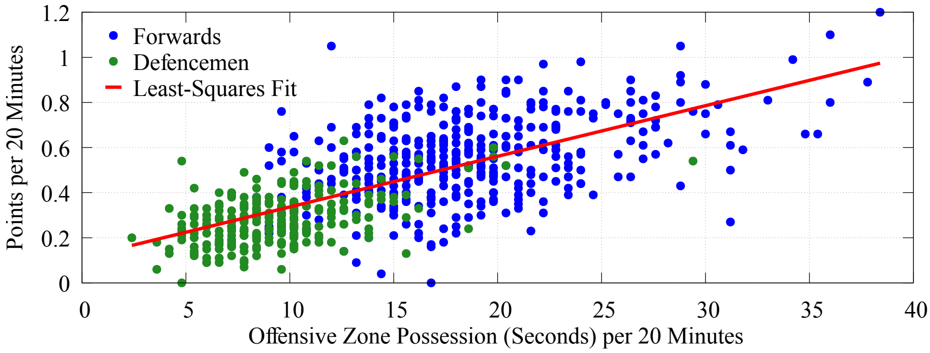

Fig. 1. Offensive Zone Possession vs Points (r = 0.70).

Offensive Production Metric	Forwards r-value	Defencemen r-value	Combined r-value
Points per 20	0.45	0.49	0.70
Goals per 20	0.30	0.30	0.64
Assists per 20	0.40	0.42	0.54
SOG per 20	0.42	0.55	0.69

Table 8. Correlations between 5v5 oZone possession per 20 minutes and offensive production.

5.4 Possession in Different Game Situations

We now turn our attention to special teams situations. Figure 2 plots CDF's of player possession time per 20 minutes for shorthanded, even strength, and power play situations, separated by position. Note that even strength differs from 5v5 as it includes 3v3, 4v4, and empty net situations. We can see that, as expected, players have less possession when shorthanded than during even strength situations, and much more possession time when on the power play. We can also see that possession time on power plays is dominated by a small number of players and that there is a larger difference among players playing the power play than in short handed or even strength situations.

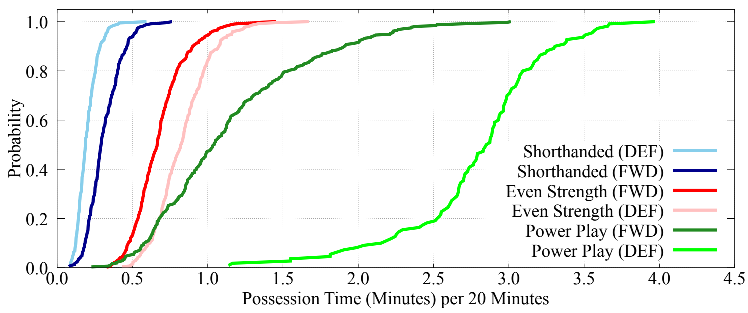

Fig. 2. Possession by strength CDF (DEF refers to defencemen and FWD refers to forwards).

Table 9 shows the correlations between offensive zone possession per 20 minutes and player offensive production metrics in both even strength and power play situations.

Offensive Production Metric	Even Strength r-value	Power Play r-value
Points per 20	0.72	0.42
Goals per 20	0.66	-0.09
Assists per 20	0.57	0.57
SOG per 20	0.61	0.14

Table 9. Correlations between oZone possession per 20 and offensive production by strength.

The correlation with shots on goal and goals goes from a moderate correlation in even strength to no correlation on the power play. However, the correlation with assists is moderate in both even strength and on the power play. Perhaps average player possession lengths is less important for scoring goals on the power play than in even strength situations. For even strength offensive zone possessions, we found that the average duration was 1.15 seconds for all possessions and 0.99 seconds for those that resulted in a goal. In contrast, power play offensive zone possessions averaged 1.32 seconds for all possessions and 0.84 seconds for possessions that resulted in a goal. This indicates a greater difference in duration, for the power play, between all possessions and possessions resulting in a goal, which may explain why there is no correlation between total offensive zone possession time and goals on the power play.

6 Metric Evaluation

We now evaluate the possession metrics using the "Meta-Analytics" framework introduced by Franks *et al.* [4], which can be used to determine the usefulness of the metrics in terms of their discriminatory power, stability over time, and independence from existing metrics. Unless otherwise stated, all metrics are for 5v5 situations and normalized per 20 minutes of ice time.

Discrimination: The discriminatory power of a metric tells us how well that metric is able to differentiate between players. We calculate discriminatory power using the method introduced by Franks *et al.* [4]:

$$\mathcal{D} = 1 - \frac{\text{Average Intrinsic Variation}}{\text{Total Between Player Variation}} \tag{1}$$

Stability: The stability of a metric tells us how stable a metric is over a period of time. We study stability by examining the correlation between the first and second halves of the season. This method differs from that proposed by Franks *et al.* [4] but maintains the philosophy of the metric.

Independence: The independence score of a metric tells us how much new information is provided by the metric compared to the others existing metrics. We compute independence scores for our metrics using the method proposed by Franks *et al.* [4]. Note that possession metrics are only included if they are the metric of interest. For example, we do not include possession metrics when calculating the independence score for points and other preexisting metrics. In addition, when computing the independence score for each possession metric we do not include any other possession metrics.

Table 10 shows the discriminatory power, stability, and independence scores for our and several other metrics. All of the possession metrics have very high discriminatory power, roughly equal to the discriminatory power of time on ice. This indicates that the possession metrics are very good at differentiating between individual players. There is a strong correlation for all possession metrics between the first half and second halves of the season. These metrics are therefore stable and could be used for predictions. Table 10 also shows that possession metrics provide quite good independence scores. For example, possession duration has a score of 0.60 indicating that only 40% of the variation in possession duration can be explained by the other metrics. This indicates that possession metrics provide information not available from the other metrics examined.

Metric	**Disc.**	**Stab.**	**Ind.**	**Metric**	**Disc.**	**Stab.**	**Ind.**
Possession Duration	0.97	0.92	0.60	Giveaways	0.73	0.55	0.80
Possession Count	0.96	0.90	0.50	Missed Shots	0.87	0.79	0.48
oZone Possession	0.98	0.93	0.44	Shot Attempts Blocked	0.89	0.79	0.56
nZone Possession	0.95	0.87	0.60	Takeaways	0.81	0.64	0.66
dZone Possession	0.99	0.97	0.53	Plus Minus	0.48	0.33	0.77
Hits	0.97	0.90	0.70	Time On Ice	0.97	0.61	0.39
Corsi Against	0.79	0.65	0.17	Points	0.73	0.55	0.07
Corsi For	0.86	0.68	0.14	Goals	0.72	0.55	0.18
Fenwick For	0.85	0.68	0.15	Assists	0.53	0.30	0.15
Fenwick Against	0.74	0.56	0.19	Shots on Goal	0.95	0.89	0.36
Blocked Shots	0.92	0.82	0.62	Penalty Minutes	0.73	0.57	0.46

Table 10. 5v5 Discrimination, stability, and independence score of the metrics in our metric set. Time On Ice and Plus Minus are not normalized. Notice that points, goals, and assists all have very low independence scores. This is expected as they are all dependent on one another. Points has an independence score of 0.50 if we exclude goals and assists from the metric set.

7 Conclusions

In this paper we introduce and analyze individual possession metrics across different zones and game situations. We find a statistically significant difference between forwards and defencemen in regard to possession time per zone, average possession length, and total possession duration. Additionally, we find that offensive zone possession time is strongly correlated with points ($r = 0.70$) and that the correlation with goals is moderate in even-strength situations ($r = 0.67$) but disappears on the powerplay ($r = -0.09$). Finally, we find that all of these possession metrics are able to effectively discriminate between players, are stable across the season, and introduce new information not captured by existing metrics.

In the future we hope to study outcomes of individual possessions (e.g., success rates) and to determine which players begin or create new possessions for their team. Additionally, we hope to examine relationships between our possession metrics and other metrics not already considered in this paper. Examples include drawn penalties, expected goals (since goals may depend somewhat on luck), zone entries, and zone exits. Computing some of these metrics requires access to per player game data that may not be available in the PPT data or via the NHL API. This data is needed to ensure that same subset of games are used when computing all metrics (due to our data cleaning and filtering process).

Acknowledgements

We thank Ben Resnick from the National Hockey League's Research and Development Team and Jonah Eisen from Rogers Communications for fruitful discussions related to this work. We also thank the anonymous reviewers for their constructive and helpful feedback. We thank Rogers Communications for providing partial funding for this research. This project is partially funded by the Natural Sciences and Engineering Research Council of Canada (NSERC) and a University of Waterloo Undergraduate Research Fellowship. We thank Neel Dayal from Rogers Communications and the National Hockey League's Information Technology, and Stats and Information Teams for making this research possible.

References

1. CASAL, C. A., MANEIRO, R., ARDÁ, T., MARÍ, F. J., AND LOSADA, J. L. Possession zone as a performance indicator in football: The game of the best teams. *Frontiers in Psychology 8* (2017), 1176.
2. COLLET, C. The possession game? A comparative analysis of ball retention and team success in European and international football, 2007-2010. *Journal of Sports Sciences 31*, 2 (2013), 123–136.
3. FERIOLI, D., RAMPININI, E., MARTIN, M., RUCCO, D., LA TORRE, A., PETWAY, A., AND SCANLAN, A. Influence of ball possession and playing position on the physical demands encountered during professional basketball games. *Biology of Sport 37*, 3 (2020), 269–276.
4. FRANKS, A. M., D'AMOUR, A., CERVONE, D., AND BORNN, L. Meta-analytics: Tools for understanding the statistical properties of sports metrics. *Journal of Quantitative Analysis in Sports 12* (2016), 151–165.
5. JAIN, R. *The art of computer systems performance analysis*. Wiley - Interscience, April 1991.
6. JAIN, R., CHIU, D., AND HAWE, W. A quantitative measure of fairness and discrimination for resource allocation in shared computer systems. *Technical Report, Eastern Research Laboratory, Digital Equipment Corporation, Hudson, MA* (September 1984).
7. LINK, D., AND HOERNIG, M. Individual ball possession in soccer. *PLOS ONE 12*, 7 (2017), 1–15.
8. LIU, H., GOMEZ, M. Á., LAGO-PEÑAS, C., AND SAMPAIO, J. Match statistics related to winning in the group stage of 2014 Brazil FIFA World Cup. *Journal of Sports Sciences 33*, 12 (2015), 1205–1213.
9. LODHI, F., NEGULESCU, S., PITASSI, M., IABONI, E., AND BRECHT, T. Individual puck possessions Part II: Speed bursts and possession times within teams. In *Proceedings of the Linköping Hockey Analytics Conference, Research Track (LINHAC)* (June 2025).
10. PITASSI, M., BRECHT, T., AND XIE, M. Puck Possessions and Team Success in the NHL. In *Proceedings of the Linköping Hockey Analysis Conference* (2024), pp. 51–66.
11. RESNICK, B. Personal communication. National Hockey League, Research and Development Team, 2024.
12. SAMPAIO, J., MCGARRY, T., CALLEJA-GONZÁLEZ, J., JIMÉNEZ SÁIZ, S., SCHELLING I DEL ALCÁZAR, X., AND BALCIUNAS, M. Exploring game performance in the national basketball association using player tracking data. *PLOS ONE 10*, 7 (2015), 1–14.
13. THOMPSON, H. The power of practice, 2014.

14. ZHANG, S., LORENZO, A., GÓMEZ, M.-A., LIU, H., GONÇALVES, B., AND SAMPAIO, J. Players' technical and physical performance profiles and game-to-game variation in NBA. *International Journal of Performance Analysis in Sport 17*, 4, 466–483.
15. ÁNGEL GÓMEZ, M., TSAMOURTZIS, E., AND LORENZO, A. Defensive systems in basketball ball possessions. *International Journal of Performance Analysis in Sport 6* (2017), 98–107.

Position Paper: New Views of Shots – Towards Measures of Net Visibility and Reachability

Evan Iaboni†, Sebastian Negulescu†, Miles Pitassi†, Amir Nazemi†, Jerrin Bright†, Vasyl Chomko†, David A. Clausi†, Sven Dickinson‡, and Tim Brecht†

University of Waterloo†
University of Toronto‡

Abstract. In this position paper, we define two new metrics: net visibility (the fraction of the net that can be seen from the perspective of the puck) and net reachability (the fraction of the net that could be reached by the puck). Reachability is slightly different from visibility because even though there might be a small portion of the net visible in a certain area (a hole), that hole may not be large enough for the puck to pass through and reach the net. We describe a framework for computing our metrics using a combination of puck and player tracking (PPT) data and video analysis (image processing). We use data and video from an NHL game to provide a proof of concept for computing net visibility and reachability. We also describe areas where more work can be done to improve the accuracy of the results and allow the computations to be fully automated. Our position is that these metrics would be valuable in studying shooter decisions and skills, goaltender and player locations and that the technologies could be used to create virtual reality images or videos.

1 Introduction

Ice hockey players often score by shooting through small spaces that appear for only a fraction of a second. We propose that one way to characterize this space is through the concept of *net visibility* which we define to be the fraction of the goalmouth that is visible from the perspective of the puck. We define net reachability to be similar to net visibility with the difference being that it accounts for the size of the puck and the fact that the puck may not be able to *reach* all areas of the net that are visible. For example, the goaltender may expose a hole that is visible but smaller than the puck.

The key insight in this paper is that we can use a combination of puck and player tracking (PPT) data from devices embedded in the players' sweaters and the puck, and video analysis to detect player locations and poses to construct a 3D-model of players and the net. Once that model is constructed, we can generate a projected image of the players onto the goalmouth (from the point of view of the puck). That resulting image can be used to determine which parts of the net are visible and which parts are obstructed. From that image we can calculate the portion of the net that is visible and, considering the size and shape of the puck, determine which portion of the net is reachable. Additionally, because we have a 3D-model of the players and net we can generate images or videos from any point of view. Two views that we think are particularly useful are the shooter's view (which can be quite different from the puck's view) and the goaltender's view. The goaltender's view allows one to understand the impact of traffic on their ability to locate and track the puck. We believe these metrics could help with coaching and player development.

2 Related Work

To our knowledge, we are the first in any sport to propose metrics that determine and quantify how much of the net is visible and reachable.

Prior work in football (soccer) examines the impact of defensive players positioned between the shooter and the goal (sometimes called traffic). That work often incorporates such information into shot prediction, or expected goals (xG) models [11] [9]. López-Valenciano *et al.* [12] examine the goaltender's perspective during free kicks. Using virtual reality to simulate and study the impact of occlusions, they found that defensive walls during free kicks impair goalkeeper performance. In contrast, we calculate metrics for shots in actual game situations, enabling realistic analysis.

Recent work in hockey analytics uses puck and player tracking (PPT) data to determine the amount of traffic in front of the net and study the impact of that traffic on shot attempts [15]. After controlling for shot angle and distance from the net, this work shows that traffic has a significant impact on the number of blocked shots and as a result, the likelihood of the shot being on goal, and the shot resulting in a goal. Interestingly, they find that most goals are scored when there is no traffic and that when shooting through traffic, the chances of scoring increase if the shot makes it through the traffic. That work uses the location of all players on the ice to determine if they would be considered in the traffic lane and does not consider how players that are closer to the puck may have a more significant impact. Additionally, the PPT data does not provide information about player orientation or pose. In contrast, our work in this paper recovers player locations and poses and can produce images to show what the traffic looks like. This includes the larger impact of players that are closer to the puck. Most importantly, we quantify how much of the net is visible and reachable.

For us to construct a 3D model of the scene, players must be detected in the video image and their stance (or pose) must be determined. Player pose estimation has been extensively explored with applications to player performance analysis and game understanding. A variety of sports, including ice hockey [2], [19], [13], [14], baseball [4], [3], and soccer [20], [21], [1], have used human pose estimation techniques. For ice hockey, GoalieNet [19] and HyperstackNet [14] are two monocular 2D techniques to estimate the poses of goaltenders and players. Recently, TokenCLIPose [2] examines pose estimation methods for players *and their stick*. 3D parametric human models are widely used for robust 3D human reconstruction and understanding [8], [6].

We utilize *parametric human models* to estimate the 3D position and shape of players, enabling the identification of the visible region of the net from the puck's perspective. These and related future contributions (e.g., more accurately recognizing player and goalie poses and equipment) could improve the accuracy of our metrics.

3 A Framework for Computing Net Visibility and Reachability

Our framework computes net visibility and reachability by explicitly reconstructing the scene at the time of the shot. We first utilize the PPT data to determine when the shot occurred and then obtain the puck and player on-ice locations and align this with the video frame in which the shot occurs. The PPT data contains x, y, and z coordinates at high frequencies (60 times per second for the puck and 12 times per second for each player on the ice). Additionally, it contains information about shots and other events that are derived from physics-based algorithms. We then construct a 3D parametric model of all the players, scaled and positioned according to the PPT data. We use the 3D scene to simulate a virtual camera to position it at the puck and then compute visibility via rasterization and reachability by simulating direct trajectories to the net.

Figure 1 shows the four main steps in our framework, each containing several sub-steps. Below, for each sub-step we: provide a description, explain how that step can be implemented (labelled Current State), and point out areas where more work could either improve the accuracy of the techniques or help in automating the computations (labelled Opportunities). Opportunities are omitted if existing approaches seem sufficient.

Although we use some manual intervention in our proof of concept computation in Section 4, technologies exist to fully automate all of these steps. Mature technologies exist for camera calibration, 3-D body pose recovery from images and video, 3-D garment modelling and recovery, and 3-D scene rendering from an arbitrary viewpoint. However, these technologies would benefit from fine-tuning for this particular application. We have not found open source solutions for all steps so we have not yet fully automated the framework. Additionally, we believe that the fine-tuning required for this application is an important consideration before publishing metrics that might be used to evaluate and compare teams and players.

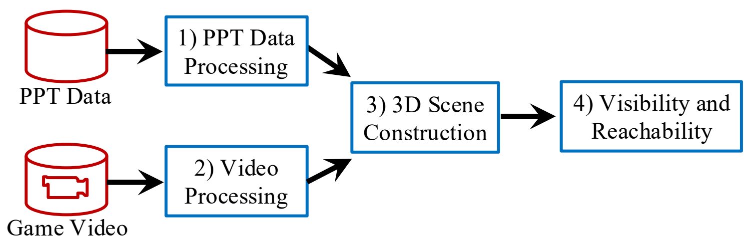

Fig. 1. The framework for computing net visibility and reachability.

As a running example, we use the shot from Alex Ovechkin's 856th career goal (scored at 16:37 of the 1st period on October 29th 2024). The frame in the broadcast for the beginning of the shot is shown in Figure 2. We use a video sequence of the goal, along with puck and player tracking data as input, to compute net visibility and reachability. For our example computation, some sub-steps were performed manually (where noted) because either existing techniques are insufficient or we were not able to find open source software that could be easily used. Note that we selected this shot from publicly available game video based on the ability of existing technologies to recognize players and poses. While existing technologies work in many cases, scenes where players are heavily obstructed can be more challenging. The steps in our framework are:

1) PPT Data Processing The first step in our framework is to process the puck and player tracking (PPT) data (unofficial data from the NHL) for the game of interest.

(a) Determine the time of the shot in the PPT data (call this T_s). ***Description:*** To determine what parts of the net are visible and reachable, we need to know the precise time of the shot, T_s. ***Current State:*** Shot events and precise UTC times are labelled in the PPT data, however, some shots are undetected and for some shots adjustments to the shot time is required (to ensure that the puck is not too far from the shooter at the time of the shot) [17]. ***Opportunities:*** Automating shot detection and precise release times are continually being improved.

(b) Obtain on-ice location for players in the PPT data. ***Description:*** To correctly place the 3D models on the rink we need their precise on-ice locations. ***Current State:*** The on-ice (x, y, z) coordinates of all players at the time of the shot is available from PPT data. ***Opportunities:*** Although the precision of the PPT location coordinates, is typically within a few inches, improving that fidelity could increase the accuracy of our metrics, as could the use of tracking devices on sticks or improved optical tracking.

2) Video Processing The second step involves finding the exact frame of game video that corresponds to the shot release time and then identifying players in the video.

(a) Find the frame of the shot in the game video (call this F_s). ***Description:*** We need to synchronize the PPT data for the time of the shot, T_s, with the frame of the shot in the game video. There may be some ambiguity here because the PPT data contains data for every one-hundredth of a second (interpolated), while game video is typically recorded at 30 frames per second. ***Current State:*** In our example, we manually determine F_s by looking for the last frame where the puck is touching the stick before release. ***Opportunities:*** For our approach to scale, we require a method to automatically determine F_s and to synchronize T_s and F_s.

(b) Identify the players in the game video. ***Description:*** We need player identification to place each player's pose at the correct on-ice coordinates. We also use Player identifications to appropriately scale players using their height from the PPT data. ***Current State:*** In our example, we use a player tracking algorithm developed by Prakash *et al.* [16] to get approximate on-ice locations of players using the game video. These approximate locations are then matched to our precise PPT locations by calculating the Euclidean distance between the video-based estimates and the PPT coordinates. Player matching allows us to scale each player properly and provides a check that T_s and F_s are matched. We believe that this step can be automated to find the appropriate frame in the game video using the time of the shot event in the PPT data. The player tracking algorithm also generates bounding boxes for each player, which we later match with the bounding boxes provided by the 3D player modelling software, allowing us to assign the proper mesh (i.e., player model) to each player. ***Opportunities:*** Player tracking is an active area of research and new techniques are being developed to better handle player occlusions in the game video. 3D models could be tailor-made for each player. Placing the player's model in their on-ice location, in the proper pose, could eliminate scaling.

3) 3D Scene Reconstruction In the third step, we utilize our processed PPT and game video data to reconstruct a 3D representation of the scene.

(a) Build a 3D model of players and the goaltender. ***Description:*** The idea of net visibility relies on the assumption that we can reconstruct accurate 3D models of the players and goaltenders using game video. These models capture the full shape of the body, which is crucial for determining occluded parts of the net from the puck's perspective. ***Current State:*** Using the game video frame at the time of the shot (F_s) and an image recognition tool, we retrieve players and their poses. In our example, we use open-source software called 4DHumans [8]. This provides us with a 3D model for all players in the frame, including the goaltender. Figure 2(a) shows the original game video frame and Figure 2(b) shows the estimated 3D model of all the players. Players shown in grey do not impact net visibility or reachability. We currently omit the players' sticks since, to our knowledge, the only work in pose reconstruction for hockey sticks has been in 2D [2]. ***Opportunities:*** Human 3D pose reconstruction is an active area of research. The inclusion of stick positions would also benefit our metrics. Due to constraints on how sticks can be held, and because we have information about whether a player is a left or right handed shot (in the PPT data), we believe that this should not be too difficult.

(a) Original video image (b) Image with player and pose recognition

Fig. 2. (a) Alex Ovechkin's 856th career goal. (b) approximate meshes of the players and goaltender overlaid. The goaltender is coloured in red, the defender in blue, and Ovechkin in orange. These players reappear in later visualizations as the same colour. NHL EDGE visualization for this goal: nhl.com/ppt-replay/goal/2024020151/172.

(b) Scale the players. ***Description:*** The reconstructed player models are not properly scaled relative to the size of the rink. As a result, the size of the players needs to be scaled. ***Current State:*** We scale each player mesh using the z-coordinate (the distance from the ice surface), from the PPT data. This value is relative to the player's right shoulder. 4D-Humans outputs each player as a SMPL mesh [10] allowing us to obtain coordinates for the right shoulder. SMPL stands for "Skinned Multipurpose Linear Model" and it produces models that include bodies with clothing, rather than stick figures. We then scale the mesh so that its right shoulder aligns with the height of the player's in-sweater device in the PPT data. Each player's height is available in the PPT data. ***Opportunities:*** Systems like SMPL could be fine-tuned to include player's equipment (including goaltenders).

(c) Determine the camera parameters. ***Description:*** The reconstructed 3D human models generate 3D meshes (of each player) relative to the camera model (using the angle of the camera used to capture the video). To correctly orient players on the rink, we need to determine several camera parameters (some examples include the focal length, as well as the angle and height relative to the ice surface). ***Current State:*** In our example, we approximate these camera's parameters manually. However, techniques to calibrate sports broadcast cameras do exist [5][7][18]. See Figure 3 for a comparison of the broadcast video image with players and poses recognized and the view from the puck's perspective with orientations corrected. ***Opportunities:*** Sports broadcast camera calibration is an active area of research.

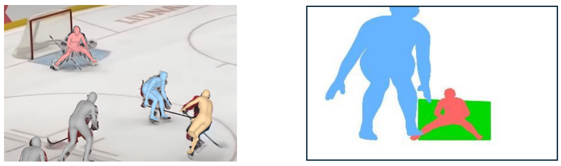

(a) (b)

Fig. 3. (a) The view from the broadcast camera's perspective. (b) The view from the puck's perspective with orientations corrected.

4) Net Visibility and Reachability In the final step of our framework, we compute measures of net visibility and reachability using the reconstructed 3D scene.

(a) Add the net to the 3D model. ***Description:*** To determine which parts of the net are visible and reachable we need to place the net into the 3D scene. ***Current State:*** The net's size and location are known, making it straightforward to add the net to the scene. For the purpose of visibility and reachability, we only need to construct the net opening. In our proof-of-concept implementation, the net opening is comprised of 10,000 non-overlapping, equal sized polygons, which are used to compute visibility and reachability. Figure 4(a) provides an example of how a smaller number of polygons could be mapped to the net opening.

(b) Adjust the camera view to that of the puck. ***Description:*** To determine net visibility, we need to be able to view the scene from the perspective of the puck (using the centre of the puck). ***Current State:*** We position the camera in 3D space at the location of the puck. We then perform perspective projection using a pinhole camera model with the optical axis aligned with the centre of the net opening.

(c) Determine the Visibility of the Net. ***Description:*** We want to calculate the percentage of the net that is open from the perspective of the puck. Obstructions between the puck and the net, like players and the goaltender, reduce net visibility. ***Current State:***

Our proof-of-concept implementation uses perspective projection to map polygons to the puck's view of the plane formed by the net opening. We perform rasterization to convert these polygons into pixels. By examining the resulting image we can determine which pixels are the open net and which pixels are obstructions. Net visibility is the number of net polygons mapped to unobstructed pixels, divided by the total number of net polygons.

(d) Determine the Reachability of the Net. ***Description:*** We recognize that there can be areas of the net that are visible but too small for the puck to fit through. Therefore, we introduce a new metric, net reachability, defined as the fraction (or percentage) of the net that the puck can pass through unobstructed. ***Current State:*** We construct an algorithm to determine net reachability by dividing the net into 10,000 non-overlapping, equal sized polygons. For each polygon, we calculate the puck's trajectory when aimed at the polygon's centre, assuming a linear path with no puck tilt or flutter. If there are any obstructions in the trajectory, we deem the polygon "not reachable". Finally, we define net reachability as the percentage of polygons that are reachable. See Figure 4(b) for a visualization of net visibility and reachability calculated for Ovechkin's goal. Notice there is a small opening between the goaltender's left arm and his chest (labelled R_3) which is visible but not reachable (because the puck is too small to fit through that gap). This highlights the importance of net reachability.

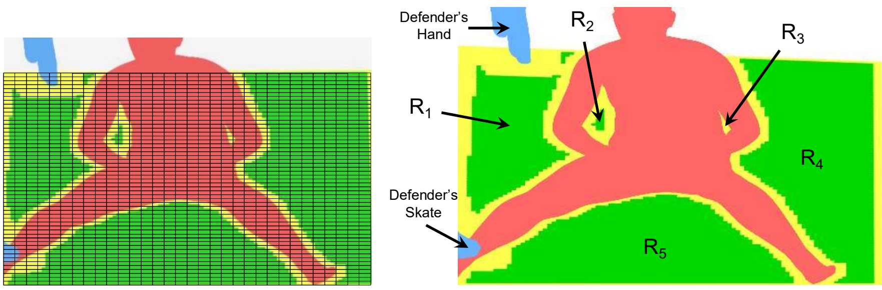

(a) Finding reachable portions (b) Visible and reachable regions

Fig. 4. View from the puck (zoomed and cropped version of Figure 3(b)). In image (b) the net is not rectangular because the shot is coming from an angle to the left of the net and from that perspective the right post appears shorter than the left post because it is farther away. Image (a) has been modified to show a straight on view of the net to more easily illustrate the concept of dividing the net into polygons of equal sizes. Green indicates unobstructed net, yellow signifies visible but not reachable, and red denotes neither visible nor reachable. The puck entered the net in Region R_5. One can see the impact of the defender's hand on reachability in R_1. It has a larger influence on reachability than the post because it is closer to the shooter. The position of the defender's skate may be incorrect due to inaccuracies in the pose recognition software and the fact that it does not know that a person may be wearing skates. Additionally, the player and pose recognition software does not handle player equipment. Research is being conducted to recognize goaltender poses and equipment that could be used to augment current approaches [19]. Building 3D models of goaltenders with their equipment and placing the model in the 3D scene at the specified location may be another way to improve accuracy.

4 Example Computation of Net Visibility and Reachability

For our proof-of-concept implementation we use a simplifying assumption that if any part of the puck hits a post or obstruction, it will not reach the net. This could be easily modified to assume, for example, that shots with half of the puck inside the post reach the net. Interesting future work would be to better understand and model the interaction between the puck and a post and the impact of the puck's spin (spinning towards the interior or exterior of the net). Naturally one would also want to study interactions with all types of obstructions (e.g., the crossbar, players, goaltenders, and equipment).

Figure 4(b) shows the different regions of the net that are visible and reachable, along with labels for each of the regions. Green signifies areas that are visible and reachable, yellow denotes visible but not reachable areas and red shows areas that are obstructed and therefore, are not visible or reachable. Table 1 shows the results of our computations for the percentage of the net that is visible and reachable in each region as well as the overall values. Notice that regions R_4 and R_5 are separate regions for reachability but become merged for visibility. This occurs because there is a small area beneath the goalie's left foot that is visible but not reachable, connecting R_4 and R_5.

To provide a high-level understanding of the idea of reachability and how it could be computed, we superimpose a grid of equal sized polygons onto the image of the net (see Figure 4(a)). To compute reachability we count the number of polygons that are green and compare that with the total number of polygons comprising the net. Table 1 shows that for this shot, 65.97% of the net was visible and 50.32% was reachable (Overall).

Region	R_1	R_2	R_3	R_4	R_5	$R_4 + R_5$	Overall
Visible	14.97 %	0.97 %	0.03 %	—	—	50.0 %	65.97 %
Reachable	8.82 %	0.15 %	0.00 %	22.92 %	18.43 %	—	50.32 %

Table 1. Percentage of the net that is visible and reachable for each region, as well as overall. Note: regions R_4 and R_5 are part of the same visible region, but form separate reachable regions. There are two separate reachable areas in visible region R_1; they are both part of R_1.

5 Potential Applications

The technology required to implement these metrics, a 3D model of players along with their and the puck's locations, could be used to produce virtual reality video simulations of any window of time (not just shots) from any desired point of view (or continually changing points of view). This could be used to increase fan engagement or the construction of new metrics that take advantage of the 3D scene. In Figure 5 we show the 3D scene for our running example from the point of view of: the puck (Figure 5(a)), the shooter (Figure 5(b)), and the goaltender (Figure 5(c)). Notice the difference in visibility between the puck's view and Ovechkin's view.

The methods used in this paper to construct a 3D scene are generalizable to other sports. We envision that our net visibility and reachability metrics could also be applied to sports like football (soccer) and lacrosse, with additional considerations such as individual player's ability to bend (curve) the ball.

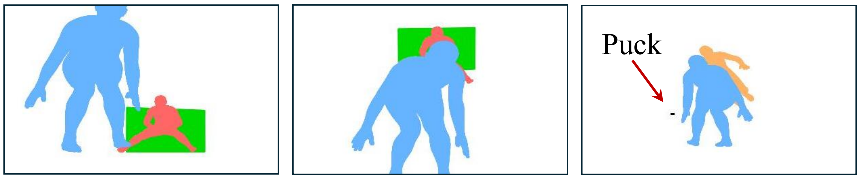

(a) Puck's view (b) Shooter's view (c) Goaltender's view

Fig. 5. Using the 3D model of the scene to generate different views. Note that in (c) Ovechkin (the light orange player) is barely visible, due to the defender's (blue player's) position, highlighting the value of this view.

Our position is that, given enough samples for each player and goaltender, our net visibility and/or reachability metrics could be beneficial to players, coaches, and executives to improve offensive and defensive tactics and overall team performance. The metrics provide information about how much of the net each defender is obstructing, aiding in defensive positioning and decision making. It can be used to evaluate which body positions (poses) are most effective for shot blocking, such as standing versus kneeling. Metrics could even be adjusted to account for whether players are attacking or defending. For example, one could assume that attacking players would move out of the way of a shot, and by removing them from the scene, make the portion of the net they are occluding reachable.

Additionally, our metric could be used to evaluate and provide insights into a player's shooting decisions and precision. Because hockey is dynamic and since we can construct a 3D model at any point in time (and view it from any viewpoint), we could examine whether players are shooting at appropriate times. This could be done by comparing the portion of the net that is reachable at instances in time prior to the actual shot where the player could shoot with that portion of the net that is reachable from the time of the shot. This would allow one to study, for example, if more or less of the net was reachable if the shot had been taken sooner. Similarly, one could reposition the shooter (and other players) to show how much of the net would have been reachable had the shooter taken a different path, and comparing that with the reachable portion of the net for the shot the player took. Examining whether players are shooting at smaller or larger reachable regions would also be informative, especially if a player often shoots at and misses smaller regions when there are larger regions that are reachable. Another possibility would be to study shooting skills by determining which players are able to score when there are only small reachable regions available (e.g., on the short side between the goaltender's head and the post).

These metric could also be used to analyze goaltender positioning to identify if they consistently leave regions of the net open and vulnerable to exploitation, determine if they "cheat" by presenting areas where they want a shooter to shoot because they believe they can make the save, and understanding whether they are successful. A simple but powerful example would be to show how the reachable portion of the net changes if the goaltender more aggressively moves towards the shooter. This could provide quantitative insights into goaltender positioning. Additionally, generating the goaltender's view would permit us to compute the amount of time (or portion of the shot duration) that the goaltender would have been able to see the puck. For example, the puck was visible for 80 milliseconds from the time of release to the end of the shot (or 44% of a shot that took 180 milliseconds to reach the goaltender). This would provide insights into goaltenders' abilities to make stops on shots through traffic or possibly absolving them of fault for not stopping shots that they could not see.

Some simple examples of these applications are provided in the Appendix. There we demonstrate how potential changes in goaltender positioning impact net visibility and reachability. Additionally, we show how these metrics would be impacted if the defender were positioned more directly in the shooting lane.

6 Discussion

We recognize that 3D human pose estimation, sports camera calibration, and video player tracking are active areas of research. Our method of calculating net visibility and reachability relies on the precision of these tools and their ability to generalize to hockey. In particular, 3D human pose estimation does not identify the player's equipment or stick, which is a limitation we would like to address in future work. Moreover, in many shots, influential players or the goaltender are obstructed from the broadcast camera's view or positioned outside the frame. This may hinder our ability to capture accurate 3D poses. Furthermore, for our net reachability metric, we assume the puck's trajectory is linear (with no rise, fall or curve). In reality, a puck's trajectory is parabolic, but from close distances or with the high speed of most shots, it is likely sufficiently close to linear. The puck may also wobble or reach the net tilted off axis, violating our no tilt assumption used when computing reachability. There are also factors not currently captured by net visibility and reachability. One such factor is the difficulty of the shot. For example, 5% of the net being visible from 60 feet away might be thought of as "more difficult" than 5% of the net being visible from 10 feet away. Likewise, shots from sharp angles may be considered differently than from the slot. Imagine an open net from the slot and an open net from a sharp angle, both would have net visibility values of 100%. However, from the sharp angle there may be only a small sliver of the net opening to shoot at. In this example, the size of the net opening differs between the two shot locations and that may not be fully reflected in our metrics. In the future, we hope to add a metric for "scoreability" to accounts for such factors. One final challenge is devising techniques to validate the values obtained from our computations.

Acknowledgements

We thank Ben Resnick from the National Hockey League's Research and Development Team for his insightful discussions, comments and feedback on several drafts of this paper. We thank Jonah Eisen from Rogers Communications and Craig Kaplan from the University of Waterloo for fruitful discussions related to this work and the anonymous reviewers for their helpful feedback. We thank Rogers Communications and Stathletes, Inc. through the Mitacs Accelerate Program for providing partial funding for this research. This project is also partially funded by the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Cheriton School of Computer Science at the University of Waterloo via Undergraduate Research Fellowship funding. We thank Neel Dayal from Rogers Communications and the National Hockey League's Information Technology, and Stats and Information Teams for making this research possible.

References

1. AFROUZIAN, R., SEYEDARABI, H., AND KASAEI, S. Pose estimation of soccer players using multiple uncalibrated cameras. *Multimedia Tools and Applications 75* (2016), 6809–6827.
2. BALAJI, B., BRIGHT, J., CHEN, Y., RAMBHATLA, S., ZELEK, J., AND CLAUSI, D. Seeing beyond the crop: Using language priors for out-of-bounding box keypoint prediction. *Advances in Neural Information Processing Systems 37* (2024), 102897–102918.
3. BRIGHT, J., BALAJI, B., CHEN, Y., A CLAUSI, D., AND ZELEK, J. S. Pitchernet: Powering the moneyball evolution in baseball video analytics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 3420–3429.
4. BRIGHT, J., CHEN, Y., AND ZELEK, J. Mitigating motion blur for robust 3D baseball player pose modeling for pitch analysis. In *Proceedings of the 6th International Workshop on Multimedia Content Analysis in Sports* (2023), pp. 63–71.
5. CHEN, J., AND LITTLE, J. J. Sports camera calibration via synthetic data, 2018.
6. DWIVEDI, S. K., SUN, Y., PATEL, P., FENG, Y., AND BLACK, M. J. Tokenhmr: Advancing human mesh recovery with a tokenized pose representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 1323–1333.
7. FALALEEV, N. S., AND CHEN, R. Enhancing soccer camera calibration through keypoint exploitation. In *Proceedings of the 7th ACM International Workshop on Multimedia Content Analysis in Sports* (Oct. 2024), MM '24, ACM, p. 65–73.
8. GOEL, S., PAVLAKOS, G., RAJASEGARAN, J., KANAZAWA, A., AND MALIK, J. Humans in 4D: Reconstructing and tracking humans with transformers. In *ICCV* (2023).
9. LINK, D., LANG, S., AND SEIDENSCHWARZ, P. Real-Time Quantification of Dangerousity in Football Using Spatiotemporal Tracking Data. *PLOS ONE 11*, 12 (2016).
10. LOPER, M., MAHMOOD, N., ROMERO, J., PONS-MOLL, G., AND BLACK, M. SMPL: a skinned multi-person linear model. *ACM Transactions on Graphics 34* (11 2015).
11. LUCEY, P., BIALKOWSKI, A., MONFORT, M., CARR, P., AND MATTHEWS, I. Quality vs Quantity: Improved Shot Prediction in Soccer using Strategic Features from Spatiotemporal Data. In *MIT Sloan Sports Analytics Conference* (2015).
12. LÓPEZ-VALENCIANO, A., MCROBERT, A., SARMENTO, H., REINA, M., PAGE, R., AND READ, P. A goalkeeper's performance in stopping free kicks reduces when the defensive wall blocks their initial view of the ball. *European Journal of Sport Science* (2021).

13. Neher, H., Fani, M., Clausi, D. A., Wong, A., and Zelek, J. Pose estimation of players in hockey videos using convolutional neural networks. In *2017 Ottawa Hockey Analytics Conference (OTTHAC), Ottawa, Canada* (2017).
14. Neher, H., Vats, K., Wong, A., and Clausi, D. A. Hyperstacknet: A hyper stacked hourglass deep convolutional neural network architecture for joint player and stick pose estimation in hockey. In *2018 15th Conference on Computer and Robot Vision (CRV)* (2018), IEEE, pp. 313–320.
15. Pitassi, M. Puck Possession and Net Traffic Metrics in Ice Hockey. M.Math. Thesis, Cheriton School of Computer Science, University of Waterloo, May 2025. https://cs.uwaterloo.ca/ brecht/theses/Pitassi-MMath.pdf.
16. Prakash, H., Shang, J. C., Nsiempba, K. M., Chen, Y., Clausi, D. A., and Zelek, J. S. Multi player tracking in ice hockey with homographic projections, 2024.
17. Radke, D., Lu, J., Woloschuk, J., Le, T., Radke, D., Liu, C., and Brecht, T. Analyzing passing metrics in ice hockey using puck and player tracking data. In *Proceedings of the Linköping Hockey Analytics Conference 2023 Research Track, LINHAC 2023, Linköping, Sweden, June 7-9, 2023* (2023), pp. 25–39.
18. Sha, L., Hobbs, J., Felsen, P., Wei, X., Lucey, P., and Ganguly, S. End-to-end camera calibration for broadcast videos. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2020), pp. 13624–13633.
19. Shahi, M., Clausi, D., and Wong, A. GoalieNet: A multi-stage network for joint goalie, equipment, and net pose estimation in ice hockey. *arXiv preprint arXiv:2306.15853* (2023).
20. Skyttner, A. Multi-person pose estimation in soccer videos with convolutional neural networks, 2018.
21. Yeung, C., Ide, K., and Fujii, K. Autosoccerpose: Automated 3D posture analysis of soccer shot movements. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2024), pp. 3214–3224.

Appendix

Figure 6 is used to provide an example of one type of application for our framework and metrics. This example focuses on one aspect of goaltender positioning, depth relative to the net. The figure shows an example of different views from the puck's perspective with the goaltender located in slightly different positions. The left image, labelled "-4 feet", shows net visibility and reachability with the goaltender moved 4 feet closer to the net. The centre image, labelled "Original" shows the goaltender's original position. The right image, labelled "+4 feet", shows the goaltender moved 4 feet closer to the shooter. Below each image, V denotes the portion of the net that is visible R denotes the portion of the net that is reachable, D_{puck} denotes the distance from the goaltender to the puck and D_{net} denotes the distance from the goaltender to the centre of the net. Note that the -4 feet and +4 feet are changes in the x value in the x,y coordinate system, which is why D_{net} does not change by exactly 4 feet.

As one expects, if the goaltender is positioned closer to the net they obstruct less of the net, resulting in larger visibility and reachability values. If the goaltender is positioned closer to the puck visibility decreases as does reachability. Note that in reality, a goaltender may change their stance when they are farther into or out of the net, which could also alter these metrics. Note that the defender's impact does not change as their location remains the same in each case.

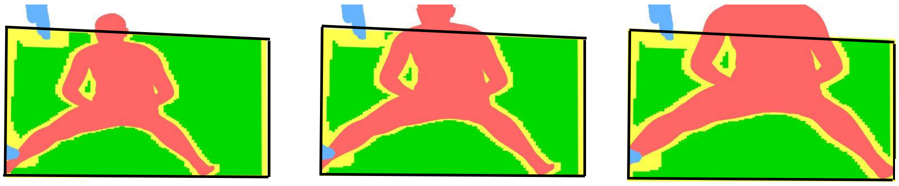

-4 feet
V = 71 %
R = 57 %
D_{puck} = 34.3 ft
D_{net} = 2.4 ft

Original
V = 66 %
R = 50 %
D_{puck} = 30.0 ft
D_{net} = 6.5 ft

+4 feet
V = 62 %
R = 45 %
D_{puck} = 25.7 ft
D_{net} = 10.9 ft

Fig. 6. Example application of net visibility and reachability. Comparing metrics with different goaltender locations.

Figure 7 provides another example of a potential application. In this case we demonstrate how net visibility and reachability would change if the defender were positioned more directly in the shooting lane. The left image shows the original position of the defender with only a part of their hand and skate seen in the left side of the image, along with the net visibility and reachability values. The right image shows how net visibility and reachability decrease substantially if the defender is more directly in the shooting lane. Note that while this increases the chance of blocking the shot, it may also partially obstruct the goaltender's view. Since one can not see all of the goaltender's face from the puck, they may not have a clear line of sight to the puck with both eyes.

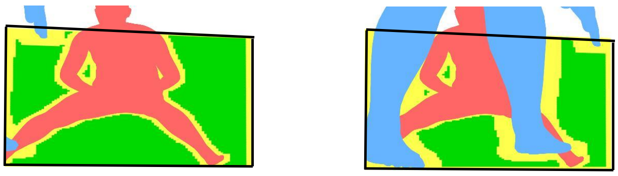

Original
V = 66 %
R = 50 %

D (#55) in shooting lane
V = 46 %
R = 26 %

Fig. 7. Example application of net visibility and reachability. Comparing metrics with a different defender location.

We believe that being able see and quantify changes in net visibility and reachability with different locations for players and goaltenders can provide valuable insights to goaltenders, defenders, shooters, coaches and fans.

Individual Puck Possessions Part II: Speed Bursts and Possession Times within Teams

Fauzan Lodhi, Sebastian Negulescu, Miles Pitassi, Evan Iaboni, and Tim Brecht

Cheriton School of Computer Science, University of Waterloo

Abstract. In ice hockey, handling and keeping control of the puck (possession) are valued skills. In this paper we study several metrics of individual player puck possessions from 2023-24 regular season NHL games. These metrics include players' speed while carrying the puck, and the distribution of puck possession times for players within their team (i.e., does a team have a few players who have a larger share of possession time or are times more equally distributed). Our goal in this paper is to examine and highlight different skills and roles related to puck possession and to design metrics that might be helpful in roster construction and/or creating line combinations.

1 Introduction

In ice hockey, being able to possess and handle the puck is a highly valued skill. Players with possession of the puck may advance the puck towards the opponent's end, set up plays, and prevent their opponents from making plays. We believe that understanding which players are able to obtain and maintain possession of the puck and what they do when they have the puck can provide critical information for valuing players and creating line combinations.

Using puck and player tracking data obtained from the National Hockey League, we utilize data from regular season games from the 2023-24 season and study individual player possessions. While puck handling skills are important for goaltenders, the types of metrics we consider are not designed to evaluate goaltenders. As a result, we do not include goaltenders in any of the analyses conducted in this paper, and henceforth, when the word "players" is used it is referring to skaters. We begin by examining the speeds with which skaters are able to carry the puck. Since many players are capable of reaching high top speeds we focus on which players are able to consistently reach high speeds while carrying the puck.

We later examine, on a per game basis, the distribution of the amount of time individual players possess the puck within their team. The objective it to understand the degree to which a smaller number of players dominate team possessions or whether possessions are distributed more equitably across players on the team. We believe that these new metrics provide insights into individual player's skills and/or roles and that these insights may be valuable when constructing rosters and/or line combinations.

From the analyses described above we make the following key contributions:

– We devise a methodology for preparing, cleaning and filtering games as well as (when appropriate) devising filters to exclude some players who may not have sufficient opportunities for us to obtain representative metrics.

- We find that, per 20 minutes, some players have significantly more bursts of 20 miles per hour (29.3 ft/s) or greater while carrying the puck (Bursts20), than others. For instance, Nathan MacKinnon averages more than 7 times as many 20+ MPH bursts per 20 minutes (3.35) than an average forward (0.44). We also find a large and significant difference between forwards and defencemen and believe that Bursts20 is a good indicator of players' roles.
- We evaluate individual contributions to team possessions by using Jain's Fairness Index to measure the distribution of possessions across all skaters within each team. For example, we find that the Florida Panthers have the most equitable distribution (index=0.85), while the Vancouver Canucks have the least equitable distribution (index=0.70). We observe significant differences between teams and believe this offers insight into roster structures and offensive styles.

2 Related Work

Much of the research studying possession in sports has focused on team possessions [1] [2][9][16][11]. Studies examining individual possessions have mainly concentrated on basketball [3][14][15] and football (soccer) [8]. These studies focus on how many times a player possesses the ball per game and how long they possess the ball.

One ice hockey study manually tracked the possession time of top players (e.g., Joe Sakic and Mike Modano) during the men's 2002 Olympic Ice Hockey games [4]. The results showed that the top players averaged one minute and seven seconds with the puck per game. Similarly they found that top players in the USA Tier 1 Youth National Championships averaged one minute and six seconds per game. They use these results to argue that youth hockey should place more emphasis on practice rather than games, to provide more opportunities for players to develop puck handling skills.

In Part I of this study, Iaboni *et al.* [5] examine the average time of each player's possession, the average time of possession per game and the average number of possessions (all in 5v5 situations). They normalized all metrics to 20 minutes of ice time because the metrics were strongly correlated with ice time. They found that the top player had possession for 1:37 (one minute and thirty seven seconds) per 20 minutes, with the league average being 0:43. They found that when considering players by position group (by grouping defencemen separately from forwards) there were only weak correlations between a player's possession time per 20 minutes and traditional measures of success per 20 minutes, measured by offensive production metrics (e.g., goals, assists, points and shots). They also examined the distance players travel with the puck during possession, finding that the top player averaged 36 feet per possession which is significantly greater than the league average of 20 feet.

In addition, they also studied offensive zone (OZ) possession time per 20 minutes and found that when considering all players combined, OZ possession time correlates strongly with points per 20 minutes ($r = 0.70$). However, this correlation may be mainly capturing differences by position group as there were only weak correlations among forwards ($r = 0.45$) and among defencemen ($r = 0.49$). The top eight players in terms of OZ possession time were statistically similar with 95% mean confidence intervals ranging from [0:35, 0:41] for the top player to [0:30, 0:35] for the eighth highest ranked player. League averages were [0:14, 0:15]. So top players averaged more than twice the OZ possession time, in 5v5 situations, than the average player.

In this paper we build on and extend the complementary work in Part I by Iaboni *et al.* [5], described above. We examine an additional set of metrics that includes the speed at which players carry the puck and we study whether a team's time of possession is concentrated among a few players or is more evenly distributed across all players.

3 Background

3.1 Definition of Individual Puck Possession

The NHL defines two types of individual puck possession. The first occurs when a player touches the puck consecutive times, with at least one of those touches occurring when the puck is on the ice. For the second type, one-touch actions are also considered possession (e.g., one-touch passes and one-timers). Each possession is credited to an individual. The time between individual possessions, such as when the puck is travelling from one player to another during a pass, is not considered part of an individual's possession. Instead, an individual possession is deemed complete: at the end of the player's final touch (e.g., a shot, pass, or area-play), when another player establishes possession (e.g., a steal), or when the puck travels a substantial distance away from the possessing player (e.g., a puck loss). We utilize these types as the definition of individual puck possessions in this paper. In prior work we studied team possessions, defining team possession as the period of time players on the same team have consecutive possessions, including the time for a pass to reach another player [11]. See that work for a more precise description of how individual and team possessions are defined.

3.2 Dataset Overview

Our research is conducted using the NHL's proprietary puck and player tracking (PPT) data, which records puck and player locations at high frequencies (60 Hz and 12 Hz, respectively). Along with the PPT data, the NHL provides individual possession models, equipped with possession information using the definitions provided in the previous section. Moreover, these datasets also include automated event detection and labelling information. These event labels include but are not limited to: shots, passes, and area plays (e.g., dumps-ins and dump-outs), This data is interpolated by the Delayed Interpolated Smoothed Hundred-Hertz (DISH) stream to provide information about puck and player locations every one-hundredth of a second. Note that this data is considered unofficial by the NHL. We also use data from the NHL API for the games included in our analysis to get official player statistics, like goals, assists and points (which are used to examine correlations between our metrics and those statistics).

4 Dataset Cleaning and Filtering

4.1 Preparing, Cleaning, and Filtering Games

In previous work we devised techniques for analyzing puck possessions by individuals and teams and examined relationships between team possession and team success [12].

We utilize the data cleaning and filtering methods in our previous work on team possessions to conduct our analysis of individual possessions [12]. This includes merging individual possession data with game information to provide additional game context and details such as power play information, score differential, and puck and player locations. We then we address several issues with that data that include: adjusting the start and end times to account for time clock resets (e.g., plays where the clock is reset due to a video review like an offside), ensuring that possessions adhere to active game play intervals, removing duplicate possessions, fixing abnormal data entries (e.g., out of sequence data) and adjusting some possessions that contain excessive distance between the puck and the possessor.

After cleaning and filtering, we found a few issues that compromised game information and data accuracy. As a result, we removed games with erroneous data for more than 4% of the game duration, or 4% of a team's possession time. After this filtering (118 games) we were were left with 91% of the league's regular season games.

4.2 Filtering Individual Players

We apply filters to the remaining 1,194 games to exclude players for which there was insufficient data. From these 1,194 games, players are excluded if they played fewer than 10 games or had less than 10 minutes of 5v5 ice time per game. In the 1,194 games studied, 921 players participated in one or more games and 250 were excluded, leaving 671 players remaining. We believe these 671 players capture a representative sample of regularly participating players since the expected number of players (i.e., excluding goaltenders) given no roster changes throughout the entire season would yield 576 players (32 teams x 18 players per team).

5 Speed with Possession

In this section we study player speeds during possessions with the goal of identifying players that carry the puck at high speeds. We focus exclusively on play during 5v5 situations because it is more indicative of regular play and avoids giving an advantage or disadvantage to players who spend more time in short-handed, power play, 4v4, 3v3, or empty net situations. We evaluate puck-carrying speeds using three metrics, inspired by the data available on the NHL EDGE website [10] that reports players top speeds and bursts of speed. Our metrics only consider player speeds when they have possession of the puck and we report all speeds in feet per second (ft/s) as we believe that this allows one to envision how much on-ice distance is being covered, given that NHL rinks are 200 feet in length and there are 50 feet between the two blue lines. Specifically, the metrics that we examine are the average number of 20+ MPH (i.e., 29.3+ ft/s) bursts reached by a player per 20 minutes (referred to as "Bursts20"), top speed across the entire season, and an average (across all games) of the top speed obtained in each game (Avg. Top Speed).

Note that Bursts20 is different from 20 MPH+ bursts reported on the NHL EDGE web site [10]. According to that site, "bursts measure the number of times a skater achieved a sustained speed above a given threshold". As noted, for Bursts20 a player must possess the puck for one second or more and we normalize the number of bursts to 20 minutes, to ensure that values are not skewed towards players with more ice time.

5.1 Data Cleaning

To capture possession speeds and draw fair comparisons, players must have enough opportunity within a possession to generate high speeds. Moreover, a player must have sufficient opportunity within a game and across the season to record high speeds. Therefore, we only consider possessions of one second or longer to capture "puck-carrying" possessions. Furthermore, we only consider games in which a player has five or more such possessions, and players with ten or more such games. Collectively, these filters exclude short possessions with insufficient puck-carrying time, and players that may not have had enough opportunites to reach high speeds in a game or over the season. After applying these filters, we are left with 663 players and an average of 53 games used per player. The PPT data provides speed computed using 12 readings per second and then "smoothed" to account for missed readings and the volatile movement possible with the tracking device over short time intervals [13]. Note that Bursts20 and average game top speed are calculated by game and then reported as an average. Also note that we record at most one burst per possession, thus if a player reaches 20+ MPH then their speed drops below and speeds up to 20 MPH (or more) during the same possession, we count this as a single burst.

5.2 Player Speeds

Table 1 shows the top 10 forwards and top 10 defencemen each sorted by Bursts20 during 5v5 situations. We sort by Bursts20 as this provides insights into which players carry the puck at high speeds more often. The ability to consistently carry the puck at high speeds (Bursts20) seems, to us, more valuable and more informative than top speed and average top game speed. When examining the data we notice that there are many well-known, highly-regarded players who average very few or zero puck carrying bursts of 20+ MPH. This is likely because those players have different roles and/or skill sets (e.g., play makers, goal scorers, or defensive-oriented players, to name a few). For example, Alex Ovechkin (WSH), Mitchell Marner (TOR), Rasmus Dahlin (BUF), Jason Robertson (DAL), and Adam Fox (NYR) have low Bursts20 averages but provide value to their respective teams in other ways. We believe that Bursts20 provides insights for teams and coaches looking to find and leverage players who can consistently carry the puck with speed when considering roster management and line combinations, however it is by no means a requirement for players to contribute to their teams (as different players may fill different roles). In addition to average Bursts20 and 95% confidence intervals for the average, Table 1 also shows the number of games used after filtering (GP⋆: Games Played and not filtered), top speed, average per game top speed, as well as league and position averages (the bottom rows).

Top speed and average game top speed both suggest that defenders can carry the puck at fairly similar speeds to forwards (see the averages shown at the bottom of the table). We find that forwards average more Bursts20 than defencemen, with a statistically significant difference. We also observe overlapping confidence intervals among

Rank	Name	Team	Pos.	GP⋆	Bursts20 20+ MPH (29.3+ ft/s)	Top Speed (ft/s)	Avg. Top Speed (ft/s)
1	Nathan MacKinnon	COL	C	79	3.35 [2.87, 3.83]	35.1	31.7
2	Denis Gurianov	NSH	RW	11	2.53 [1.59, 3.48]	32.6	30.3
3	Julien Gauthier	NYI	RW	13	2.53 [1.59, 3.47]	34.6	30.0
4	Connor McDavid	EDM	C	68	2.18 [1.73, 2.62]	35.5	30.8
5	Noah Gregor	TOR	C	51	1.91 [1.44, 2.38]	33.5	29.8
6	Mathew Barzal	NYI	C	72	1.76 [1.32, 2.19]	34.5	30.2
7	Jack Eichel	VGK	C	58	1.71 [1.32, 2.10]	34.0	30.3
8	Martin Necas	CAR	C	68	1.61 [1.16, 2.05]	34.5	30.0
9	Andreas Athanasiou	CHI	C	24	1.58 [0.93, 2.24]	32.9	29.4
10	Ryan McLeod	EDM	C	70	1.55 [1.13, 1.97]	33.4	29.3
1	Cale Makar	COL	D	74	0.70 [0.50, 0.90]	33.1	28.6
2	Jake Sanderson	OTT	D	71	0.67 [0.44, 0.89]	33.9	28.6
3	Spencer Stastney	NSH	D	19	0.61 [0.28, 0.95]	31.8	28.4
4	Luke Hughes	NJD	D	74	0.61 [0.41, 0.81]	32.7	28.9
5	Nick Leddy	STL	D	73	0.59 [0.44, 0.74]	34.6	28.9
6	Sean Walker	PHI	D	77	0.59 [0.38, 0.79]	32.0	28.1
7	Quinn Hughes	VAN	D	69	0.58 [0.41, 0.76]	33.1	28.8
8	Colton Parayko	STL	D	73	0.58 [0.42, 0.74]	31.9	28.6
9	Jalen Chatfield	CAR	D	62	0.52 [0.31, 0.74]	32.1	27.1
10	Jamie Drysdale	ANA	D	32	0.49 [0.26, 0.72]	32.6	28.0
+	League Avg.			53	0.32 [0.29, 0.35]	31.2	26.8
+	Forwards Avg.			54	0.44 [0.39, 0.48]	31.5	27.3
+	Defensemen Avg.			53	0.14 [0.12, 0.15]	30.7	26.1

Table 1. Top 10 players ranked by average 5v5 20+ MPH Bursts per 20 minutes. GP⋆ denotes the number of games used (i.e., after applying filters).

the top 3 forwards, when comparing forwards ranked 2 to 10, and between the top 10 defencemen. This suggests that many of the top players in Table 1 are not significantly different from one another. However, we point out that the differences between all 20 players in the table and their respective position averages are statistically significant.

Notably, Nathan MacKinnon ranks 1st with 3.35 Bursts20 in comparison to the forward average of just 0.44 (7.6 times more). Cale Makar ranks first among defencemen with 0.70 Bursts20, compared to the defencemen average of 0.14 (5 times more). We note that a few players in Table 1 have had relatively low numbers of opportunities to obtain high speeds in possessions of one second or longer (GP⋆). As a result, these players typically have wider 95% confidence intervals than the rest of Table 1. This illustrates that Bursts20 may be useful for identifying players in smaller roles that have demonstrated an ability to consistently carry the puck at high speeds (although with a limited sample size).

Figure 1 plots, separately, the cumulative distribution function of Bursts20 for all forwards and all defencemen. This graph shows the clear and large difference between forwards and defencemen. Namely, it shows that nearly 19% of defencemen average zero Bursts20, compared to roughly just 7% of forwards, further illustrating that many defencemen may not be expected to carry the puck at high speeds. It also shows a large disparity between top forwards, like MacKinnon, with very high Bursts20 and other middle-ranked forwards. Players with an average of 1.0 or more bursts per 20 minutes represent fewer than 10% of all forwards and about one half of the forwards average fewer than about 0.3 bursts per 20 minutes.

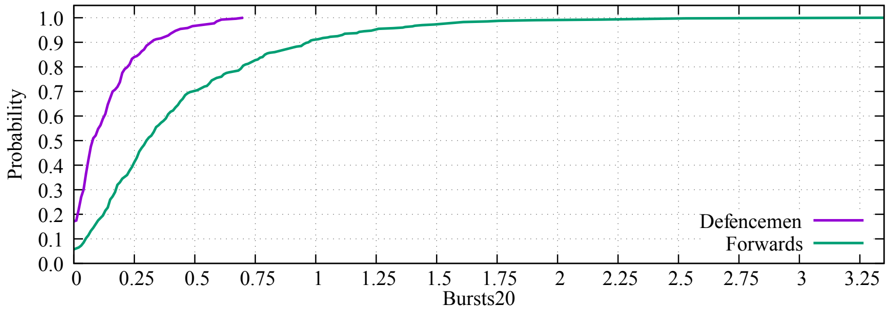

Fig. 1. CDF of Bursts20 for defencemen and forwards.

In future work, it would be interesting to consider possessions of shorter durations (e.g., half a second rather than one second) and examining the sensitivity of the results to that choice. It would also be interesting to consider bursts relative to each player's top speed. For example, studying bursts that are within p percent of a player's top speeds throughout the season. This could be useful in understanding a player's bursts relative their capability and how a player's speed changes over time. Such possibilities might include examining differences as a player ages, as their fitness level changes, or while they recover from an injury.

6 Individual Contributions to Team Possessions

In this section, we study the distribution of individual possession times across players on each team. The goal is to understand whether a team's possession time is concentrated among a few players or more evenly distributed across all players. While previous sections used individual possession data to gain insights into player roles and styles, this section focuses on how those possessions collectively shape each team's overall possession profile. However, we find a strong correlation between a player's possession time and their time on ice (TOI) ($r = 0.73$), thus our findings may also reflect underlying patterns in TOI distribution. While fairness has been used to study talent distribution in the NHL (in the context of strong-link and weak-link team structures [6]), to our knowledge, it has not previously been publicly used to study puck possession or TOI.

To examine how evenly teams share puck possession across their lineup, we compute an "equity score" based on ranked possession contributions. For each game, players on a team are sorted by their total possession duration (for 5v5 situations) and assigned a rank from 1 to 18. In this analysis, each player's possession duration is taken as-is, without normalizing for ice time. The filters described in Section 4.2 are not applied, so all players who appear in a game are included. This allows us to capture the full distribution of possession time across the entire lineup for each game.

For each team, we aggregate the possession durations across all games by rank. We sum the total time held by the top-ranked player across all games, then repeat this for ranks two through eighteen. This rank-based approach avoids bias from injuries or roster changes over the season. Each rank's total is divided by the team's overall possession time to obtain a share vector. This vector describes the proportion of total possession held by each rank from 1 to 18. We then compute Jain's Fairness Index on this vector to determine the team's equity score [7]. The equity scores range from 0 to 1 with higher values indicating a more even distribution. The Equity score (Jain's Fairness Index) is defined as:

$$E(t) = \frac{\left(\sum_{i=1}^{n} x_i\right)^2}{n \sum_{i=1}^{n} x_i^2} \quad (1)$$

Where t is the team, $E(t)$ is its equity score, x_i is the proportion of possession time held by rank i, and n is the number of ranks (18).

Table 2 ranks teams by their equity score with 95% confidence intervals computed by bootstrapping (resampling each team's games with replacement). Despite differences at the extremes, many teams have overlapping 95% confidence intervals, suggesting that possession distribution is similar across many teams.

Rank	Team	Equity Score	Rank	Team	Equity Score
1	FLA	0.85 [0.84, 0.86]	17	MIN	0.80 [0.79, 0.82]
2	VGK	0.85 [0.84, 0.86]	18	TBL	0.80 [0.79, 0.81]
3	DAL	0.84 [0.83, 0.85]	19	STL	0.80 [0.78, 0.81]
4	NSH	0.83 [0.82, 0.85]	20	BUF	0.79 [0.77, 0.80]
5	DET	0.83 [0.82, 0.84]	21	CBJ	0.79 [0.78, 0.80]
6	SEA	0.83 [0.82, 0.84]	22	WSH	0.79 [0.78, 0.81]
7	CAR	0.83 [0.82, 0.84]	23	EDM	0.78 [0.77, 0.79]
8	ARI	0.83 [0.82, 0.84]	24	NYR	0.78 [0.77, 0.79]
9	PHI	0.83 [0.82, 0.84]	25	NJD	0.78 [0.77, 0.80]
10	LAK	0.82 [0.80, 0.83]	26	ANA	0.77 [0.75, 0.78]
11	WPG	0.82 [0.81, 0.83]	27	MTL	0.77 [0.76, 0.79]
12	TOR	0.82 [0.80, 0.83]	28	NYI	0.76 [0.74, 0.77]
13	SJS	0.82 [0.81, 0.83]	29	PIT	0.75 [0.73, 0.76]
14	CGY	0.81 [0.80, 0.82]	30	OTT	0.73 [0.71, 0.74]
15	BOS	0.81 [0.80, 0.82]	31	COL	0.72 [0.70, 0.73]
16	CHI	0.81 [0.80, 0.82]	32	VAN	0.70 [0.68, 0.72]

Table 2. Equity score (Jain's Fairness Index) in 5v5 situations for all teams in the NHL.

Figure 2 plots each team's equity score against their average 5v5 goal differential. We use goal differential as the primary measure of team success because it is adaptable across game situations (e.g., 5v5). Interestingly, the results show that both balanced and unbalanced possession strategies can lead to strong team performance. The Florida Panthers (FLA) rank first in equity score, while the Vancouver Canucks (VAN) rank last, yet both are among the top four teams in average goal differential. This lack of relationship is reflected in the near-zero correlation between equity scores and average goal differential ($r = 0.02$).

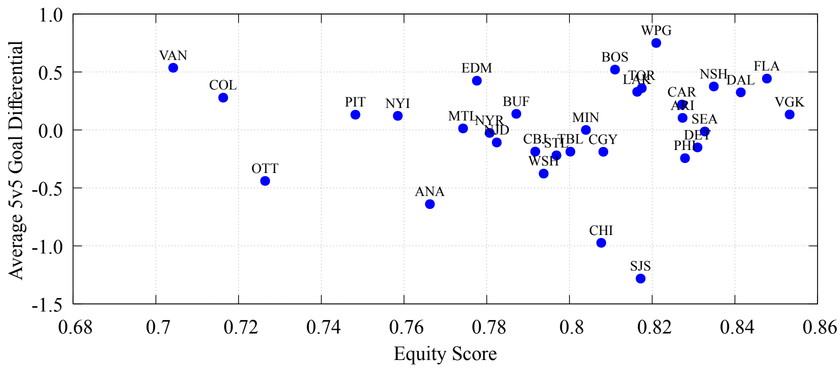

Fig. 2. Equity score versus average 5v5 goal differential ($r = 0.02$).

To illustrate how possession is distributed differently among successful teams, Figure 3 compares the Florida Panthers (FLA, 1st in fairness), Boston Bruins (BOS, 15th), and Vancouver Canucks (VAN, 32nd), who all rank in the top four in average goal differential but differ significantly in equity score. The figure shows that, on average, a smaller number of players account for a larger share of possession time on Vancouver compared to Florida and Boston (this can be seen by the steeper rise in Vancouver's curve over the first few players). This is primarily due to the top individual player on Vancouver averaging approximately 18% of the team's possession, while the top individuals for Florida and Boston each accounted for about 11%.

While fairness is computed per game and the specific top ranked player may vary, Vancouver's curve reflects a pattern of consistently high concentration at the top rank. In 64 of Vancouver's 69 games included in our dataset (92.8%), Quinn Hughes led the team in possession time. He had one minute and thirty-seven seconds of possession time per 20 minutes in 5v5 situations (the top ranked player in the league in that category [5]). The remaining five games were led by Filip Hronek (4) and Tyler Myers (1). Notably, no players from Boston or Florida rank among the top 15 in that category. Florida's top player, Mike Reilly, had 1 minute and 7 seconds of possession per 20 minutes, while Boston's top player, David Pastrnak, had 56 seconds of possession per 20 minutes.

After the top player, the rate of possession accumulation across subsequent ranks is comparable across all three teams, and in fact, the jump from the first to second player is slightly smaller on Vancouver than on Florida. This confirms that Vancouver's lower fairness score is mainly driven by Quinn Hughes high possession time in most games.

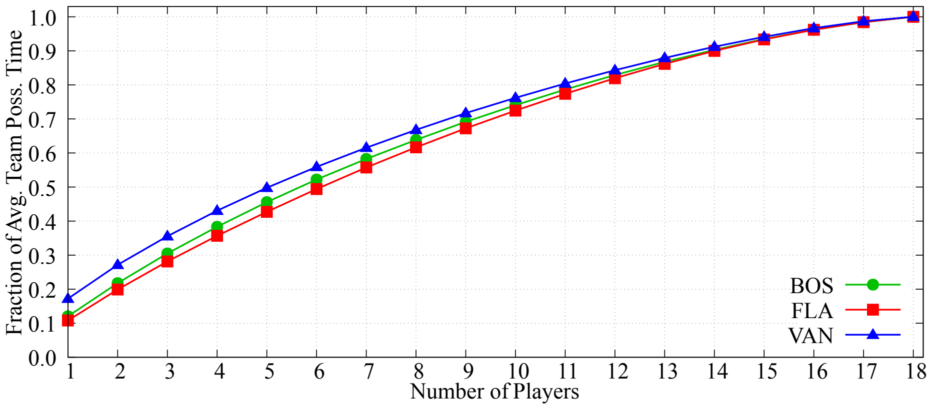

Fig. 3. Cumulative share of team possession held by players ranked 1-18 in 5v5 situations, aggregated across all games. Florida (1st in equity), Boston (15th), and Vancouver (32nd) all rank top-4 in goal differential but significantly differ in how evenly possession is distributed among team members.

7 Conclusions

In this paper we utilize unofficial NHL puck and player tracking data to introduce and analyze metrics related to player speed while in possession of the puck. We determine the number of times per game a player carries the puck for one second or more at a speed of 20+ MPH, normalize that value to 20 minutes of ice time and compute each player's per game average. We call this metric Bursts20 and find that top ranked players significantly outperform their position group averages. We believe this metric can be useful for studying and identifying players with different skills, playing styles, or roles, and that they may be useful for constructing line combinations and rosters.

We also devise a method for analyzing possession distributions within a team using Jain's Fairness Index to compute an "Equity Score". This measures how equally puck possessions are spread among players on the same team in each game. We believe that this metric provides information about team structures and playing styles and that it offers value in team analysis and scouting. We find no evidence that equitable distribution of possessions within a team influences average goal differentials.

An interesting direction for future work would be to investigate which players create or begin new possessions for their team. Additionally, we plan to examine the outcomes of individual possessions. For example, possessions that end in a pass, dump-in, shot on net, or whistle, how the outcomes vary across players, as well as the success rate of a player's possessions. Finally, we hope to examine relationships between Bursts20 and other metrics. Some examples include: zone entries, zone exits, drawn penalties, expected goals (since goals may somewhat depend on luck) and other possession outcomes. For some of these metrics it requires access to individual game data from alternative sources (i.e., data that is not available in the PPT data or via the NHL API). This is needed to ensure that only the same set of games used to compute Bursts20 are included (due to the cleaning and filtering process).

Acknowledgements

We thank Ben Resnick from the National Hockey League's Research and Development Team and Jonah Eisen from Rogers Communications for fruitful discussions related to this work. We also thank the anonymous reviewers for their helpful feedback. We thank Rogers Communications and the Natural Sciences and Engineering Research Council of Canada (NSERC) for partial funding for this project, as well as the Cheriton School of Computer Science at the University of Waterloo, for Undergraduate Research Fellowship support. We thank Neel Dayal from Rogers Communications and the National Hockey League's Information Technology, and Stats and Information Teams for making this research possible.

References

1. Casal, C. A., Maneiro, R., Ardá, T., Marí, F. J., and Losada, J. L. Possession zone as a performance indicator in football: The game of the best teams. *Frontiers in Psychology 8* (2017), 1176.
2. Collet, C. The possession game? A comparative analysis of ball retention and team success in European and international football, 2007-2010. *Journal of Sports Sciences 31*, 2 (2013), 123–136.
3. Ferioli, D., Rampinini, E., Martin, M., Rucco, D., La Torre, A., Petway, A., and Scanlan, A. Infuence of ball possession and playing position on the physical demands encountered during professional basketball games. *Biology of Sport 37*, 3 (2020), 269–276.
4. Harry Thompson. The Numbers Game. Retrieved from https://questhockey.com/2012/05/15/the-numbers-game-usa-hockey-study-reveals-importance-of-skill-development-training/, 2002.
5. Iaboni, E., Negulescu, S., Pitassi, M., Lodhi, F., and Brecht, T. Individual puck possessions Part I: Frequency, duration, and distance travelled. In *Proceedings of the Linköping Hockey Analytics Conference, Research Track (LINHAC)* (June 2025).
6. Iyer, P. Strong and weak links: Talent distribution within teams. Hockey-Graphs, 2017.
7. Jain, R., Chiu, D., and Hawe, W. A Quantitative Measure Of Fairness And Discrimination For Resource Allocation In Shared Computer Systems. 5.
8. Link, D., and Hoernig, M. Individual ball possession in soccer. *PLOS ONE 12*, 7 (2017), 1–15.
9. Liu, H., Gomez, M. Á., Lago-Peñas, C., and Sampaio, J. Match statistics related to winning in the group stage of 2014 brazil fifa world cup. *Journal of Sports Sciences 33*, 12 (2015), 1205–1213.
10. NHL EDGE. https://edge.nhl.com, 2025. Accessed March 2025.

11. PITASSI, M., BRECHT, T., AND XIE, M. Puck Possessions and Team Success in the NHL. In *Proceedings of the Linköping Hockey Analysis Conference* (2024), pp. 51–66.
12. PITASSI, M., BRECHT, T., AND XIE, M. Puck possessions and team success in the NHL. In *Proceedings of the Linköping Hockey Analytics Conference, Research Track (LINHAC)* (2024).
13. RESNICK, B. Personal communication. National Hockey League, Research and Development Team, 2024.
14. SAMPAIO, J., MCGARRY, T., CALLEJA-GONZÁLEZ, J., JIMÉNEZ SÁIZ, S., SCHELLING I DEL ALCÁZAR, X., AND BALCIUNAS, M. Exploring game performance in the national basketball association using player tracking data. *PLOS ONE 10*, 7 (2015), 1–14.
15. ZHANG, S., LORENZO, A., GÓMEZ, M.-A., LIU, H., GONÇALVES, B., AND SAMPAIO, J. Players' technical and physical performance profiles and game-to-game variation in nba. *International Journal of Performance Analysis in Sport 17*, 4, 466–483.
16. ÁNGEL GÓMEZ, M., TSAMOURTZIS, E., AND LORENZO, A. Defensive systems in basketball ball possessions. *International Journal of Performance Analysis in Sport 6* (2017), 98–107.

Dropping the Gloves, Driving the Play? Reassessing the Role of Fighting in Modern NHL Games

Rodney J. Paul, Nick Riccardi, Hunter Kuchenbaur

Syracuse University, Syracuse, NY 13244 USA

Abstract. This paper investigates the evolving role of fighting in NHL hockey by analyzing over one million play-by-play events from the 2021–22 to 2023–24 seasons. Using Corsi as a proxy for offensive activity, we find that fights are associated with short-term increases in game intensity—particularly for trailing teams. A logistic regression model further shows that fights are more likely in games with more hits and when score differentials are large. These findings suggest that fighting continues to shape game momentum and fan experience in subtle and measurable ways.

Keywords: Fighting, Possession, Corsi

1 Introduction

Fighting in hockey has long been a polarizing and enduring topic within the sport. Passionate arguments exist on both sides of the debate. Opponents of fighting often cite the significant health risks—both immediate and long-term—that players face, while also questioning whether fighting influences the outcome or momentum of a game. On the other hand, supporters argue that fighting serves a protective role, particularly for star players who may be targeted by cheap shots and point out that fans appear to respond positively to fights during games.

Historically, research supported the idea that fans enjoy fighting, with several studies showing a positive relationship between the number of fights and game attendance in both the NHL [1]–[4] and minor leagues [5]–[7]. However, more recent findings challenge this perspective. Fortney [8], for example, reported a negative and significant relationship between fighting and NHL attendance, suggesting a potential shift in fan preferences. Yet, public reaction to the highly publicized fights during the 2025 4 Nations Hockey Tournament indicates that fan appreciation for fighting may still be alive and well.

One complicating factor in assessing the relationship between fighting and attendance in the NHL is the use of dynamic ticket pricing, where prices fluctuate based on demand. This can obscure the effect of game-specific factors—such as fighting—on attendance, as the pricing models may already account for these elements. In contrast, minor league hockey, where ticket prices remain static, continues to show a positive relationship between fighting and attendance [5]-[7].

Given the overall decline in fighting in the NHL, the shift toward more skilled players (over a more physical "enforcer" role, and the growing uncertainty about whether fans are still influenced by fights when deciding to attend games, it is worth reevaluating the role of fighting in modern hockey. Several recent studies, such as those by Goldschmied [9], Leard [10], and Coates et al. [11], have failed to find any significant link between fighting and positive team outcomes, including winning games or scoring the next goal. These findings suggest that fighting may contribute little to competitive success in today's game. However, previous research has not examined the effect of fighting on in-game possession metrics. It is possible that fights may indirectly contribute to increased game activity—measured via possession statistics—even if they do not lead directly to goals. With access to detailed play-by-play and possession data, it becomes feasible to test whether fights result in tangible changes in game flow that fans might find exciting.

Using data from the 2021–22 through 2023–24 NHL seasons, this paper investigates the short-term effects of fighting on offensive activity using Corsi (shot attempts) as a proxy for possession. Our findings show that fighting is followed by an increase (decrease) in offensive activity, particularly for trailing (winning) teams by as much as 16% (-30%). This suggests that fights may act as catalysts for more dynamic gameplay with more shot attempts —an aspect that could help explain continued fan interest.

In addition, we construct a predictive model of fight occurrences based on in-game factors. Our model reveals that fights are more likely to occur when there is a two-goal or greater score differential (by a factor of 1.6), and they tend to happen earlier in games. The likelihood of fighting also increases with the number of hits and penalties in a game.

The remainder of the paper is organized as follows: Section II presents a review of the relevant literature. Section III analyzes changes in Corsi and other variables following fights. Section IV outlines our predictive model of fight occurrence. Section V discusses the implications of our findings, and Section VI concludes the paper.

2 Related Literature

This literature overview synthesizes existing academic work on fighting in the National Hockey League (NHL), with a focus on its ethical implications, impact on attendance, strategic utility, and relevance within the modern context of sport analytics. A particular emphasis is placed on empirical findings and methodological approaches that inform current debates surrounding the role and value of fighting in professional hockey.

Ethical critiques of fighting are central to the discourse. Lewinson [12] evaluates fighting through a universal code of athlete conduct derived from the NHL, the International Olympic Committee (IOC), and the Canadian National Minor

Hockey Association (NMHA). He argues that fighting largely contradicts core sporting virtues such as discipline and integrity, even while acknowledging that some defend it under the virtues of courage and loyalty. Drawing on both utilitarian and deontological frameworks, Lewinson concludes that fighting ought to be banned in order to reduce harm and promote a morally consistent code of conduct for athletes.

This ethical framing intersects with ongoing questions about the appeal of fighting from a spectator standpoint. Historically, fighting was believed to drive fan engagement, with several early studies indicating a positive relationship between fighting frequency and game attendance [1]–[4]. However, more recent evidence challenges this assumption. Fortney [8], using data from 2000 to 2020, finds a significant negative correlation between fights per game and average attendance. His results suggest that fans may now prefer high-scoring games over violent ones, signaling a shift in fan preferences that mirrors the league's own emphasis on speed and skill.

Attendance research has also historically considered the role of outcome uncertainty. Rottenberg [13] first proposed the Uncertainty of Outcome Hypothesis (UOH), suggesting that fans are more likely to attend games between evenly matched teams. However, Coates and Humphreys [14] critique the UOH in the NHL context, proposing a behavioral model centered on reference-dependent preferences and loss aversion. Paul et al. [7] found no significant support for outcome uncertainty influencing attendance in junior hockey leagues, and similar inconclusive results have been reported across European leagues, including those in Finland, Sweden, and Russia [15]. These mixed findings indicate that factors such as competitive balance and entertainment value—of which fighting is a debated component—may interact more dynamically with attendance than previously assumed.

The relationship between fighting and attendance appears to differ across league contexts. In Europe, where fighting is strictly penalized or banned, penalty minutes have a limited or inconsistent impact on spectators' interest. For instance, in Germany's DEL, penalty minutes were positively associated with attendance, while Finland's SM-Liiga showed no such effect [15]. In Canada, fighting did not significantly impact attendance in the Quebec Major Junior Hockey League, though it did in the broader Canadian Hockey League. In North American minor leagues—such as the American Hockey League (AHL), ECHL, and Southern Professional Hockey League (SPHL)—fighting continues to be positively associated with attendance [5]–[7], suggesting that its draw may be more pronounced in smaller markets or lower-tier professional contexts.

While fighting is often assumed to energize teams or sway game momentum, empirical evidence undermines this belief. Goldschmied [9] and Leard [10] both find no significant correlation between winning a fight and winning the game or scoring the next goal. Coates et al. [11] further demonstrate a negative relationship between fighting and team success, adding strategic doubt to its on-ice utility.

Research by Sirianni [16] supports this by illustrating how the role of the "enforcer" has evolved into a niche function, where players who fight do so in highly structured, often premeditated scenarios—typically against one another in controlled contexts. From a behavioral standpoint, Goldschmied [17] analyzes fight timing and concludes that players are significantly less likely to fight late in games or during the postseason, suggesting that the decision to engage is calculated rather than impulsive. Part of this is due to the instigator rule and/or the possibility of demotion if the decision to fight hurts the team.

This calculated nature, however, does not translate to tangible momentum. Studies by Steegar [18] using entropy analysis, and Kniffin [19] in collegiate hockey series, find little evidence for momentum between or within games, even in situations where teams achieve blowout victories or short-term winning streaks. Vesper [20] adds that perceived "hot hands" are not statistically supported in hockey and may, in fact, lead to decreased shot selectivity and efficiency.

The cultural normalization of violence in hockey has also drawn concern from injury prevention researchers. Cusimano [21], through qualitative interviews with youth players, finds that aggressive behavior is socially reinforced by parents, coaches, and teammates, particularly as a demonstration of loyalty or retaliation. Hutchinson [22] connects this culture of contact to concussion rates, reporting that 88% of diagnosed concussions in NHL games involved direct player contact, often occurring along the boards and early in games. Still, Goldschmied [23] reports no significant association between frequent fighting and reduced life expectancy among players from 1957 to 1971, suggesting that the most serious health effects may be short-term or not easily measurable via mortality.

Referee behavior further complicates the picture. Schuckers [24] finds that referees are less likely to call penalties in close or late-game situations, and that visiting teams are penalized more frequently than home teams. Guerette [25] expands on this by studying games without fans during the COVID-19 pandemic, showing that the typical home-ice advantage in penalty calls disappeared in empty arenas, indicating the influence of crowd pressure on officiating.

From a methodological standpoint, these studies draw on a wide range of tools. Researchers have used logistic regression [24], survival analysis [23], entropy modeling [18], exponential graph networks [16], and time series forecasting [26, 27] to explore fighting's place in the game. Metrics such as Fenwick% and xG are increasingly applied to study game flow and momentum, though their predictive power on short-term outcomes remains limited. The broader takeaway is that while fighting may be calculated and deeply entrenched in hockey's cultural history, it has little effect on outcomes, waning influence on attendance, and is increasingly at odds with the ethical and safety priorities of modern sport.

Recent scholarship has continued to refine our understanding of fighting's strategic role and broader effects. Farrington [28] presents a paradox in NHL dynamics, showing that increased fighting correlates negatively with team success, suggesting that rather than serving as a motivator, frequent fighting may hinder performance. Rockerbie [29] extends previous attendance models and finds that fighting has a small but statistically significant negative impact on NHL attendance, casting further doubt on the assumption that violence is a profitable fan draw. Meanwhile, Goldschmied and Espindola [30] explore whether hockey fights are driven by impulse or strategy. Their analysis reveals that fights occur significantly less often late in games or during the playoffs—supporting the notion that these confrontations are calculated decisions rather than spontaneous acts, with timing influenced by potential penalties and team tactics. Pitassi, Brecht, and Xie [31] contribute further by showing that a novel possession-based metric—Average Offensive Zone Possession Time Differential—strongly correlates with goal differential, outperforming traditional shot-based statistics. Despite the volume and diversity of existing research, several gaps remain. For instance, a few studies have incorporated real-time player tracking or high-resolution event data to evaluate the immediate tactical implications of fights. Moreover, while fan sentiment is often implied through attendance data, qualitative or survey-based studies on contemporary fan attitudes toward fighting are sparse. As the NHL and other leagues move toward data-driven player evaluation and league governance, there remains substantial room for new research that integrates ethics, fan behavior, and advanced analytics to better understand fighting's evolving role in the sport.

3 Data and Methodology

We scraped play-by-play data for every game played over three recent NHL seasons (2021-22, 2022-23, 2023-24) from `api-web.nhle.com`. Play-by-play data contain timestamped events throughout a game with additional game details and event descriptors. Each event contains details such as the score of the game at the time of the event, the number of skaters on the ice and whether the goalie is on the ice for both teams, and x- and y-coordinates for where the event took place (if applicable). The different events that get recorded throughout a game include starts and ends of periods, ends of shootouts, faceoffs, hits, stoppages, takeaways, giveaways, penalties, delayed penalties, shots, failed shot attempts, and goals. Shots are broken up into three categories: blocked, missed, and on-goal. Additionally, there are details given as to what type of penalty is committed. Across the three seasons of data and all games, there are 1,324,038 total events.

The Effect of Fighting on Offensive Production

We first conduct an exploratory analysis of the impact that fighting has on team offensive production post-fight. Since goals are infrequent events in hockey ($\sim$2% of all recorded events), measuring offensive production in terms of goals scored paints an incomplete picture. Instead, goal scoring opportunities, measured by shot attempts, can be a better proxy for how well a team is performing. We use Corsi, which sums all shot attempts taken by a team, to measure offensive production for each team in a game.

To analyze the impact of fighting on offensive production, we create post-fight windows of time and compare offensive production within these windows to offensive production from the start of the game to the time of the fight.

Table 1. Summary of Fights Per Game Across Seasons.

Games with:	2021-2022	2022-2023	2023-2024
0 Fights	1,037 (79%)	1,037 (79%)	1,047 (80%)
1 Fight	228 (17%)	233 (18%)	227 (17%)
2+ Fights	47 (4%)	42 (3%)	38 (3%)

The post-fight windows include: two minutes post-fight, five minutes post-fight, 10-minutes post-fight, and the duration of time from the fight occurring and the end of the period. For each occurrence of a fight, we compare the offensive production of both teams during the windows after the event to their offensive production before the event occurred. The pre-fight window that is compared to all post-fight windows encompasses all events from the start of the game to the time of the fight.

We restrict the sample to regular season games and regulation periods as overtime periods are played with fewer players on the ice. Initially, our sample includes 972 fights. To allow for enough time to pass, we only analyze fights that occur at least two minutes into the game and at least two minutes before the end of the game, reducing the sample to 877 fights. Additionally, we only analyze fight occurrences which resume play at even strength immediately post-fight to avoid entangling results with the impact of a team having a power play, which further reduces the sample size to 630 (∼65% of all fights).

In the before and after windows, we calculate Corsi rate by summing shot attempts during the window and dividing by the duration of the window in seconds. If there was not enough time after the event occurring until the end of the game to cover the calculated duration range (e.g., a fight happening with three minutes left in the game would not have a complete five-minute post-fight window), we divide the Corsi sum by the actual time elapsed. Equation one displays Corsi rate where t is the duration of the window in seconds (e.g., 120 seconds for the two-minute window for a fight that occurs before the last two minutes of a game):

$$CorsiRate = \frac{(Goals + ShotsonGoal + MissedShots + BlockedShots)}{t} \tag{1}$$

Table 2 provides an array of paired Welch's t-test results comparing the post-fight Corsi rates to the pre-fight rates for fights. We do not assume equal variance across the compared post-fight and pre-fight windows, hence the choice of Welch's t-tests. Since each fight has post-fight windows that correspond to a pre-fight window, we use paired tests. The rates were multiplied by 60 before conducting the tests for interpretability, and the rates represent Corsi per minute. Eight t-tests were specified for fights that occur during each of the following game score scenarios: all scenarios, tie games, home team losing by one goal, home team winning by one goal, home team losing by two or more goals, and home team winning by two or more goals. The eight t-tests include four for both the home and away teams for their rates in the four post-fight windows. For each test, the mean difference in Corsi rate is reported with t-statistics in parentheses.

Table 2. Summary of Fights Per Game Across Seasons.
The *-notation notes statistical significance of t-test at 1% (***), 5% (**) and 10% (*) levels.

		N	After 2	After 5	After 10	Until EOP
All scores		630				
	Home		-0.010	0.019	0.020	0.017
			(-0.296)	(0.869)	(1.115)	(0.789)
	Away		-0.027	-0.031	-0.029	-0.017
			(-0.892)	(-1.390)	(-1.611)	(-0.753)
Tie game		197				
	Home		0.033	0.058	0.054	0.042
			(0.514)	(1.324)	(1.470)	(1.056)
	Away		0.035	-0.016	-0.028	0.009
			(0.660)	(-0.416)	(-0.826)	(0.225)
Home losing by 1		101				
	Home		0.117	0.158	0.099	0.108
			(1.459)	(2.722)***	(2.251)**	(1.899)*
	Away		-0.180	-0.192	-0.088	-0.073
			(-2.380)**	(-3.804)***	(-2.269)**	(-1.547)
Home winning by 1		111				
	Home		-0.079	-0.041	0.000	0.019
			(-1.007)	(-0.798)	(0.001)	(0.322)
	Away		0.146	0.136	0.129	0.091
			(2.069)**	(2.952)***	(3.511)***	(1.916)*
Home losing $\geq$ 2		99				
	Home		-0.001	0.042	0.079	0.036
			(-0.014)	(0.842)	(1.949)*	(0.795)
	Away		-0.283	-0.234	-0.208	-0.260
			(-4.440)***	(-4.210)***	(-4.600)***	(-5.455)***
Home winning $\geq$ 2		122				
	Home		-0.127	-0.121	-0.129	-0.115
			(-1.803)*	(-2.782)***	(-3.561)***	(-2.357)**
	Away		0.052	0.093	0.022	0.090
			(0.750)	(1.748)*	(0.556)	(1.463)

Table 2 provides an array of paired Welch's t-test results comparing the post-fight Corsi rates to the pre-fight rates for fights. We do not assume equal variance across the compared post-fight and pre-fight windows, hence the choice of Welch's t-tests. Since each fight has post-fight windows that correspond to a pre-fight window, we use paired tests. The rates were multiplied by 60 before conducting the tests for interpretability, and the rates represent Corsi per minute. Eight t-tests were specified for fights that occur during each of the following game score scenarios: all scenarios, tie games, home team losing by one goal, home team winning by one goal, home team losing by two or more goals, and home team winning by two or more goals. The eight t-tests include four for both the home and away teams for their rates in the four post-fight windows. For each test, the mean difference in Corsi rate is reported with t-statistics in parentheses.

Results of the t-tests suggest that teams might benefit from fighting. In scenarios where the game is not tied, there appears to be evidence the losing team at the time of the fight benefits from either or both an increase in offensive production themselves and a decrease in offensive production for the winning team. These results might suggest that it can be strategic to fight in certain game scenarios. While the reported mean differences might appear miniscule, mean Corsi rates per minute in the dataset are roughly 0.98 and 0.95 for home and away teams, respectively. Therefore, a result such as the 0.158 increase for the home team and -0.192 decrease for the away team in the five minutes post-fight window when the home team is losing by one goal is rather substantial. For this example, approximate percentage changes of Corsi rate in the post-fight window are +16% for the home team and -20% for the away team. Coupling these results suggest a major post-fight advantage for the home team.

4 Fighting Probability Model

We model the occurrence of a fight in an NHL game based on game characteristics. Namely, whether the score differential, time remaining in the game and period number, Corsi differential, and hit and penalty running totals impact the probability of a fight breaking out. We specify logistic regression models with a dependent variable of a game event being a fight ($y = 1$). Regressors include score differential (Home – Away), time remaining, period number, Corsi differential (Home – Away), Hit count, and Penalty count. Corsi differential is measured using the cumulative sum of Corsi for each team at the time of the event while hit and penalty counts are the cumulative sum of these events across both teams.

Without any transformations, the functional form of the models shows issues with heteroskedasticity and autocorrelation. We correct these issues in two ways. First, in Model I, we present results using clustered standard errors, clustered by the individual game. In Model II, results are presented using heteroskedasticity-consistent standard errors. We attempted using heteroskedasticity- and auto-correlation- consistent standard errors, but alas, the methods utilized in R were computationally costly and did not converge due to the sample size and complexity of the regressors.

Models I and II are presented in Table 3. The robust standard errors calculated for Model II are consistently smaller compared to those for Model I, increasing the significance of the explanatory variables across the board. The lack of correction for autocorrelation in Model II is likely leading to a misspecification of the model. However, the only difference between the two models is $ScoreDiff_1$ is significant in Model II but not in Model I. The results of Model I suggest that fights are more likely to occur when the score differential in a game is two goals or greater in either direction. The time remaining and period number variables suggest that fights are more likely to occur earlier in the game. Lastly, the number of hits and penalties increases the likelihood of a fight occurring, which is to be expected, as these variables control for overall aggression.

Table 3. Results of Models I and II.
The *-notation notes statistical significance of t-test at 1% (***), 5% (**) and 10% (*) levels.

Variable	I	Odds Ratio	II	Odds Ratio
Intercept	-8.009***		-8.009***	
	(0.361)		(0.273)	
$ScoreDiff_{-1}$	0.022		0.022	
	(0.112)		(0.076)	
$ScoreDiff_{1}$	0.117		0.117*	1.124
	(0.101)		(0.070)	
$ScoreDiff_{-2}$	0.404***	1.498	0.404***	1.498
	(0.108)		(0.077)	
$ScoreDiff_{2}$	0.470***	1.600	0.470**	1.600
	(0.101)		(0.070)	
Time Remaining	0.001***	1.001	0.001***	1.001
	(0.0001)		(0.0001)	
Corsi Differential	0.030		0.030	
	(0.035)		(0.021)	
Second Period	-0.955***	0.385	-0.955***	0.385
	(0.118)		(0.077)	
Third Period	-2.427***	0.088	-2.427***	0.088
	(0.201)		(0.126)	
Hit Count	0.015***	1.015	0.015***	1.015
	(0.005)		(0.003)	
Penalty Count	0.220***	1.246	0.220***	1.246
	(0.015)		(0.005)	
Home Team Fixed Effects	Yes		Yes	
Away Team Fixed Effects	Yes		Yes	

Table 4 presents the results of models using score differential as a continuous variable and its squared term. Again, Model III uses clustered standard errors and Model IV uses heteroskedasticity-consistent standard errors. Operationalizing score differential as a continuous variable with its squared term in Model III-IV provides the same suggestions as Models I-II: fights are more likely to occur as the score differential grows, regardless of whether the home team is winning or losing.

Table 4. Results of Models III and IV.
The *-notation notes statistical significance of t-test at 1% (***), 5% (**) and 10% (*) levels.

Variable	I	Odds Ratio	II	Odds Ratio
Intercept	-7.979***		-7.979***	
	(0.362)		(0.274)	
Score Differential	0.006		0.006	
	(0.020)		(0.013)	
Score Differential^2	0.024***	1.024	0.024**	1.024
	(0.005)		(0.003)	
Time Remaining	0.001***	1.001	0.001***	1.001
	(0.0001)		(0.0001)	
Corsi Differential	0.040		0.040*	1.041
	(0.036)		(0.021)	
Second Period	-0.868***	0.420	-0.868***	0.420
	(0.115)		(0.075)	
Third Period	-2.343***	0.096	-2.343***	0.096
	(0.197)		(0.125)	
Hit Count	0.014***	1.014	0.014***	1.014
	(0.005)		(0.003)	
Penalty Count	0.215***	1.240	0.215***	1.240
	(0.015)		(0.005)	
Home Team Fixed Effects	Yes		Yes	
Away Team Fixed Effects	Yes		Yes	

5 Discussion

The findings of this study contribute new dimensions to our understanding of fighting in hockey, particularly in terms of its on-ice effects and contextual likelihood. While earlier research has questioned the strategic value of fighting—often concluding that it does not lead to improved outcomes such as winning a game or scoring the next goal—our analysis suggests that fights can serve as a catalyst for increased offensive activity, at least in the short term. Specifically, fights appear to boost Corsi rates (i.e., shot attempts), particularly for the team behind on the scoreboard. These increases in offensive zone activity may not always translate into goals, but they do indicate a measurable shift in game tempo that could enhance the excitement and momentum perceived by players and fans alike.

This potential for fights to energize gameplay may partially explain the lingering fan interest in fighting, even as its frequency declines and its role as a performance tool diminishes. The results align with the hypothesis that fighting can be situationally beneficial—less as a deterministic event, and more as a psychological or momentum-shifting mechanism, particularly for teams attempting to disrupt an opponent's control or revive their own effort.

This insight complicates the narrative that fighting is purely detrimental or antiquated, suggesting that its role is more nuanced and possibly adaptive to specific game states. The predictive modeling further enriches this perspective by identifying the conditions under which fights are more likely to occur. Score differential—particularly when it reaches two goals or more—emerges as a key driver, suggesting that fights often occur as responses to perceived imbalance rather than in tightly contested games. The timing of fights also matters: they are more likely to happen earlier in games, likely due to teams' hesitancy to incur penalties or lose players during decisive moments. In addition, higher counts of hits and penalties are strongly associated with fight occurrence, reinforcing the idea that fights emerge from escalations in physicality and game intensity.

Importantly, the relationship between fighting and Corsi metrics complicates earlier conclusions from studies such as Goldschmied [9], Leard [10], and Coates et al. [11], which focused largely on scoring and winning. Our study highlights that there may be more subtle, immediate effects on gameplay that are not captured by goals alone. This underscores the importance of incorporating advanced possession metrics and high-resolution event data when assessing the tactical or entertainment value of fighting in contemporary hockey.

From a policy standpoint, these findings walk a middle line. They neither fully vindicate fighting as an essential tool nor entirely discredit its relevance. Instead, they suggest that fighting continues to exert situational effects on game dynamics that may hold residual value for teams, players, and spectators—particularly in terms of psychological tone and energy on the ice.

6 Conclusion

This paper revisits the complex and controversial role of fighting in hockey through the lens of modern sport analytics. Drawing on play-by-play data from three NHL seasons (2021–22 to 2023–24), we examined both the in-game effects of fighting on offensive production and the contextual conditions under which fights are most likely to occur. Our analysis shows that fighting is associated with short-term increases in offensive activity, especially for trailing teams—suggesting that fights can act as momentum shifts even if they do not translate directly into scoring outcomes.

Furthermore, our predictive modeling indicates that fights are more likely when games are physically intense, involve higher penalty counts, or feature a notable score differential. These insights reinforce the idea that fighting often emerges not randomly, but as a strategic or emotional response to in-game dynamics.

Together, these results lead to a more nuanced understanding of fighting's place in modern hockey. While fighting may no longer be central to winning games or building rosters, it retains the ability to influence gameplay intensity and spectator experience. Future research should continue to explore these short-term effects using additional tracking data, fan sentiment surveys, and cross-league comparisons. As the NHL and other leagues continue evolving toward faster, more skilled styles of play, understanding how legacy elements like fighting affect the game's rhythm and perception will remain essential to informed policymaking and fan engagement.

Acknowledgments: We thank the anonymous referees for their helpful comments.

Disclosure of Interests: The authors have no competing interests.

References

1. Jones, J.C.H. (1984). Winners, losers, and hosers: Demand and survival in the National Hockey League. *Atlantic Economic Journal*, 12(3), 54–63.
2. Jones, J.C.H., Ferguson, D.G., & Stewart, K.G. (1993). Blood sports and cherry pie: Some economics of violence in the National Hockey League. *American Journal of Economics and Sociology*, 52(1), 87–101.
3. Jones, J.C.H., Stewart, K.G., & Sunderman, R. (1996). From the arena into the streets: Hockey violence, economic incentives, and public policy. *American Journal of Economics and Sociology*, 55(2), 231–249.
4. Paul, R.J. (2003). Variations in NHL attendance: The impact of violence, scoring, and regional rivalries. *American Journal of Economics and Sociology*, 62(2), 345–364.
5. Paul, R.J., Weinbach, A.P., & Robbins, D. (2013). American Hockey League attendance: A study of fan preferences for fighting, team performance, and promotions. *International Journal of Sport Finance*, 8(1), 21–38.
6. Paul, R.J., Weinbach, A.P., & Robbins, D. (2015). Fighting, winning, promotions, and attendance in the ECHL. *Sport, Business and Management*, 5(2), 139–156.
7. Paul, R.J., Weinbach, A.P., & Riccardi, N. (2019). Attendance in the Canadian Hockey League: The impact of winning, fighting, uncertainty of outcome, and weather on junior hockey attendance. *International Journal of Financial Studies*, 7(1), 1–12.
8. Fortney, T. et al. (2022). National Hockey League fights per game and viewership trends: 2000–2020. *Frontiers in Sports and Active Learning*, 1, Article 932608.
9. Goldschmied, N. (2017). Do not drop your gloves: Winning a fight in the NHL does not improve chances to win the game. *Psychology of Sport and Exercise*, 56, 101986.
10. Leard, B., & Doyle, J. (2015). The effect of home advantage, momentum, and fighting on winning in the National Hockey League. *Journal of Sports Economics*, 16(5), 531–553.
11. Coates, D., Battre, M., & Deutscher, C. (2012). Does violence in professional ice hockey pay? Cross-country evidence from three leagues. In P. Rodriguez, S. Késenne, B.R. Humphreys (Eds.), *Sports Economics: Current Research* (Vol. 4, pp. 47–63). Springer.
12. Lewinson, R.T. (2012). The morality of fighting in ice hockey: Should it be banned? *Journal of Sport and Social Issues*, 36(1), 96–113.
13. Rottenberg, S. (1956). The baseball players' labor market. *Journal of Political Economy*, 64(3), 242–258.
14. Coates, D., & Humphreys, B.R. (2012). Game attendance and competitive balance in the National Hockey League. *Journal of Sports Economics*, 13(4), 364–377.
15. Paul, R.J. et al. (2016). Attendance in European Hockey Leagues: Evidence from Finland, Sweden, and Russia. *Journal of Sports Economics*, 17(1), 25–47.
16. Sirianni, A. (2021). The specialization of informal social control: Fighting in the National Hockey League. SocArXiv. `https://doi.org/10.31235/osf.io/fjx7p`
17. Goldschmied, N. (2013). Fighting in hockey: Evidence of an enforcer role. *Sport Management Review*, 16(4), 460–465.

18. Steegar, G. et al. (2019). Winning and losing streaks in the NHL: Are teams experiencing momentum or are games a sequence of random events? *Journal of Quantitative Analysis in Sports*, 15(3), 213–227.
19. Kniffin, K.M. (2015). Are hockey winning streaks real? *CHANCE*, 28(3), 27–33.
20. Vesper, A. (2015). Putting the hot hand on ice. *CHANCE*, 28(3), 3–8.
21. Cusimano, M.D. (2013). Addressing violence in hockey: Beyond penalties and fighting. *British Journal of Sports Medicine*, 47(1), 1–2.
22. Hutchinson, M.G. (2015). Head contact in men's ice hockey: Comparing injury rates and situational factors in professional and amateur games. *Injury Epidemiology*, 2(1), 1–8.
23. Goldschmied, N. (2013). Mortality rates among NHL players and the implications of frequent fighting. *Sport Sciences for Health*, 9(2), 97–101.
24. Schuckers, M. (2011). Referee bias in the NHL: Evidence from penalty call data. *Journal of Quantitative Analysis in Sports*, 7(4), Article 2.
25. Guerette, K. (2021). No fans, no bias? NHL refereeing during the COVID-19 pandemic. *Sport in Society*, 24(10), 1609–1622.
26. Weissbock, J. (2014). Predicting the NHL using advanced statistics and machine learning. *MIT Sloan Sports Analytics Conference.*
27. Johansson, F. (2019). A time series approach to analyzing hockey game momentum. *Journal of Sports Analytics*, 5(1), 45–57.
28. Farrington, D. (1995). Fighting behavior and winning National Hockey League games: A paradox. *Perceptual and Motor Skills*, 81(3), 1230–1230.
29. Rockerbie, D.W. (2016). Fighting as a profit-maximizing strategy in the National Hockey League: More evidence. *Applied Economics*, 48(4), 323–333.
30. Goldschmied, N., & Espindola, S. (2013). I went to a fight the other night and a hockey game broke out: Is professional hockey fighting calculated or impulsive? *Sports Health*, 5(3), 232–234.
31. Pitassi, M., Brecht, T., & Xie, M. (2024). Puck possessions and team success in the NHL. *Proceedings of the Linköping Hockey Analytics Conference LINHAC 2024 Research Track*, 51-66.

Community Notes

Merging Shot Data and Measuring Player Importance to Team's Offensive Flow

Louis Boulet[0009−0000−3432−2604]

Hockey Data Analyst, LB-Hockey, Ottawa ON, Canada
lbhockey.social@gmail.com

1 Bridging the Data Gap

For the last few years, the quality of available shot data features in the public hockey analytics sphere has really fallen behind what is being used within teams and companies. Although the NHL provides various parameters for each shot attempt, passes leading to them are not registered anywhere. This prevents the use of situational expected goals to measure skills such as expected assists for one's playmaking strength, or expected points from entry shots for the ability to create off the rush.

MoneyPuck (MP) does an excellent job of cleaning and formatting what the NHL offers with precise coordinates, previous events, and their own expected goal model attached to every shot attempt. But this still leaves room for more. Now it does not mean that this type of tracking is entirely unavailable. Corey Sznajder manually tracks hundreds of games every season as part of his incredibly insightful AllThreeZones (A3Z) project. He records more situation context like offence type (cycle, forecheck, rush), screen presence, and most notably, the three most recent passes leading to a shot.

While the NHL's unique shot ID is not present in Sznajder's shot tracking sheets, the MP and A3Z datasets share several features that allow for a potential merging process. The shot time, type (backhand, slap, wrist, tip, etc.), outcome (miss, on net, goal), and shooter-goalie pair are noted in both sources. This is enough for roughly 90% of the tracked A3Z shot attempts to be matched with their equivalents on the other side of the data pond, taking under 25 seconds for a roughly 400-game sample to be generated. The result is a database that links coordinates and expected goal values from MoneyPuck to the contextual indicators and pass tracking from AllThreeZones.

2 Network Analysis Modelling in Hockey

Network analysis concerns the evaluation of relationships within a structure made of nodes that are connected through edges. Analyzing a passing network within a team felt like a natural way to implement this into my work, with the players and passes between them acting as the nodes and edges respectively. These edges will be directed, from sender to receiver, and weighted, using separation from the play's eventual shot (primary, secondary, tertiary assist) and that shot's threat level (expected goal value).

Extensive research has focused on quantifying importance within these networks, using metrics known as centrality measures. And as is the case throughout data science, there is no end-all-be-all to every solution. We will use four centrality measures to build an offensive puck-movement profile, capturing a diversity of contributions and aiming to outline any given player's inclinations with the puck and role when on the attack.

The first is importance through flow, using Betweenness Centrality. Betweenness is calculated by going over how often a player lies on the shortest path between nodes. Players ranking highly here are crucial facilitators who maximize their team's passing routes toward shots. They tend to act as the main link between any two teammates, resulting in their line's offensive movement flowing through them.

We can estimate a player's influence with PageRank Centrality: an iterative process that considers the entire team dynamics. This benefits skaters who connect highly valued players, usually acting as an intermediary between central portions of the network.

By considering all possible paths within a network, Information Centrality finds skaters who enhance the efficiency of their team's offensive system. This is where offensive support pieces tend to shine. They make themselves valuable by creating pathways that allow for flexible progression even if they are not the most direct options, rendering possessions more robust to interruptions should they arise.

Keeping the simplest for last, Weighted Degree Centrality favours high-volume players as it sums the total edge weights coming in and out of a node. Heavily involved passers stand out here even if their plays do not result in groundbreaking connections.

3 Chemistry Interactions

Using these profiles, we can conduct a quick case study of how the different areas of puck-moving interact with each other to establish chemistry at the line level. We'll be taking the Ottawa Senators' first line as an example here. All the ranks mentioned in this section are league-wide among forwards during the 2023-24 season.

First, Brady Tkachuk is the identity piece in Ottawa. He plays a very high-danger-centric style of game and so much of the Senators' offence is geared towards getting him a chance in-tight whether that is through a screened tip, rebound, bumper play, or general net scramble. As a result, Tkachuk has a lot of sway over how the team's plays unfold, giving him a top 3 Influence (PageRank) score.

Next, we have Tim Stützle, who is the primary carrier and dynamic mover. Many of the Senators' possessions are made possible by virtue of the center connecting plays thanks to his playmaking and skating abilities. With such a high share of Ottawa's offensive pathways going through him, Stützle garners a top 3 Flow (Betweenness) mark.

Lastly on the right wing is Claude Giroux. He tends to play more of a tertiary role here, acting as the reliable veteran piece for his younger linemates. But Giroux goes beyond that, with his anticipation of potential passing routes ensuring Ottawa's puck possessions can progress towards a chance no matter their current state. The adaptability and support he provides rank him in the top 10 on the Efficiency (Information) side.

Each player's distinct specialization in a different facet of play-driving and the meshing of these puck-movement abilities allows the line to consistently produce at a high level. Interactive visualizations of these offensive puck-movement profiles are available to all at LB-Hockey.com along with the full-length article.

xReboundsPlus, Creating a Statistic to Predict Rebound Quality

Evan Lawing

elawing1018@gmail.com
https://medium.com/@elawing1018

Rebounds are one of the most important aspects to any team's success when it comes to hockey. Scoring off of rebounds is becoming increasingly important and starting to play a larger role in hockey every season. In an analysis by Neil Pierre-Louis, he found that the number of rebounds per game increased dramatically from the 2017-18 season to the 2023-24 season, rising from 3.53 rebounds per game to 6.17. The increase in rebounds highlights the fact that teams are using rebounds more than ever in an effort to score goals where they catch the goalie and defense off guard. To show the value of each rebound, a new statistic, xReboundsPlus, will quantify the quality of the rebound of a given shot. This will be calculated in three steps: first, modeling the probability of a rebound (xRebounds); second, modeling the expected shooting angle of the rebound (xAngle); and finally, combining xRebounds and xAngle to determine the goal probability.

The shot data used for these models was collected from moneypuck.com; specifically, the shot data from 2023-24 were used to train these models and applied to the data from the 2024-25 season. The variables that were used for this analysis were if the shot resulted in a goal, if the shot was a rebound, if the shot generated a rebound, if the shot was on a rush, the handedness of the shooter, the shot type, the angle of the shot, and the arena adjusted shot distance. For this analysis, a rebound was defined using the Moneypuck standard: a subsequent shot occurring within three seconds of the initial shot.

The first step in creating the xReboundsPlus statistic was to create an expected rebound model that outputs the likelihood that a given shot will result in a rebound. Extreme gradient boosting (XGBoost) was selected as the model to be used for rebound classification due to its ability to classify well and create the model rather quickly. The dependent variable for the model was whether the shot generated a rebound and the independent variables were the angle of the shot, the shot distance adjusted to the arena, the type of shot, the handedness of the shooter, and if the shot came off the rush.

K-fold cross validation with five folds was used to determine the best XGBoost model and the parameters were tuned to prevent overfitting. These tuned parameters included adjusting for a class imbalance, changing the maximum depth, and adjusting the learning rate of the model. After tuning the parameters, the final xRebounds model had an area under the ROC curve of 0.9017 and an accuracy of 0.8097. Additionally, the model had a specificity of 0.8418 and a sensitivity of 0.8071, making this a usable piece of the xReboundsPlus model.

The next step is creating the xAngle model, unlike the previous model for xRebounds, this model needs to output a continuous variable: the rebound shot angle in degrees. To account for this, multiple different types of models were created and compared to determine the best possible fit. The models that were considered for this expected angle prediction were a neural network, another XGBoost model, a random forest model, multiple linear regression, and Bayesian regression. A switch to Bayesian regression was made after the other methods failed. Using Bayesian regression allowed the model to be able to make predictions based on prior events and is better for modeling uncertainty. The Bayesian regression model did converge with all Rhat values being 1 meaning there were no signs of divergence or sampling issues. The selected Bayesian regression model had a low Bayesian R-Squared value of 8%, but this was comparatively better than the other models.

Once the models were created for xRebounds and xAngle, the xReboundsPlus model was ready to be created. Before the model was created, additional data manipulation was needed to bring in the xAngle and xRebound predictions for each shot. Additionally, if the next shot was a goal and the distance of the next shot were brought over to each shot observation as well. Finally, the data was filtered to only shots that generated rebounds to have the proper sequencing for the training.

XGBoost was selected as the machine learning method with K-fold cross-validation. The independent variable were xRebounds, xAngle and the distance of the rebound shot. After tuning, the model had an area under the ROC curve of 0.9242, which means that it was able to classify well but was not at risk of overfitting as the original model. A confusion matrix was created as well achieving an accuracy of 80.99%.

After fully creating xReboundsPlus, the model was implemented on shot data for the 2024–25 NHL up until the morning of March 26, 2025. Since there was no actual distance that could be used for each individual shot, the average distance of a rebound was used for each calculation. Once the statistic was applied to each shot, rankings were created based on the sums of xReboundsPlus, xRebounds, and shots. Brady Tkachuk led the league in the statistic with 94.85 xReboundsPlus, followed by Matt Boldy and Nathan Mackinnon with 89.39 and 86.09 respectively.

This metric provides a framework for measurement of rebound opportunities, but there are still many ways to improve the model. With data that is able to track where the shot would hit the net and the positioning of the goalie, the prediction could become even more accurate as the likelihood of a rebound and the expected angle could both become more accurate as the model will more accurately depict hockey in three dimensions. Additional shot velocity or goalie specific data would improve the model as well due to differences in reaction time. Overall, the project was able to meet the original goal by creating a new rebound-based statistic, but still has room for improvement. The full code used to create this project can be found on my Github, https://github.com/elawing40.

Exploring Leadership and Cultural Training Experiences of Ice Hockey Coaches: Lessons Learned from Nova Scotia

Dr. Lynn LeVatte[1], Dr. Christina Phillips[2], Dr. Shaun Ranni[3], Dr. Kristin O'Rourke[1], Sarah Timmons[1], Sarah MacRae[1], Michelle Hoffman

[1] Education Department, Cape Breton University, Nova Scotia, Canada
[2] Department of Curriculum, Teaching & Learning University of Toronto, Ontario, Canada
[3] Sport Leadership & Active Living Department, Cape Breton University, Nova Scotia, Canada

Abstract. Within the competitive sport of ice hockey, coaches are instrumental in motivating, communicating with and supporting players to reach their full potential. To enhance coaching skills and training to meet the needs of an evolving sport, many minor and professional associations examine various professional development types necessary to respond to culture and diversity. This study examined ice hockey coaches' coaching leadership and cultural training experiences in the province of Nova Scotia, Canada. Leadership is a fundamental factor influencing the performance of sports teams. Leadership can be provided by coaches, assistant coaches or other staff on sports teams. This leadership capability is essential for hockey coaches who are tasked with providing cohesion and effective communications within the team unit, thus impacting overall performance. In many cases, this team unit may comprise players representing two or three languages, various skill abilities, and interest or motivation to play hockey. This qualitative study (n=25) included 19 male and 6 female coaches. Participants coached youth aged 5 to 18 years. The study aimed to investigate the leadership and cultural training experiences of minor hockey coaches who volunteered with Hockey Nova Scotia, a member of the national association, Hockey Canada. Participants completed semi-structured, in-person and/or online interviews, consisting of 5 open-ended questions. The findings indicated that participants did not recall receiving professional training specific to leadership and cultural development. Results demonstrated specific areas of training need, such as supports for behaviour, communication and disability may enhance performance. Additionally, results indicated that a preferred format of leadership and cultural training would be in-person sessions compared to online delivery. These interventions could potentially support overall team performance, communication and engagement. Findings were analyzed using a thematic approach, and the research team assisted in developing gaps in training to prepare and implement future professional development opportunities for coaches with Hockey Nova Scotia.

Keywords: Leadership, coaching, ice hockey, culture, professional development training, performance

1 Introduction

Youth sports participation such as hockey has been shown to correlate with many developmental assets in youth, including physical, cognitive and social competencies (Hansen et al., 2003). In addition, research within the participation in youth sports has documented such benefits as healthy self-esteem, higher rates of pro-social engagement and academic achievement, and the development of character and life skills (Horn, 2008).

2 Related work

Professional development training is necessary to maintain skillsets, enhance personal growth, continuous learning and involves expanding one's expertise within their field. However, evidence indicates that very few training programs teach staff to integrate cultural assets, beliefs or values (Gomez-Hurtado, 2021). Furthermore, managing cultural diversity is therefore both a necessary and an essential task, particularly for management and coaches who must oversee a greater number of tasks, roles, and responsibilities (Rönnström & Scott, 2019). Youth program staff have the potential to have a positive and lasting impact on youth. As such, high-quality professional development for staff is vital to youth programs' success (Rana et al., 2013). Hockey Canada explained that coaches are caring & enthusiastic, and that a well-trained coach can be a positive influence on the experience of players, parents and other coaches (Hockey Canada, 2025). Hockey coaches play a pivotal role within youth ice hockey players and can impact players and overall team performance.

Exploring the terminology of leadership and coaching several common elements exist, such as a one-on-one relationship, raising self-awareness, performance, learning and development, and behavioural change (Grant et al., 2019). Additionally, research conducted by Jones et al.(2016), revealed core features of coaching, which included providing a supporting relationship, setting personal development objectives, achievement of these objectives through focusing on inter- and intra-personal issues, and helping the player develop and be more effective by providing the player with the tools, skills and opportunities they need (de Haan et al., 2023).

3 Methods

A qualitative research design was deliberately chosen with the aim of facilitating open and engaged discussion through semi-structured interviews. This methodological structure involves qualitative methods that are accommodating and flexible, allowing for the combination and adjustment of different qualitative approaches to fit the sample, phenomenon, context, and epistemological paradigm of the researcher (Denzin & Lincoln, 2017).

Number of Interviews: 25; Youngest Participant: 20; Oldest Participant: 63,
Average Age of Participants: 40.4, Minimum Number of Coaching Years: 1,
Maximum Number of Coaching Years: 50, Average Number of Years Coaching: 11.35

4 Findings

The preliminary findings are categorized within a thematic analysis approach. Using Atlas Ti, several themes emerged from the data and are summarized to represent common ideals within ice hockey coaching leadership and cultural diversity. The themes were identified as coach training, coaching barriers to professional development for cultural diversity and leadership, mentoring opportunities, and formats for leadership training and hockey performance.

The topic of mentorship was another emerging theme from our data. The critical importance of obtaining leadership development skills, communication skills, growing personal coaching abilities and understanding the impact of coach-player relationships are vital for team performance. As participant 10 explained: "For me personally, a leadership piece maybe like a mentorship program, and that would be great... here you can have like a like a senior hockey Nova Scotia Rep come around and, you know, once a month or whatever and, go through the with the individual coaches and maybe give them some kind of like tips on what they can change".

5 Discussion

The aim of the current study was to examine the leadership and cultural training experiences of ice hockey coaches in Nova Scotia, Canada. Results revealed that some coaches utilized a mentoring process with experienced coaches to support additional informal training practices, many faced challenges or barriers regarding accessibility in receiving formalized training programs to support cultural diversity or leadership and the lack of access to training may have impacted player and team performance. Coaches noted that communication with players about personal issues could support how they played and how their confidence was impacted by the communication. These results compliment findings from research conducted by Athletic Insight Research (2025) who determined that "effective communication is crucial in team sports. It directly impacts performance, morale, and cohesion" (p.1). The current study strengthens previous literature within coaching, communication and leadership training needs.

6 Conclusion

The aim of the present research study was to explore cultural leadership training experiences of minor hockey coaches. This study was limited by gender, as 19 of 25 participants were male, thus providing disproportionate numbers of male responses, the study also was limited as it was conducted within a smaller province in Atlantic Canada. Suggestions for future study within the realm of culture and diversity training for hockey coaches and delivery of training types would be recommended. Embracing leadership opportunities to support diversity can assist minor hockey coaches to create a more equitable mindset and environment which can foster deep connections and relationships with players, families and caregivers. Finally, the current study reinforces the need to provide support for volunteer coaches who require additional training to improve knowledge of culture, diversity and leadership within hockey.

References

Athletic Research (2025). Team dynamics and communication in sport: Boosting cohesion in sport. https://www.athleticinsight.com/sports-psychology/team-dynamics\-and-communication

de Haan E. Nilsson V.O. (2023). What can we know about the effectiveness of coaching? a meta-analysis based only on randomized controlled trials. *Academy of Management Learning & Education*. https://doi.org/10.5465/amle.2022.0107

Denzin N. K., Lincoln Y. S. (2017). *The SAGE handbook of qualitative research* (5th ed.). Sage Publications.

Gómez-Hurtado, I., González-Falcón, I., Coronel, J.M. (2018). Perceptions of secondary school principals on management of cultural diversity in Spain. The challenge of educational leadership. *Educational Management, Administration & Leadership*, 46(3), 441-456. https://doi.org/10.1177/1741143216670651

Gomez-Hurtado, I., Gonzalez-Falcon, R. V. I., Vargas, F. J. (2021). Inclusive Leadership: Good Managerial Practices to Address Cultural Diversity in Schools. *Social Inclusion*, 9(4), 69–80. https://doi.org/10.17645/si.v9i4.4611

Hansen, D.M., Larson, R.W., Dworkin, J.B. (2003). What adolescents Learn in Organized Youth Activities: A Survey of Self-Reported Developmental Experiences, *Journal of Research on Adolescence*, 2003, 13, 25-55.

Hockey Canada (2025). Hockey Canada coaching program. https://www.hockeycanada.ca/en-ca/hockey-programs/coaching/essentials.

Horn, T.S., Coaching Effectiveness in the Sport Domain, in: Horn, T.S., ed. (2008). *Advances in Sport Psychology*, Human Kinetics, Champaign, pp. 239-268.

Rana, S., Baumgardner, B., Germanic, O., Graff, R., Korum, K., Mueller, M., Randall, S., Simons, T., Stokes, G., Xiong, W., Peterson, K. (2013). From youth worker professional development to organizational change. *New Directions for Youth Development*, 139, 27–57. https://doi.org/10.1002/yd.20068

Rönnström, N., Scott, P. (2019). Lideres escolares, maestros de la complejidad. In J. Weinstein & G. Muñoz (Eds.), *Liderazgo en escuelas de alta complejidadsociocultural.* Diez Miradas [Leadership in schools of high sociocultural complexity. Ten looks] (pp. 11–152). Ediciones Universidad Diego Portales.

Industry papers


BEARMIND
www.bearmind.tech

49ing

Vertical AI for Ice Hockey: from raw video to decisions that move the scoreboard

49ing builds AI systems that turn multi-angle game video into coaching-grade insights in minutes. Our **Data Cockpit** ingests venue and broadcast feeds, auto-tracks players and puck, detects events, and surfaces **actionable clips** with context. Coaches, scouts, and refereeing departments use 49ing daily for game prep, player development, opposition scouting, and post-game review.

What's new

– **Natural-language video search**: "Show all controlled DZ breakouts with possession exit that lead to OZ entries"
– **Attack Scenarios library**: Off the rush, forecheck, OZ faceoffs, sustained OZ possessions—each with KPIs and teachable clips.
– **Goalie intelligence**: xG/xGOT shot maps with pre-shot movement, traffic, and rebound control.
– **Mobile coaching assistant**: On-bench retrieval of tagged sequences and set-play reminders.
– **AI Officiating support (pilot)**: Multi-angle incident triage to speed up review workflows.

Why teams choose 49ing

– **Speed**: Seconds to minutes from ingest to indexed video with auto-tags.
– **Precision**: Multi-camera tracking + event models trained on elite-level data.
– **Workflow fit**: Exports to common video tools; shareable playlists.
– **Trust**: Transparent definitions, versioned models, and human-in-the-loop review.

What we're studying next

Possession value under different forecheck schemes, chain-of-events causality (entry → OZ time → shot quality), and load-to-fatigue links that predict execution late in games.

Contact: andreas.haenni@49ing.ch — 49ing AG, Zürich, Switzerland

Hat-TriQ

by Stretch On Sense

The magic of Hat-TriQ

Hat-TriQ is a unique analysis platform assisting elite teams with detailed player analysis. How the players practice, how the players play the game, and the player's well-being are all gathered on one platform. Focused on empowering decision-makers with data, Hat-TriQ gives you detailed player and team analysis so your organization can make decisions based on true facts. As we combine all systems into one, you will reach your full potential for your entire organization. **With Hat-TriQ, there are no silos. IT's one platform.**

Hat-TriQ – The sports analytics' Spotify

Before Spotify, iTunes, and even Napster, we listened to music on a CD player.

If you wanted to listen to Bruce Springsteen's Born To Run, the song, not the whole album, you needed to buy that CD. If you wanted to listen to Glory Days from the album Born in the U.S.A., you needed to buy a second album from Bruce Springsteen. In that case, you now have 20 songs, but you only wanted to listen to 2 songs. You have too much music for your appetite and desire, which you also can compare to too much unfiltered data.

With Spotify, you have one platform with all artists, bands, and songs from the past, and what's fresh today and what will be cool and fresh tomorrow. You can choose one song from one album and quickly switch to a new artist and song. You can access all the music in the world and easily filter and listen to what you like.

With Hat-TriQ, you have the same option. We offer you all the systems on the market into one platform, and we can easily be the filter of your choice. We customized Hat-TriQ for you as Spotify creates your playlists.

HAT-TRIQ

STRETCH ON SENSE

Qlik

Some questions to hockey analytics experts

Miika Arponen, Ässät Pori.

Who are you and what is your connection to hockey analytics?

I am Miika Arponen and I work as a data analyst for Ässät Pori in the Finnish Liiga. I have been doing this for four seasons now, and before that I worked as a hockey journalist for over a decade. My background is in software development.

How do you use hockey analytics in your job?

As a data analyst, hockey analytics pretty much *are* my job. I gather the data from multiple sources, research and analyze them and then provide conclusions to the GM and the coaching staff.

How do you communicate hockey analytics findings to your customers/viewers/players/coaches?

Mostly just verbally in conversations, but I have also created some dashboards and tools for the GM and the coaches to use, and occasionally I also produce written reports and graphs.

What hockey question would you like hockey analytics to answer next?

Anything related to tactics and playing formations, like "what is the most efficient and secure way to bring the puck up the ice on the power play". Obviously no one conclusive answer will ever be found, but that is the interesting next step I think.

Which hockey analytics method/notion is the most important/influential in your job?

While being a bit boring, I'd still say the expected goals as a lot of the more complicated and complex methods and models are ultimately based on them after all one way or another.

Where is the hockey analytics field going? What do you envision for the next 10 years?

In 10 years the AI will very probably be a lot more relevant than it already is. I'm fairly sure human analysts can focus more on the interpretation and innovation and less on how to get the numbers and how to transform them.

What was your main take-away from LINHAC 2025?

As always, the people are by far the biggest asset the conference has. Getting to know the brilliant minds in the business is very valuable and interesting.

Robert Bandemer, d-fine.

Who are you and what is your connection to hockey analytics?

As a member of the Sports Analytics team at d-fine, I specialize in leveraging sports data for various applications in my daily work. Having a background in Football Analytics, transitioning to projects with sports clubs and federations in the world of ice hockey was a natural progression. In this space, we utilize available data for use cases in the areas of match analysis, scouting and more.

How do you use hockey analytics in your job?

In my role, hockey analytics involves leveraging data to provide actionable insights across various areas of the game. We analyze large datasets collected from matches, tracking data, and player performance to identify patterns, trends, and key performance indicators. For example, in match analysis, we dive deep into data to assess team strategies, measure puck possession, shot quality and zone entries/exits. This helps teams evaluate their performance and prepare for future opponents.

How do you communicate hockey analytics findings to your customers/viewers/players/coaches?

When communicating hockey analytics findings, I prioritize keeping things simple, clear, and actionable. I focus on delivering easy-to-understand, tangible insights rather than overwhelming them with overly detailed or technical explanations.

What hockey question would you like hockey analytics to answer next?

There are still many fundamental questions in hockey analytics left to explore, but one area I'm particularly interested in is understanding player decision-making. Gaining deeper insights into why players make certain choices (whether it's passing, shooting or positioning) could provide a new dimension to analytics. However, this is incredibly challenging to measure, as the data required to truly capture decision-making is still quite limited.

Where is the hockey analytics field going? What do you envision for the next 10 years?

The hockey analytics landscape, particularly in Europe, needs to grow and establish itself further before fully unlocking its potential. Over the next 10 years, I see tracking data becoming more widely utilized and skeleton data playing a larger role, providing deeper insights into player movements and mechanics. With advancements in data infrastructure and acceptance of analytics, the field will drive significant improvements in player development, strategy, and performance optimization.

What was your main take-away from LINHAC 2025?

My main takeaway from LINHAC 2025 was the variety of use cases and innovative approaches in hockey analytics, particularly the growing application of skeleton data and computer vision. I also found it fascinating to see how hockey analytics is being combined with insights from other sports, like football, to share knowledge and learn from each other's advancements.

AJ Bernstein, San Jose Sharks.

Who are you and what is your connection to hockey analytics?

AJ Bernstein, Strategy and Data Analyst at the San Jose Sharks.

How do you use hockey analytics in your job?

Analytics act as a piece of the puzzle in our larger evaluation process, providing insights that may not be captured by our other information sources.

How do you communicate hockey analytics findings to your customers/viewers/players/coaches?

The way coaches and management use data differs, so we tailor our delivery accordingly. Sometimes a simple PowerPoint is best; other times a detailed discussion is necessary.

What hockey question would you like hockey analytics to answer next?

I am particularly interested in leveraging detailed tracking data to better measure individual skills.

Which hockey analytics method/notion is the most important/influential in your job?

The most important analytic method in my job is not solely data-related, but a disciplined, evidence-based approach to problem-solving.

Where is the hockey analytics field going? What do you envision for the next 10 years?

We are in a transitional phase in the growth of hockey analytics. The value of a data-driven approach is established, and the task now is to determine which data is truly useful. As access to and demand for data expand, it will become even more important to identify and prioritize the information that actually matters for each organization.

What was your main takeaway from LINHAC 2025?

I really appreciated the ability to discuss the similarities and differences among all the teams present on my panel.

Andreas Hänni, 49ing.

Who are you and what is your connection to hockey analytics?

49ing is a Swiss sports-tech company focused on vertical AI for ice hockey. We combine computer vision, tracking, and sequence modeling to turn video into decisions for coaches, scouts, players, and officials.

How do you use hockey analytics in your job?

Day-to-day we: (1) ingest track multi-angle video; (2) detect events; (3) value sequences; and (4) deliver searchable clips, reports, and dashboards in our platform **Data Cockpit** and our mobile app **Hockey AI**.

How do you communicate hockey analytics findings to your customers/viewers/players/coaches?

Three layers:

– Clips-first: short, labeled videos with "why it matters" overlays.
– Dashboards: scenario KPIs (rush vs. forecheck vs. OZ), goalie shot maps, and matchup drill-downs.
– Narratives: auto-generated scouting notes that coaches can edit.

Which hockey analytics method/notion is the most important/influential in your job?

Reliable multi-camera tracking & re-identification (foundation), plus sequence models that score chains of actions rather than single events. For coaches, explainability (transparent definitions, comparable baselines) is just as important as model accuracy.

Where is the hockey analytics field going? What do you envision for the next 10 years?

From post-hoc reporting to real-time, agentic assistants: on-bench retrieval, automated between-period prep, opponent-specific set-play suggestions, and officiating triage - underpinned by standardized data schemas, privacy-by-design pipelines, and cross-league video exchange.

What was your main takeaway from LINHAC 2025?

The gap between **research prototypes** and **coach-ready workflows** is narrowing fast - especially around possession modeling and tracking.

Anything else the LINHAC audience should know?

We're open to research collaborations (method benchmarks, shared definitions, and evaluation datasets). If you're building methods that need video + ground truth events, we're happy to explore structured access under the right agreements.

Martin Lundholm, Skellefteå AIK.

Who are you and what is your connection to hockey analytics?

My name is Martin Lundholm, and I work as an Analyst and Scout for Skellefteå AIK Hockey in the Swedish Hockey League (SHL).

How do you use hockey analytics in your job?

Hockey analytics is the foundation and core of the work I do for the club. I use analytics - data analysis to gain insights and make decisions, both at the team and individual player level. This includes improving our own play, scouting opposing teams and players, and generally providing data-driven support to decision-makers in the organization, such as the GM and coaching staff.

How do you communicate hockey analytics findings to your customers/viewers/players/coaches?

Communication varies depending on the role within the organization. With the GM, I use verbal discussions and reports/dashboards, along with numbers, graphs, tables, etc. With the coaching staff, the messaging is more simplified, concise, and to the point.

Each year, I've focused more on how the message is delivered and less on the details themselves. A coach is probably not interested in knowing whether we had 30.7 or 32.4 controlled zone entry attempts - they care more about whether we're following the system we agreed to play.

What hockey question would you like hockey analytics to answer next?

Measuring actions by players without the puck - such as positioning and its impact on performance on the ice. This includes defensive presence, backchecking, and pressuring puck-carrying opponents etcetera.

Which hockey analytics method/notion is the most important/influential in your job?

There isn't a single best method - it depends on who you're communicating with. The key is turning raw data and numbers into visually accessible dashboards and reports, preferably with color, images and video to enhance understanding. It's also crucial to adapt the language from overly technical terms or numbers into familiar hockey terminology.

Where is the hockey analytics field going? What do you envision for the next 10 years?

I believe analytics is becoming integrated into every part of a hockey club's organizational structure - not just with the men's senior team, but also in data-driven decision-making for women's and junior teams, as well as in non-sporting departments like communications and sales.

What was your main takeaway from LINHAC 2025?

Unfortunately, there's still a significant gap between academia, public work, and professional hockey clubs. We need more accessible data for the general public and the average hockey enthusiast to explore, experiment with, and learn from - "trial and error".

Another takeaway is that there's a lot of interesting knowledge and inspiration to be drawn from other sports, such as football (soccer).

Is there something else related to hockey analytics that you would like the LINHAC audience to know?

If you're truly passionate about hockey analytics, there is a path to a role within the field. Educate yourself, learn, dive deep - and don't be afraid to show what you can do. Whether through your X (Twitter) account, a personal website, published reports, or simply by reaching out to your local club and offering your knowledge and services. Learn, and then don't hesitate to raise your hand and get involved!

It doesn't need to be perfect, just start doing it.

I wish you all good luck.

Albin N Maelum, Stretch On Sense.

Who are you and what is your connection to hockey analytics?

We are a sports analytics company specialized in transforming hockey data into actionable insights. Our connection to hockey analytics is direct — we work with professional leagues and clubs to help them leverage data for decision-making both on and off the ice.

How do you use hockey analytics in your job?

We use hockey analytics to identify patterns, measure performance, and provide clear recommendations to teams, coaches, and organizations. That means everything from game analysis and tactical insights to organizational reporting and fan engagement.

How do you communicate hockey analytics findings to your customers/viewers/players/coaches?

Communication is key. We focus on making complex data simple, visual, and tailored to the audience. Coaches and players get actionable performance insights; managers and directors receive clear reports for decision-making; fans see the bigger picture in an engaging way.

Which hockey analytics method/notion is the most important/influential in your job?

Contextualization. Numbers alone don't create value - it's when data is put into the right context, whether tactical, physical, or financial, that analytics truly influences decisions.

Where is the hockey analytics field going? What do you envision for the next 10 years?

We see hockey analytics moving toward real-time, integrated solutions. In 10 years, teams will rely on live data during games, predictive models for player development and health, and fully connected systems that align sporting performance with business outcomes.

What was your main take-away from LINHAC 2025?

The main takeaway was that collaboration across research, technology, and the hockey community is stronger than ever. Analytics is no longer just a niche – it's becoming a central part of how hockey is played, coached, and managed.

Is there something else related to hockey analytics that you would like the LINHAC audience know?

Yes – that the future of hockey analytics isn't only about better models or bigger datasets. It's about accessibility. If we want analytics to make a real impact, it must be presented in ways that players, coaches, executives, and fans can all understand and use.

Josh Pohlkamp-Hartt, Boston Bruins.

Who are you and what is your connection to hockey analytics?

Josh Pohlkamp-Hartt, Associate Director Of Hockey Analytics with the Boston Bruins.

How do you use hockey analytics in your job?

I work primarily as an analyst and data scientist in my role with the Bruins. These 2 roles directly use hockey analytics: designing models and metrics, then evaluating them to help our team make better decisions.

How do you communicate hockey analytics findings to your customers/viewers/players/coaches?

The key is to find common language and focus on actionable information. We usually convert our analysis from being numerical to more descriptive then provide examples through video of past experiences. An example might be to convince a player to shoot more from dangerous areas, we could talk about their XG created with their shots on ice but that will not resonate. Instead when talking with them, we are better suited to provide some video examples of low danger shot patterns they take and situations where they did not act to take more dangerous shot opportunities.

Additionally, it is important to build trust by working collaboratively and being open to feedback. Feedback from hockey experts is invaluable in improving analytical quality.

What hockey question would you like hockey analytics to answer next?

Defensive attribution and injury prevention are my biggest interests. The quality of a defensive specialist is tough to measure and most metrics fail to account for the opportunities they suppress by occupying space correctly. Injuries, especially non-contact ones, can be reduced with smart training and deployment. I believe that there are improvement that can be made through the evaluation of the biomechanics of a player with more detailed pose data.

Which hockey analytics method/notion is the most important/influential in your job?

Opportunity is more repeatable than execution. If as a team you have an effective strategy to optimize your share of the offensive opportunities, you can more often repeatably win this battle. Whereas, if you rely on goaltending and shooters to save and score above expected, you have a smaller set of events (there are less shots than puck possession touches in a game) and are more susceptible to randomness. The true NHL contenders are tuned for opportunity.

Where is the hockey analytics field going? What do you envision for the next 10 years?

Our data detail is growing at a rate that is outpacing the research community and over the next 10 years we will see some exciting growth in terms of team-side financial commitments and public analysis to fill in the gaps. This should open up exciting areas of research into decision making and multi-player strategies like we have seen in soccer.

What was your main takeaway from LINHAC 2025?

The concept of AI is an accepted norm now and we should be using this to our advantage by pushing the envelope on the ways we can integrate hockey analytics into our decision making and improve the easy of use for our users.

Is there something else related to hockey analytics that you would like the LINHAC audience to know?

No, they seem well informed!

Freddie Sjögren, Freddie Sjögren Consulting AB.

Who are you and what is your connection to hockey analytics?

I am the founder of *Performance Facilitator*, a service that helps elite hockey clubs build and optimize their performance departments. My background combines sports science, applied performance, and research in how physical tracking data can be connected to tactical and technical AI-generated data. My focus is to raise the standard in hockey by making data more meaningful and actionable for decision makers, players, and organizations.

How do you use hockey analytics in your job?

I use analytics to identify patterns that link physical performance with tactical behavior and technical execution. This means combining tracking data with video and contextual factors to better understand availability, readiness, and ROI on player investments. Analytics are not just numbers for me - they are tools to help performance staff, coaches, and executives make smarter, evidence-based decisions.

How do you communicate hockey analytics findings to your customers/viewers/players/coaches?

I translate complex data into clear, actionable insights. For executives and GMs, that often means ROI-focused dashboards and reports that connect performance data to financial impact. For coaches and players, it means practical visuals, clips, and simple metrics that can guide training, recovery, and game preparation. The key is always to adapt the language of analytics to the audience.

What hockey question would you like hockey analytics to answer next?

I would like to better understand the true connection between *physical load, tactical decision-making, and technical execution.* For example, how fatigue or intensity in a certain shift directly affects passing accuracy, puck battle efficiency, or tactical choices. This integration would allow teams to manage players more precisely and improve both performance and longevity.

Which hockey analytics method/notion is the most important/influential in your job?

Contextualized tracking data – not just raw skating metrics, but tracking that is tied to tactical situations and technical skills. When analytics are layered with context, they shift from being descriptive to being predictive and prescriptive, which is where the real value lies.

Where is the hockey analytics field going? What do you envision for the next 10 years?

I believe we are moving from siloed analytics (physical, tactical, technical, medical) toward integrated performance ecosystems where AI helps combine all streams. In 10 years, hockey organizations will rely on connected data platforms that allow real-time decision-making on player availability, tactics, and investment strategies. The clubs that embrace this holistic, evidence-based approach will be the most successful.

What was your main takeaway from LINHAC 2025?

The main take-away for me was the need to bridge the gap between cutting-edge analytics research and real-world application. There is so much innovation happening, but the key challenge is implementation – making sure that the insights actually improve decision-making in the daily environment of clubs and players.

Is there something else related to hockey analytics that you would like the LINHAC audience know?

Yes – that analytics should not only be about performance enhancement but also about *sustainability*. Player health, career longevity, and smarter financial decisions are equally important outcomes of analytics. If we want hockey to evolve, analytics must serve the bigger picture of availability, well-being, and return on investment.

Erik Wilderoth, Färjestads BK.

Who are you and what is your connection to hockey analytics?

I'm the assistant general manager and hockey analyst for Färjestad BK. My role is to translate raw data and statistical insights into practical, strategic decisions that improve our team, from player recruitment to game preparation. I serve as a bridge between the world of numbers and the daily operations on the ice.

How do you use hockey analytics in your job?

I use hockey analytics in several ways: Player Recruitment, Opponent Analysis, Evaluating and developing our own players.

How do you communicate hockey analytics findings to your customers/viewers/players/coaches?

Communication is key as Lignell talked about at the 2024 conference. I avoid overwhelming coaches and players with complex terms or numbers with decimals. Instead, try to use a common language. Story telling is in most cases a better way to get facts to stick.

I use visualizations, such as graphs and video clips with overlays, to strengthen the narrative. Our goal is to give the players and coaches tools and insights to get better, not just numbers.

What hockey question would you like hockey analytics to answer next?

Defense is still the hardest part to quantify in my opinion. How good the players are with the puck, that's easy. But some players have the puck like < 1 minute per game. What does he spend the rest of the time and to what success rate?

Which hockey analytics method/notion is the most important/influential in your job?

For me, the most influential method is Expected Goals (xG), especially when combined with individual player performance. xG gives us a much more nuanced picture of chances created than traditional shot metrics. By looking at xG, we can measure performance. Performance leads to results.

Where is the hockey analytics field going? What do you envision for the next 10 years?

In Europe I must say that it's on a plateau for the moment. The software companies are good at convincing teams that the software is of such ease to use that analysts don't get hired. This is a trend we need to break and much is up to teams to understand the potential of building stuff on our own.

I would love to say that all top leagues in Europe have at least one analyst hired in 10 years and we in SHL have at least 2, but that feels like utopia today.

What was your main take-away from LINHAC 2025?

My biggest insight from LINHAC 2025 is that North America is moving fast. Europe slower. We need to step up. The student competition gets better for every year and also the research papers.

Morgan Zeba, Spiideo.

Who are you and what is your connection to hockey analytics?

I'm Director of Sales Europe at Spiideo, focused on ice hockey. I work with leagues, clubs, and federations worldwide, helping them implement video and analytics solutions. Our ambition is to be the **global leader and natural choice** for hockey.

How do you use hockey analytics in your job?

Spiideo delivers a full ecosystem:

- **Spiideo Replay** for officiating, coach's challenge, and **player safety reviews**
- **Spiideo League Exchange** so leagues can seamlessly share video between teams and, referees/officials
- **Spiideo Perform** for every team to develop their players with video and analytics
- **Spiideo Play** so every league can stream their games fully automatically and in multiple angles.

By combining automated capture with integrated data, we help organizations improve decisions, increase safety, and create long-term value.

How do you communicate findings?

Always by connecting analytics to real workflows:
Coaches → player development with Perform
Referees → clear decisions with Replay (our VAR solution)
Safety staff → evaluate hits, head contact, and return-to-play
Leagues → collaboration through League Exchange
Broadcasters → richer storytelling with synchronized video + metrics

What's next for hockey analytics?

The future lies not only in moving from "what happened" to **why it happened**, but also in making data **easier to access and work with**. With AI, complex insights will become intuitive so that coaches and players can seamlessly use analytics every day. At the same time, analytics will play a growing role in reducing risks and strengthening **player safety**.

Where is the field going?

We already provide real-time analytics today, uniquely with multiple camera angles synchronized live — something no one else can deliver at this scale. What's next is to offer live player tracking into games, adding a new dimension for coaching, officiating, and player safety. Combined with AI that identifies patterns automatically, this will transform how the game is understood and managed.

Final note

With Replay, League Exchange, Perform and Play, Spiideo offers an **automated, multi-angle, cost-efficient, and high-quality solution** for leagues, referees, and teams. We are committed to being the **global leader in hockey video, analytics, and player safety**.

Some questions to football analytics experts

Ola Lidmark Eriksson,
Football Analytics Sweden AB.

Who are you and what is your connection to football analytics?

Ola Lidmark Eriksson, first and foremost working as CTO at Football Analytics Sweden AB where I am CTO and responsible for our analytics platform used by 61 of 64 elite football clubs in Sweden. I also have a long history as a pundit on Swedish TV focusing on data and analytics within football.

How do you use football analytics in your job?

I use analytics every day and have done so for the last 8 years. All the best insights and know-how I have collected over the years we try to incorporate into our analytics platform. It can be new KPIs, how to aggregate data, or a new visualization. The idea is always to make it accessible for everyone using our platform - no coding needed. Honestly, it is also for my own convenience sometimes - not having to use Jupyter notebooks or some Python script. It should be easier and faster.

How do you communicate football analytics findings to your customers/viewers/players/coaches?

We have our website where we try to post news and insights. But apart from that I'd say that the main area for myself is that I hold presentations on several courses and seminars held by the league organisation and the federation every year in Sweden.

What football question would you like football analytics to answer next?

I am very interested in football's pretty rigid idea about forwards, midfielders and defenders. And goalkeepers, of course. What I would like to work more on is how, using data, one could start to question why we must have these denominations of players. Shouldn't everyone on the pitch collectively attack and defend and hence not be called different things? I'd like to see the same being discussed in hockey. PP, for instance, is a moment where I really cannot understand any other concept than that all 6 players on the PP side all attack.

Which football analytics method/notion is the most important/influential in your job?

In football I'd say that over the years, as our modelling has improved as well as the data collection, I more and more lean towards all-in-one methods/KPIs. It can be clustering based on players' playing styles, but even more the "value" ones measuring players' abilities as unabstract as possible with simple 1-5 grades.

Where is the hockey analytics field going? What do you envision for the next 10 years?

I hope it will catch up with football, as I right now definitely can see that football is ahead, especially when it comes to actually being used in organisations.

What was your main takeaway from LINHAC 2025?

I think that was my main takeaway - the former statement. That hockey is behind football right now. At least in Europe.

Is there something else related to football analytics that you would like the LINHAC audience to know?

Maybe that I have started to try to take our football platform to hockey.

Pieter Robberechts, KU Leuven.

Who are you and what is your connection to football analytics?

I'm Pieter Robberechts, currently a PhD researcher within the DTAI research group at KULeuven, specializing in data science and machine learning applied to sports (especially football).

How do you use football analytics in your job?

I do research in data science and machine learning, where football often serves as an application domain. In some projects, we start from a concrete problem in football (e.g., Can we quantify how a player performs under high mental pressure?) and develop AI methods to solve it. In other cases, we use football data as a rich, real-world use case to explore and validate new AI and machine learning techniques developed in our group.

How do you communicate football analytics findings to your customers/viewers/players/coaches?

Our work focuses more on methodological development than on directly delivering insights to practitioners. However, we place strong emphasis on making our methods understandable and accessible to non-technical audiences. We do this in three ways: First, we try to design our models around concepts that make sense within the context of the sport. For instance, when creating features for a model, we prioritize domain-specific variables rather than relying on abstract, auto-generated features or deep learning. This makes the outputs easier for coaches and analysts to interpret. Second, we maintain a blog (https://dtai.cs.kuleuven.be/sports/blog/) where we translate our research into more accessible content for a broader audience. Third, beyond building models, we often take a step back to critically assess how these models should be interpreted. As such, we try to contribute to a more thoughtful and transparent use of AI in sports.

What football question would you like football analytics to answer next?

One key question I'd like football analytics to tackle is: "How would the model's prediction change if we replaced one player with another?" This kind of counterfactual reasoning is often the first thing practitioners ask when presented with model outputs. It highlights a major challenge in current models: their lack of player- and team-specific understanding. Most models today struggle to capture the differences between, say, building an attack around Kylian Mbappé versus Robert Lewandowski. This limits both the trust practitioners place in the models and the insights we can derive. I hope the next big step will be developing models that better disentangle individual player characteristics.

Which football analytics method/notion is the most important/influential in your job?

I hate to give the obvious answer, but it's probably still xG. It fundamentally changed how we think about evaluating actions in context. It introduced a probabilistic, data-driven way to assess the quality of actions and most new metrics that we develop today are still based on the same core idea. Also, even though xG has been around for over a decade, we're still discovering better ways to model, apply, and interpret it.

Where is the football analytics field going? What do you envision for the next 10 years?

The evolution of football analytics is closely tied to the evolution of data: from basic statistics to event stream data, and more recently to tracking data. Looking ahead, I believe the next major shift will be driven by the increasing availability and use of 3D pose tracking data, which captures detailed full-body movement. This type of data opens the door to a much deeper understanding of decision-making, biomechanics, and physical performance. Also, we'll likely see analytics evolve more from analyzing what happened to explaining why and how it happened. Furthermore, over the next decade, I believe we'll see analytics become more personalized and grounded in the physical, technical and perhaps cognitive capabilities of players.

What was your main takeaway from LINHAC 2025?

The presentation that stood out most to me was by Mikael Svarén. The main thing I took from it is that there's still a lot of untapped potential in collaboration between sports scientists and data scientists, particularly around the use of 3D pose tracking data. Also, as Devin Pleurel mentioned during one of the panels: this is by far the largest dataset ever in biomechanics. To fully realize its potential, we need much tighter collaboration between the domain expertise of sports scientists and the methodological tools of data scientists.

Is there something else related to football analytics that you would like the LINHAC audience to know?

Yes - one of the key drivers of progress in football analytics in recent years has been the increasing availability of open-source data and software packages, which have made the field far more accessible and collaborative. The availability of open datasets has significantly lowered the barrier to entry, and the availability of high-quality community-driven tooling (e.g., kloppy, mplsoccer) accelerates the development of new methods. I believe hockey analytics could really benefit from a similar push.

Student competition papers

Identifying and Analyzing SHL Ice Hockey Match Styles Based on Event Data Aggregation and K-Means Clustering

Yanjie Lyu, Qingxuan Cui, Huaide Liu, Han Xia, Yi Yang

Linköping University, Linköping, Sweden

1 Introduction

As a high-speed and dynamic team invasion sport, ice hockey's game strategy and tactical choices have a crucial impact on the result of the game [1]. While traditional post-match statistics such as goals, assists, +/- values, etc., provide information about the result, they often fail to reveal the overall strategic tendencies adopted by the team during the game, the so-called "game style" [2]. Understanding the game styles of different teams not only helps to assess team characteristics and prepare for games, but also provides a reference for player recruitment and roster construction [3].

With the development of data collection technology, detailed event data makes it possible to quantitatively analyze the style of the game [4]. At present, relevant institutions or organizations have begun to explore the use of these data for analysis. For example, clustering based on the positioning of players during the game to distinguish player roles and styles [5], research that focuses on specific behaviors such as passing networks or space utilization patterns [6], or dedicated to evaluating the value of players or actions through machine learning models [7]. In recent years, a promising direction has been to draw on the experience of other team projects (such as field hockey), and utilize clustering algorithms to process statistical data so that macroscopic game styles can be identified [2].

This paper aims to apply and extend this methodology to SHL hockey game data with the following contributions:

1. Identify distinct playing styles in the SHL by applying a K-Means clustering algorithm to **game-level aggregated event statistics**, where styles are characterized separately for the two teams participating in each match, thereby capturing the variability of tactical patterns across individual games rather than season-level aggregates.

2. Characterize the identified styles using statistical metrics and visualization techniques, such as radar charts and spatial heatmaps.

3. Rigorously validate the effectiveness of these styles and their matchup dynamics through statistical testing, including Bootstrap resampling and Chi-squared tests, while correcting for sample size imbalances using Bayesian averaging.

4. Introduce a game-theoretic framework to model the strategic interactions between styles, identifying dominant strategies and the resulting Nash Equilibrium within the league's tactical meta.

2 Background

The data basis for this paper is the Swedish Ice Hockey League (SHL) detailed game event data provided and licensed by Sportlogiq Inc. This kind of refined event data records the details (such as time, player, team, coordinates, result, etc.) of every pass, dribbled, shot, zone transition, etc., during the game. However, it is challenging to extract meaningful patterns from high-dimensional, time-series event data, especially macro team tactical styles [8].

Subsequently, K-Means clustering, a widely used unsupervised algorithm, is employed to partition matches into distinct tactical style clusters [9]. To clearly interpret and present these tactical styles, the study utilizes visualization methods: radar charts illustrate the multi-dimensional performance profiles of different styles, and heatmaps visualize the winning rates between competing tactical styles, revealing potential interactions and constraints among them.

3 Methodology

This study adopts a multi-stage analysis process, aiming to identify, quantify and interpret the game styles at the team level from the SHL match event data, and ultimately evaluate the relative effectiveness of different styles. This process integrates methods such as data processing, feature engineering, unsupervised clustering, supervised learning interpretation, and adversarial analysis.

3.1 Data Preparation and Feature Extraction

The research data is derived from the SHL race event log provided by Sportlogiq. The original data was first preprocessed and sorted by 'gameid' and 'compiledgametime'. This was followed by meticulous feature engineering designed to translate discrete match events into continuous or count-type variables capable of capturing a team's tactical intent and execution efficiency. The key steps include: (a) Created numerical or categorical identifiers for core game events such as passing, shooting, puck possession entering, blocking, clearing, etc.; (b) Combined the event outcome field to quantify the successfully executed actions such as the number of successful passes and the number of successful area entries; (c) For the carry event, by calculating the spatio-temporal differences between adjacent events, the carry_duration and carry_distance were extracted to reflect the mode of advancing with the puck. The goal of this part is to build a rich feature set to lay the foundation for the subsequent aggregation of team performances.

3.2 Team Performance Aggregation and Metric Construction

To achieve the transformation of the analytical scale from micro events to macro team performance, aggregating the variables extracted by feature engineering according to team (teamid) and game (gameid) is a proper way. In order to be able to calculate the key efficiency indicators and capture the performance in other dimensions, the aggregation strategy is determined based on the nature of the variables: the number of event occurrences (such as num_passes, num_successful_passes, which are the basis for calculating the efficiency indicators) can be obtained by summation the corresponding indicator variables, The expected goals (xg_allattempts), stick handling time and distance, etc. can be calculated as the average value of a single game. Subsequently, using these aggregated results, efficiency indicators such as pass_success_rate were constructed by calculating the corresponding ratios.

3.3 Game Style Identification using K-Means Clustering

To explore and discover the potential, data-driven game style in the SHL competition, K-Means clustering algorithm was adopted. Given that K-Means is sensitive to the scale of input features, 13 aggregated features covering aspects such as puck control, offense, defense and efficiency were selected , and they were standardized to ensure that each feature has zero mean and unit variance. The optimal number of clusters was determined by assessing the performance of various k values using the Silhouette Score in conjunction with the Elbow Method. Although all Silhouette Scores for k were below 0.2, indicating relatively weak cohesion and separation, such results are common in clustering tasks involving complex, multi-dimensional behaviors such as team playing styles. The results showed that the Silhouette Score peaked at $k = 3$, and the Elbow plot of within-cluster sum of squares also exhibited a noticeable inflection at $k \approx 3 - 4$. In the case of $k = 3$, the three resulting styles—High-Pressure Offense, Defensive Counterattack, and Puck Control Play—demonstrated clear tactical semantics and presented distinct visual separations. This structure simultaneously balances the mathematical optimization of the clustering algorithm and the interpretability of ice hockey tactical analysis. Consequently, $k = 3$ was identified as the optimal number of clusters.

3.4 Evaluation of Inter-Style Effectiveness

Finally, in order to preliminarily evaluate the effect of the identified game styles in actual confrontations, the competition results under different style combinations were analyzed. By matching the two opposing teams in the same match and their style tags and combining with the game results, the average winning rate (scoring rate, win=1, draw=0.5, loss=0) of each style when facing the specific opponent's style was calculated. This pairwise comparison helps to reveal the underlying restraint relationship among styles.

4 Results and Discussion

Through K-Means clustering method in this study, three game styles with significantly different characteristics were successfully identified. An evaluation of the win rates between these styles allows for a detailed characterization of each, an identification of key driving factors, and an exploration of their practical effectiveness.

4.1 SHL Game Style Profiles and Key Features

The three main game styles discovered by K-Means clustering have been initially named "Defensive Counterattack", "High-Pressure Offense" and "Puck Control Play". Figure 1 visually presents the average performance differences of these three styles on 13 standardized aggregated features.

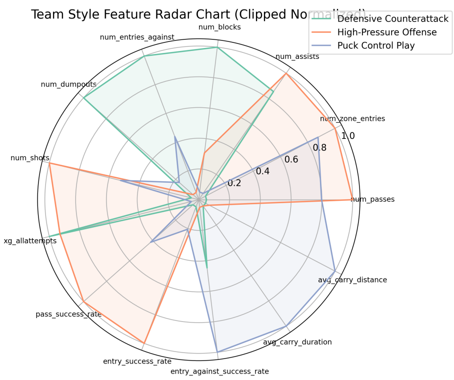

Fig. 1. Team Style Feature Radar Chart (Clipped Normalized)

By synthesizing the information from Figure 1, the profiles of the three playing styles are clearly depicted:

(a) **Defensive Counterattack** This style excels in defensive metrics, while scoring low on most offensive and possession metrics. This highlights its strategy of prioritizing solid defense and relying on quick transitions.

(b) **High-Pressure Offensive** This style leads in offensive outputs, with a notably high number of passes, reflecting its aggressive approach of high-frequency pressing and creating numerous shooting opportunities.

(c) **Puck Control Play** This style is distinguished by exceptionally high average carry distance and average carry duration, as well as the highest entry success rate, with number of passes being the most significant differentiating factor for this style (Figure 2). This indicates its core strategy of controlling the tempo of the game through long periods of possession and high success rate in advancing the puck. However, its relatively low number of shots and expected goals from all attempts (Figure 1) suggest potential limitations in converting possession advantage into concrete offensive threats.

4.2 Discovery and Validation of Style Matchup Dynamics

The performance of different playing styles was evaluated in this study, and the results show significant differences in their win-loss outcomes.

A preliminary analysis reveals a potential counter-relationship between playing styles (Figure 2):(a)The **Defensive Counterattack** style not only effectively suppresses the **Puck Control Play** style(win rate 82%), but also slightly outperforms the **High-Pressure Offensive** style(win rate 55%); (b)The **High-Pressure Offensive** style, while strongly dominating the **Puck Control Play** style(win rate 86%),is at a disadvantage when facing the **Defensive Counterattack** style (win rate 45%); (c)The **Puck Control Play** style struggles against both other styles, with win rates of 18% and 14%, respectively.

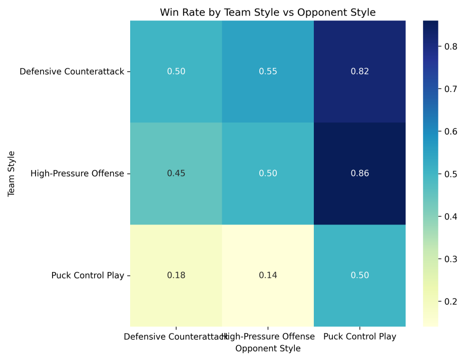

Fig. 2. Win Rate by Team Style vs Opponent Style

Subsequently, a series of rigorous statistical tests and model adjustments were performed, adhering to a "validate–correct–revalidate" framework, to ensure the robustness of the key findings and to mitigate potential biases arising from data imbalance.

4.2.1 Validation of Cluster Distinction in Tactical Features

To further validate the distinctiveness of the clustering results in the tactical dimension, we conducted Welch's independent samples t-tests on key tactical features (e.g., number of passes, number of shots, puck-carrying distance, and entries against), with the significance level set at 0.05. Welch's t-tests revealed that the majority of tactical features exhibited statistically significant differences between the identified playing styles ($p < 0.05$), with only a few indicators (e.g., dump-outs, entry success rate, and average carry distance for certain style pairs) showing no significant variation. This supports the robustness of the clustering results in differentiating teams in the tactical feature space.

4.2.2 Statistical Test of Key Findings

Before proceeding with model correction, we first conducted statistical tests on our two most central findings. First, to verify that the three playing styles identified in this study have explanatory power for match outcomes, rather than being arbitrary labels, we performed a **Pearson's Chi-squared test**. This test aims to determine whether an association exists between the two categorical variables: "combination of playing styles" and "match outcome (win/draw/loss)." The test yielded a highly significant p-value ($p < 0.0001$), leading to a strong rejection of the null hypothesis that "the two variables are independent." This fundamentally proves that our style clustering is effective and has a genuine statistical relationship with match results. Second, to address the most notable finding—the 82.14% win rate of the "Defensive Counterattack" style against the "Puck Control Play" style—we employed the **Bootstrap resampling method** to test whether this was a statistical artifact of a small sample. After 1000 resamples, we constructed a 95% confidence interval for this win rate, which was [**69.64%, 92.86%**]. This result is highly persuasive, as the lower bound of the interval (69.64%) is substantially higher than the 50% chance level, confirming that the strong tactical counter-relationship we observed is statistically robust.

4.2.3 Model Correction Considering Sample Imbalance

Although the aforementioned tests confirmed the validity of our core findings, a deeper examination of the data revealed a potential issue that could affect the fairness of the win-rate matrix: a significant imbalance in the sample sizes of style matchups. The data shows that while there were as many as 95 matches between "Defensive Counterattack" and "High-Pressure Offense," there were only 7 matches between "High-Pressure Offense" and "Puck Control Play." This imbalance poses a risk: a raw win rate calculated from only 7 matches (calculated to be 85.7%) could lead to misleading and extreme conclusions, as it is unlikely to represent the true long-term competitive relationship between the two styles. To systematically correct for this estimation bias caused by small samples, we introduced **Bayesian Averaging** to smooth the raw win rates. This method calculates a weighted win rate using the following formula:

$$\text{Weighted Rate} = \frac{C \cdot \text{prior\_mean} + N \cdot \text{original\_rate}}{C + N} \tag{1}$$

Here, N is the actual number of matches, original_rate is the raw win rate, prior_mean is the prior average we set (0.5), and C is the credibility constant, representing the strength of our confidence in the prior (dynamically calculated in this study based on the mean number of matchups). This method "pulls" the win rates from small samples toward the more credible 50% baseline, thereby generating a more robust matchup matrix. The effect of this correction is significant: the win rate of "High-Pressure Offense" against "Puck Control Play" was effectively adjusted from the raw 85.7% to a more conservative 55.0%, while the win rates of matchups with ample samples remained largely unaffected. The final win-rate heatmap presented in our study (as Figure 6) is based on this corrected and more equitable weighted win-rate matrix.

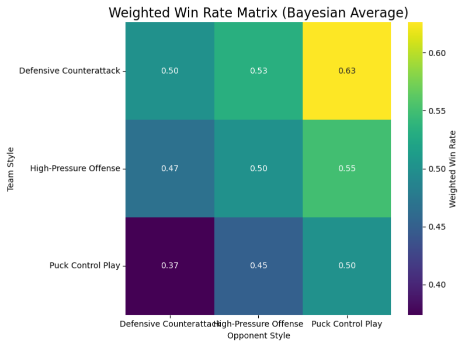

Fig. 3. Weighted Win Rate Matrix (Bayesian Average)

4.2.4 Cross-Validation of Model Generalization Ability

Finally, to test the ultimate stability and generalization ability of the corrected matchup model, we implemented a **70/30 cross-validatio**. We randomly split the entire dataset into a training set (70%) and a test set (30%) and independently constructed weighted win-rate matrices on these two subsets. The comparison revealed that the two matrices exhibit a highly consistent pattern. For instance, the advantageous position of "Defensive Counterattack" and the disadvantageous position of "Puck Control Play" were stably reproduced in both the training set and the "unseen" test set. This series of validation, correction, and re-validation steps forms a complete argumentative loop, ultimately proving that the playing styles identified in this study and their interactions are effective, robust, and possess good generalization ability.

4.3 Spatial Analysis of Possession Loss Locations

Despite the statistical corrections for sample imbalance, the Puck Control Play style's performance remains poor, with a low win rate that contrasts with its theoretical advantage in controlling game tempo. To further explore the defensive issues it may face in practice, the spatial distribution of average goals conceded by teams of each style during matches was analyzed. Through a heatmap visualization, the differences in goal-conceding areas across playing styles are revealed, providing a spatial perspective on the potential defensive weaknesses of the Puck Control Play style.

Figures 3 to 5 illustrate the spatial distribution of average puck loss locations for teams employing the three playing styles.

In Figure 3, the Average Possession Loss for Puck Control Play style teams reveals that possession posses are primarily concentrated in the defensive zone and near the neutral zone blue line. This suggests instability in puck control during defensive-to-offensive transitions or when organizing plays through the neutral zone. Such spatial patterns may expose Puck Control Play style teams to higher risks of turnovers under aggressive forechecking, thereby creating counterattack opportunities for opponents and negatively impacting overall game outcomes.

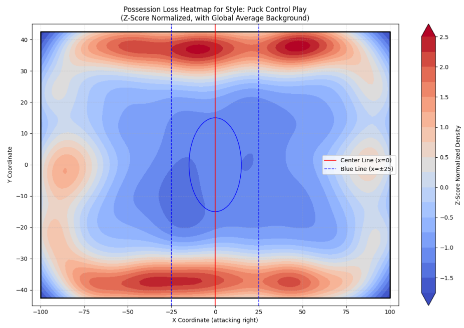

Fig. 4. Average Possession Loss Heatmap for Style: Puck Control Play

Figure 4 illustrates the Average Possession Loss for teams employing a Defensive Counterattack strategy. Compared to other styles, turnovers are more concentrated and occur farther from their own goal, primarily on both sides of the center red line. This pattern reflects the tactical focus on solid defensive positioning and swift counterattacks. Most puck losses take place during contested plays in the neutral zone and do not directly threaten the defensive zone, which may partly explain the higher win rate associated with this playing style.

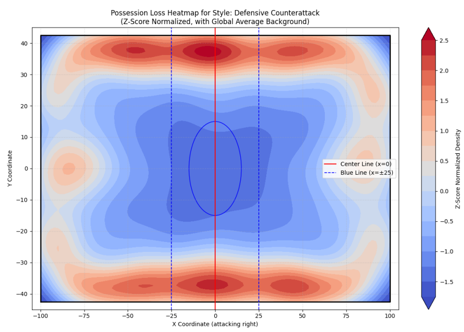

Fig. 5. Average Possession Loss Heatmap for Style: Defensive Counterattack

Figure 5 presents the Average Possession Loss for teams employing a high-pressure offensive strategy. Turnovers are concentrated in the offensive zone and around the offensive blue line. This pattern suggests that during aggressive forechecking and rapid transitions, possession may be lost due to rushed plays or passing errors, resulting in puck losses high up the ice. While this approach entails a higher turnover risk, the fact that these losses occur far from the team's own net reduces the immediate threat of conceding goals. This reflects a tactical trade-off between offensive intensity and tolerance for risk in puck management.

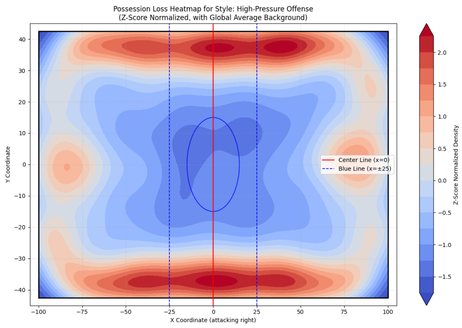

Fig. 6. Average Possession Loss Heatmap for Style: High-Pressure Offense

4.4 Game Theory-Based Analysis of Tactical Dynamics

To reveal the deep-seated interactions between different tactical styles from a theoretical perspective, we model a match as a strategic game. The players are the two opposing teams, the strategies are the three playing styles, and the payoff values are the win rates corrected by the Bayesian Averaging method. The resulting payoff matrix is as follows:

Table 1. Playing Style Game Payoff Matrix

Team A / Team B	Defensive Counterattack (DC)	High-Pressure Offense (HP)	Puck Control Play (PC)
Defensive Counterattack (DC)	(0.50, 0.50)	(0.53, 0.47)	(0.63, 0.37)
High-Pressure Offense (HP)	(0.47, 0.53)	(0.50, 0.50)	(0.55, 0.45)
Puck Control Play (PC)	(0.37, 0.63)	(0.45, 0.55)	(0.50, 0.50)

Analysis of this matrix reveals clear strategic dynamics: regardless of the opponent's strategy, "Defensive Counterattack (DC)" always yields the highest expected payoff, making it the unique **strictly dominant strategy** in this model. This finding leads directly to the game's only **pure-strategy Nash Equilibrium**, where both teams adopt the "Defensive Counterattack" strategy, resulting in a payoff of (0.50, 0.50) for both. Any unilateral deviation would lead to a reduced payoff. Although pure-strategy analysis points to "Defensive Counterattack" as the theoretical optimum, in practice, a **mixed strategy** led by this dominant strategy may have greater strategic value to avoid predictability. The game theory analysis provides profound theoretical insights into the league's tactical ecosystem. First, "Defensive Counterattack" is the core and optimal choice in the current tactical environment. Second, the (DC, DC) equilibrium point reveals a potential trend of "involution" or a conservative "meta" in the league, where adopting a conservative posture is the safest option amid uncertainty. Finally, this implies that for teams using "High-Pressure Offense" or "Puck Control Play" to succeed, they must possess superior tactical execution or player talent to compensate for their inherent disadvantage in macro-strategy.

5 Summary

This paper applied the K-Means clustering algorithm to identify three distinct playing styles from aggregated event data in SHL matches: "Puck Control Play", "Defensive Counterattack", and"High-Pressure Offense".Each style was clearly characterized through radar charts , highlighting their unique attributes in areas such as passing, shooting, dribbling, defense, and effectiveness. Preliminary analysis of style-versus-style win rates suggests that the "Defensive Counterattack" style achieved a higher win rate within this dataset. This paper provides a foundational framework for the quantitative understanding of playing styles in ice hockey.

6 Future Work

This paper serves as an initial exploration and could be extended in the following directions:

- **Methodological Refinement**: Determine the optimal number of clusters K using techniques such as the elbow method; adopt more precise definitions of match outcomes (e.g., final result); and account for stylistic dynamics across different phases of the game (e.g., by periods or overtime).
- **Incorporating Spatial Dimensions**:Integrate spatial information more thoroughly, such as computing metrics for specific pitch zones or applying spatial clustering techniques, to provide a richer description of playing styles [10].

- **Focus on Sequential Patterns**:To better align with the LINHAC competition task, future research should emphasize the **analysis of event sequences** that lead to key outcomes (e.g., goals, successful area entries) and examine how these patterns relate to the competition styles identified in this paper.
- **Model Selection**:Explore alternative clustering algorithms, such as fuzzy clustering methods that allow for partial membership across clusters [3], which may better capture the nuanced nature of playing styles.
- **Data Expansion and Validation**:Extend the analysis to additional seasons or other leagues to assess the generalizability and stability of the identified playing styles.

7 Code Access Link

The code used in this paper can be accessed here:
https://github.com/lyuuuuY/HockeyAnalysis.git

References

1. Leonardo Lamas, Junior Barrera, Guilherme Otranto, and Carlos Ugrinowitsch. Invasion team sports: strategy and match modeling. *International Journal of Performance Analysis in Sport*, 14(1):307–329, 2014.
2. Felicity Lord, David B Pyne, Marijke Welvaert, and Jocelyn K Mara. Identifying and analysing game styles and factors influencing a team's strategy in field hockey. *Journal of Sports Sciences*, 40(8):908–919, 2022.
3. Rasmus Säfvenberg, Niklas Carlsson, and Patrick Lambrix. Identifying player roles in ice hockey. In *International Workshop on Machine Learning and Data Mining for Sports Analytics*, pages 131–143. Springer, 2023.
4. Guiliang Liu, Wang Zhu, and Oliver Schulte. Interpreting deep sports analytics: Valuing actions and players in the nhl. In *International Workshop on Machine Learning and Data Mining for Sports Analytics*, pages 69–81. Springer, 2019.
5. Anton Olivestam, Axel Rosendahl, Erik Wilderoth, Niklas Carlsson, and Patrick Lambrix. Characterizing playing styles for ice hockey players. In *Linköping Hockey Analytics Conference*, pages 39–50, 2024.
6. Boyang Zhang and Tommy Löwendahl. Exploring the dynamics of ice hockey strategies using yolov8 and gephi in sports education. In *The European Conference on Education*, pages 817–826, 2023.
7. Oliver Schulte. Valuing actions and ranking hockey players with machine learning. In *Linköping Hockey Analytics Conference*, pages 2–9, 2022.
8. Felicity Lord, David B Pyne, Marijke Welvaert, and Jocelyn K Mara. Methods of performance analysis in team invasion sports: A systematic review. *Journal of sports sciences*, 38(20):2338–2349, 2020.
9. Christopher M Bishop and Nasser M Nasrabadi. *Pattern recognition and machine learning*, volume 4. Springer, 2006.
10. Claude B Vincent and Byron Eastman. Defining the style of play in the nhl: An application of cluster analysis. *Journal of Quantitative Analysis in Sports*, 5(1), 2009.

An LSTM-Based Approach to Predicting Zone Exit Success in Ice Hockey

Arun Ramji, Ethan Batiuk

University of Waterloo, Canada

Abstract. In ice hockey, successfully exiting the defensive zone with puck possession is a critical factor in determining game outcomes, as controlled exits have been shown to generate more subsequent offensive entries and shot attempts than uncontrolled dump-outs [1]. Failed zone exits and turnovers leave defenses vulnerable to counterattacks without defensive structure, increasing the likelihood of high-quality chances against them. However, limited research has explored the sequential dynamics of plays leading to successful exits versus turnovers. Additionally, current methods only credit the player completing the exit, overlooking contributions from other actions, such as puck battles or passes preceding the final exit play. This project addresses these gaps by developing a predictive model to identify play sequences that maximize successful zone exits, enabling us to quantify every action in the breakout sequence and provide actionable insights for team strategies and player evaluations.
Prior work, building on Corey Sznajder's tracking work has established the value of controlled zone exits and transition play [2]. We build upon this work by taking into account temporal dependencies in play sequences. Our approach employs a sequence-aware Long Short-Term Memory (LSTM) model, trained on event data from the SHL, to predict zone exit outcomes.

Keywords: Zone Exits, LSTM, Ice Hockey

1 Background

A zone exit occurs when the puck is moved out of the defensive zone. Controlled exits, achieved through passes or carries, are more effective than uncontrolled dump-outs because they produce more shots and goals [1]. Traditional approaches, such as zone exit tracking, emphasize the final exit event but overlook the preceding actions that shape the exit outcome. We analyzed event data from 156 Swedish Hockey League games, sourced from SportLogic. Events include passes, shots, dump-ins, with location, outcome, timestamps, and score state.

2 Algorithm

To model the sequential nature of zone exits in ice hockey, we selected an LSTM network [3], which captures temporal dynamics in play sequences. Zone exit attempts consist of events, puck retrieval, passes, dekes, and exit attempts, where success depends on their order and interaction. Unlike tree-based models that miss chronological dependencies, LSTMs inherently account for event order, handling variable-length sequences (2–30 events) effectively. The LSTM's gating mechanism prioritizes recent events while still considering earlier actions.

To prepare our data for training, we structured sequences to reflect play flow. Each sequence begins when the exiting team gains possession in the defensive zone. By default, we also include the two preceding events to provide context. Sequences were excluded if the possession ended in a faceoff or if they occurred in the final minutes of games to avoid score effects. Each sequence of events leading up to either an exit or a turnover was then tagged as controlled, uncontrolled, or failed. These processed sequences formed the dataset used to train our model. We represented categorical information using embeddings for event and type features, capturing similarities between them, and included five numeric features: spatial coordinates, score state, event outcome, and player position. Finally, our padding strategy with pack_padded_sequence ensured the LSTM ignored padded tokens, maintaining efficiency across varying sequence lengths.

3 Findings

3.1 Model Performance

The dataset was split into training, development and testing sets, with the test set comprising the following outcomes for zone exit attempts:

– **Successful Controlled Exits**: 17,502
– **Failed Exit Attempts**: 16,779
– **Successful Uncontrolled Exits**: 7,104

Evaluation protocol. We randomly split the dataset into training (70%), development (15%), and test (15%) subsets using stratified sampling on the three outcome classes (controlled, uncontrolled, failed) with a fixed random seed for reproducibility. The split was performed at the sequence level rather than grouped by game, since each sequence is self-contained and labeled independently. No sequence appeared in more than one subset, and the test set was strictly held out from training and hyperparameter tuning.

Ground truth. Labels were derived from our preprocessing pipeline (Section 3), which tagged each sequence of events leading to an exit or turnover as a controlled, uncontrolled, or failed exit. These labels served as the ground truth against which predictions were evaluated.

Evaluation metrics. We assessed performance on the held-out test set only. Accuracy was defined as the proportion of correctly classified sequences across all three classes. Cross-entropy loss was computed as the average categorical log loss. ROC-AUC was calculated using a one-vs-rest approach and macro-averaged across classes. The Brier score measured the mean squared error between predicted probability vectors and one-hot ground truth labels. Calibration was assessed using reliability curves [4].

Table 1. Evaluation metrics for the LSTM-based zone exit prediction model, assessing its ability to predict successful and failed zone exits.

Metric	Value
Cross-Entropy Loss	0.4567
Accuracy	0.7818
ROC-AUC	0.9328
Brier Score	0.2846

Results reporting. Alongside overall metrics, we present the confusion matrix and per-class precision, recall, and F1-scores to highlight strengths and weaknesses across classes. All reported numbers come from the held-out test set, ensuring that performance reflects genuine generalization. The model achieved strong performance (Accuracy = 0.78, ROC-AUC = 0.93), with calibration curves showing close alignment between predicted probabilities and observed outcomes (Figure 1).

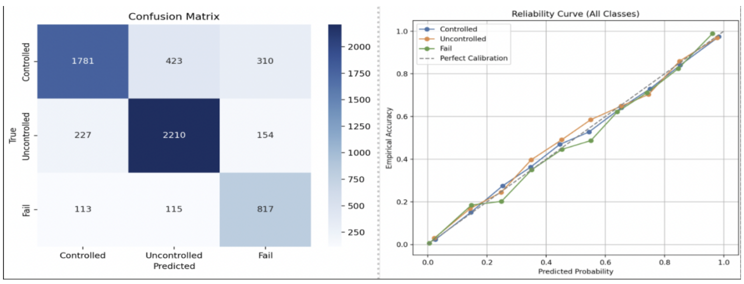

Fig. 1. Model performance evaluation. (Left) A confusion matrix comparing true labels (rows: Controlled, Uncontrolled, Fail) against predicted labels (columns), with values indicating the number of instances (e.g., 1,781 true controlled exits predicted correctly). (Right) Reliability curves for all classes (Controlled: blue, Uncontrolled: orange, Fail: green), plotting predicted probability (x-axis) against empirical accuracy (y-axis), with a dashed line representing perfect calibration.

Classification Metrics by Class. To further assess the model's performance across classes, we computed precision, recall, and F1-score for each outcome (controlled, uncontrolled, fail). Figure 2 shows that the model achieves balanced performance, with the "uncontrolled" class having the highest scores (precision: ∼0.82, recall: ∼0.84, F1-score: ∼0.83), followed by "controlled" and "fail." The slightly lower scores for "fail" (precision: ∼0.78, F1-score: ∼0.76) suggest that failed exits are more challenging to predict.

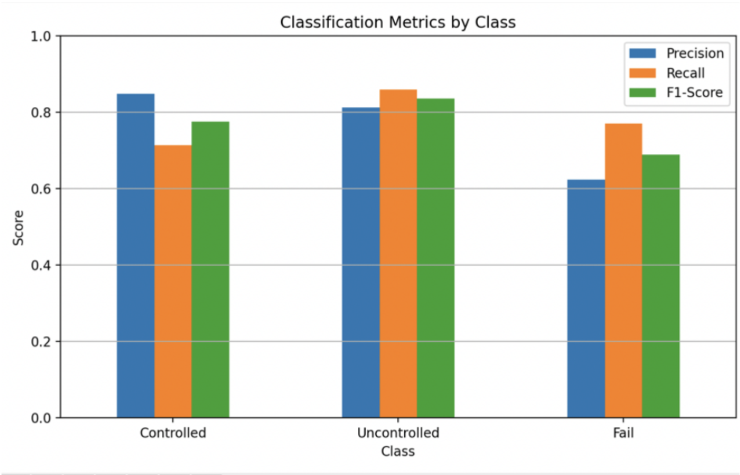

Fig. 2. Classification metrics by class. A bar chart displaying precision (blue), recall (orange), and F1-score (green) for each class (Controlled, Uncontrolled, Fail), with scores ranging from 0.0 to 1.0 on the y-axis. The "Uncontrolled" class shows the highest metrics, while "Fail" has the lowest, indicating prediction challenges.

3.2 Player Rating Metric

Our model produces, for each puck touch in an exit sequence, the predicted probabilities of three outcomes: controlled exit, uncontrolled exit, or failed exit. To translate these probabilities into a meaningful player rating, we draw on prior work by Chatel & Brosseau [1], which quantified the average number of shot attempts generated following different exit outcomes. Their detailed categories are collapsed into three coefficients: controlled (0.38), uncontrolled (0.16), and failed (0.07). These values represent the expected shot attempts associated with each outcome. For every puck touch, we compute the change in expected shot value (ΔEV) by multiplying the model's predicted probabilities of each outcome by their corresponding coefficients and summing across outcomes. Each player is credited with the ΔEV for each puck touch. Table 2 lists the top and bottom five players by EV, adjusted for touches, with full results available in *playerrating.py*.

Table 2. 5 lowest and 5 highest players for expected shots created.

Player ID (Lowest)	ΔEV	Player ID (Highest)	ΔEV
71713.0	-0.103	270437.0	0.022
316087.0	-0.102	548693.0	0.018
509158.0	-0.090	608705.0	0.016
464287.0	-0.087	561317.0	0.014
461057.0	-0.075	348173.0	0.012

We believe this model provides a stronger measure of player value in puck exits than traditional exit counts. To compare, we also applied our shot-expectancy framework in the traditional way, assigning the full coefficient to the player credited with the tagged exit (e.g., 0.38 for a controlled exit). Our model instead distributes credit across all touches by weighing each event according to predicted exit probabilities. The two approaches show a moderate positive correlation, suggesting our metric captures related information while also offering distinct insights.

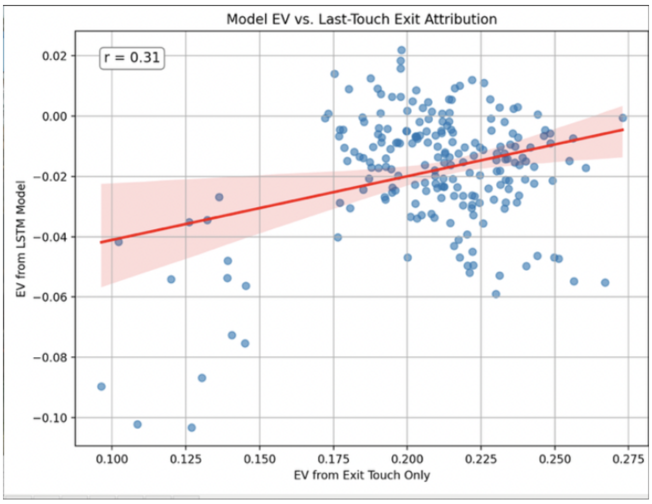

Fig. 3. Model EV Shift vs Last Touch EV Shift.

3.3 Event-Level Insights

Event-level analysis is another tool that can be used from this model, and it can be found in analysis.py. Our top takeaways are that when the team in the defensive zone has the puck, the only event that increases their expected shot value from an exit perspective is a pass. From an offensive team perspective, the worst thing you can do is dump the puck in and not be the first team to get possession, giving an average of 0.15 plus expected shot value to the opposing team. Additional takeaways and insights can be found by running analysis.py and reviewing the generated statistics.

3.4 Value Distribution Within Sequences

Using our shot-expectancy framework, we further examined when expected value (EV) gains occur during a zone exit sequence. By analyzing the distribution of EV changes across sequence positions, as shown in the table below, we demonstrate the importance of our model by highlighting how events early in the sequence contribute meaningfully to determining the outcome of a zone exit. Notably, nearly 30% of EV change occurs in the middle of sequences (50–60%), indicating that actions well before the final exit attempt often have the strongest influence on outcomes. Sequence position here refers to the percentage of events elapsed within a possession sequence (e.g., 50–60% corresponds to the middle actions). While the final actions (90–100%) still account for 17% of total EV change, the cumulative effect of earlier events is larger, underscoring the importance of modeling entire sequences rather than focusing solely on the last touch.

Table 3. : Distribution of Expected Value (EV) Changes Across Sequence Positions.

Sequence Position	% of Total EV Change
0-10%	+0.48%
10-20%	+0.79%
20-30%	+4.47%
30-40%	+11.46%
40-50%	+4.40%
50-60%	+29.81%
60-70%	+13.05%
70-80%	+11.50%
80-90%	+6.60%
90-100%	+17.43%

4 Summary of Key Ideas

This study introduces a novel approach to analyzing zone exits in ice hockey. Traditional metrics often focus solely on the final exit action, neglecting the sequential dynamics of preceding events, such as puck battles and passes. To address these gaps, the study employs an LSTM model, leveraging event data from

SHL games provided by SportLogic to predict zone exit success with high accuracy (0.7818) and ROC-AUC (0.9328). Additionally, a new player rating metric quantifies individual contributions to zone exits based on contributions of all plays leading to a zone exit. Our analysis further shows that the probabilities of exit outcomes (controlled, uncontrolled, failed) shift substantially throughout the sequence, with meaningful changes often occurring well before the final attempt.

5 Future Directions

Incorporating full player tracking data, capturing the location of every player on the ice, would significantly enhance the quality of the model and the depth of insight generated. This would allow for analysis of off-puck movement and could uncover valuable insights into areas such as which forecheck schemes most effectively reduce the odds of a successful breakout.

References

1. "Introducing Offensive Sequences and The Hockey Decision Tree," Hockey Graphs, Mar. 26, 2020. https://hockey-graphs.com/2020/03/26/introducing-offensive-sequences-and-the-hockey-decision-tree/. Accessed Aug. 1, 2025.
2.Novet, Alex. "Why Possession Is the Key to Zone Exits." Hockey Graphs, 30 July 2019, https://hockey-graphs.com/2019/07/30/why-possession-is-the-key-to-zone-exits/. Accessed Aug. 1, 2025.
3. Wikipedia Contributors, "Long short-term memory," Wikipedia, Nov. 22, 2018. https://en.wikipedia.org/wiki/Long_short-term_memory. Accessed Aug. 2, 2025.
4. Tanner, Gilbert. "Metrics." Machine Learning Explained, 2021, https://ml-explained.com/blog/metrics-explained. Accessed Aug. 1, 2025.
GIT: https://github.com/EthanAB99/LINHAC25

A Machine Learning Approach to Score Passes Based on Their Expected Offensive Contribution

Kiran Roye

Ossining High School, Ossining NY 10562, USA

1 Introduction

Passing plays a critical role in facilitating high-level ice hockey performance. Effective passing enables teams to transition from defense to offense, evade pressure, and generate scoring opportunities [1,2]. Traditional hockey analytics often quantify passing effectiveness based on assist counts or zone transitions, but these fail to isolate the true offensive value of the pass itself.

This study introduces a machine learning model that estimates the expected offensive value of a pass—specifically, the likelihood that a shot immediately following the reception will result in a goal. This value, which we refer to as **Expected Primary Assists (xPA)**, is designed to measure the contribution of a pass toward generating high-quality scoring chances. Importantly, xPA differs from traditional assist metrics in that it focuses exclusively on *primary assists*—the pass directly preceding the shot that leads to a goal.

2 Methodology

2.1 Data Processing

The dataset employed was a proprietary event dataset from Sportlogiq, encompassing 156 games from the Swedish Hockey League (SHL). The raw dataset contained 91,558 completed passes across all zones and game situations.

Each successful pass was identified from the event stream using the following criteria: event type labeled "pass" with outcome "successful", followed by a corresponding "reception" event with outcome "successful". Each pass-reception pair was assigned a randomly generated unique ID for tracking purposes.

Training Data Construction: For model training, we implemented a strict filtering process to focus on primary assist scenarios. For each reception, the algorithm searches forward chronologically through the event stream to identify the next action by the receiving player. A pass-reception pair was included in the training set only if:

1. The next event by the receiving player was a shot attempt
2. No change in possession occurred between reception and shot
3. No line change occurred between reception and shot

This filtering process yielded a training dataset of 8,247 pass-shot sequences. Each training example was labeled with the expected goals (xG) value of the shot immediately following the pass, representing the target xPA value.

Feature Engineering: The feature set was designed to capture both the spatial and temporal characteristics of each pass:

Continuous features (7 total):

- Pass origin coordinates: (x_{start}, y_{start})
- Reception coordinates: (x_{end}, y_{end})
- Temporal difference: dt (time between pass and reception)
- Pass distance: $\sqrt{(x_{end} - x_{start})^2 + (y_{end} - y_{start})^2}$
- Estimated pass velocity: distance $/ \ dt$

Categorical features (encoded as 9 binary columns):

- Reception type (6 categories): clean reception, deflection, rebound, etc.
- Manpower situation (3 categories): even strength, power play, penalty kill

All categorical variables were one-hot encoded using R's `model.matrix()` function, resulting in a final input matrix of 16 features.

2.2 Model Architecture and Training

A supervised learning approach was implemented using a feedforward neural network in R via the `keras` package. The model architecture was designed to handle the mixed continuous and categorical input features:

Architecture:

- Input layer: 16 features (7 continuous + 9 categorical binary)
- Normalization layer: z-score standardization of continuous features using training set statistics
- Hidden layers: 32 $\rightarrow$ 16 $\rightarrow$ 8 $\rightarrow$ 4 $\rightarrow$ 2 units, all with ReLU activation
- Output layer: single sigmoid unit producing xPA values in $[0, 1]$ range

Training Configuration:

- Optimizer: RMSProp (learning rate = 0.0005)
- Loss function: Mean Squared Error (MSE)
- Evaluation metric: Mean Absolute Error (MAE)
- Training epochs: 32
- Validation split: 10% (stratified sampling)
- Batch shuffling: enabled

2.3 Player-Level xPA Aggregation and Evaluation

After training, the model was applied to the complete dataset of 91,558 passes to generate xPA predictions for all completed passes, regardless of whether they resulted in shots. This application represents an important limitation discussed later, as the model is being applied to out-of-distribution data.

Player Aggregation Process: For each player who participated in at least 5 games, we computed:

- **xPA per game** $= \frac{\sum xPAfromallcompletedpasses}{gamesplayed}$
- **Primary assists per game** $= \frac{totalprimaryassists}{gamesplayed}$
- **Total assists per game** $= \frac{totalassists(primary+secondary)}{gamesplayed}$

The relationship between xPA per game and both primary assists per game and total assists per game was evaluated using Pearson correlation coefficients and linear regression analysis.

3 Results

3.1 Model Training Performance

The xPA model demonstrated strong convergence during training. Figure 1 shows the MSE and MAE metrics across 32 epochs.

In the final epoch, the model achieved:

- Training MSE: 0.002119, MAE: 0.02626
- Validation MSE: 0.001713, MAE: 0.02368

The validation metrics outperformed training metrics, suggesting effective generalization without overfitting. Both loss curves plateaued around epoch 20, indicating stable convergence.

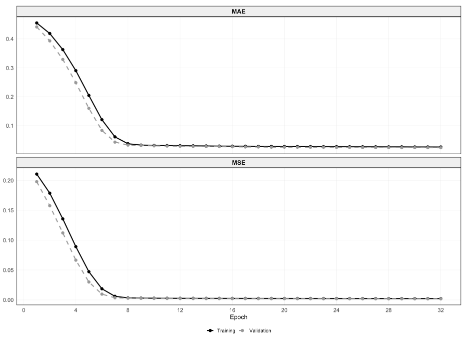

Fig. 1. Training and validation metrics for the xPA prediction model over 32 epochs. Solid black lines represent training metrics; dashed gray lines represent validation metrics. Convergence occurs by approximately 20 epochs.

3.2 Spatial Trends in xPA

A stratified sample of 2,000 passes was visualized on a scaled hockey rink (Figure 2). Higher xPA values were consistently assigned to passes originating from wide areas and received near the slot, aligning with established knowledge of high-danger shooting zones. Cross-ice seam passes and cycle plays in the offensive zone also generated elevated xPA scores.

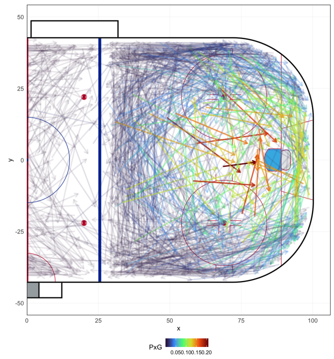

Fig. 2. Visualization of a random sample of passes. Each vector is colored based on its xPA value, with warmer colors indicating higher expected primary assist value.

3.3 Feature-xPA Relationships

Spearman rank correlations were calculated between xPA predictions and various pass characteristics across the full dataset:

- Pass velocity: $\rho = 0.21$ ($p < 0.001$)
- Longitudinal movement (endX - startX): $\rho = 0.081$ ($p < 0.001$)
- Lateral movement (|endY - startY|): $\rho = -0.097$ ($p < 0.001$)
- Pass angle relative to goal: $\rho = -0.013$ ($p < 0.001$)

While all associations achieved statistical significance due to the large sample size (n = 91,558), most correlations were weak in practical terms. Pass velocity demonstrated the strongest positive relationship with xPA, suggesting that quicker passes tend to create better scoring opportunities.

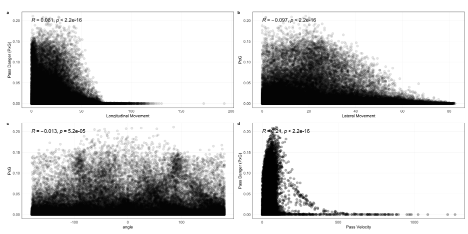

Fig. 3. Scatter plots showing Spearman correlations between xPA and different pass characteristics. Sample size limited to 5,000 random observations for visualization clarity.

3.4 xPA vs Assists Analysis

Linear regression analysis was conducted to evaluate the relationship between player-level xPA and traditional assist metrics. Results are presented for both comparisons:

xPA vs Total Assists:

- Pearson correlation: $r = 0.70$
- Linear regression: $R^2 = 0.49$
- Relationship: moderate positive correlation

xPA vs Primary Assists:

- Pearson correlation: $r = 0.78$
- Linear regression: $R^2 = 0.61$
- Relationship: moderately strong positive correlation

The stronger correlation with primary assists validates the model's focus on primary assist scenarios. However, both relationships show substantial unexplained variance, potentially due to factors such as teammate finishing ability, goaltender performance, and the distribution shift between training and application data.

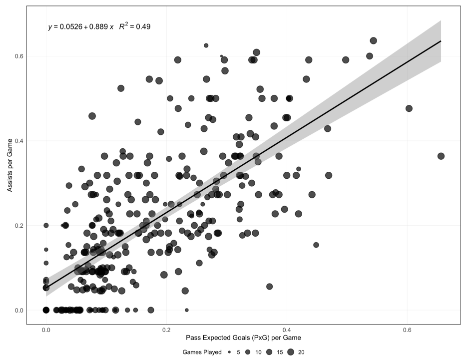

Fig. 4. Relationship between xPA per game and assists per game. Point size reflects games played; line represents the fitted linear regression. Left panel shows total assists, right panel shows primary assists only.

4 Discussion

This study introduces a method for evaluating the offensive contribution of individual passes using a supervised learning approach trained on event-level hockey data. The resulting xPA metric estimates the expected goal probability of a shot taken immediately following pass reception, providing a continuous measure of passing value.

Key Findings:

- The xPA model successfully identifies spatial patterns consistent with hockey tactical knowledge
- Player-level xPA correlates more strongly with primary assists ($r = 0.78$) than total assists ($r = 0.70$)
- Pass velocity emerges as the strongest individual predictor of pass value
- The model demonstrates stable training with good generalization properties

Applications: xPA offers several advantages over traditional assist-based metrics: it provides continuous value estimates for all completed passes, captures the quality rather than just quantity of passing plays, and can be applied for player evaluation, scouting, and tactical analysis even when passes don't result in shots.

Limitations and Future Work: Several important limitations constrain the current approach:

1. **Distribution Shift:** The model is trained exclusively on pass-shot sequences but applied to all completed passes. This creates a fundamental mismatch between training and application distributions that may limit predictive accuracy.
2. **Information Leakage:** Including reception coordinates as features may allow the model to implicitly learn shot location, potentially inflating performance metrics while reducing generalizability to truly novel passing scenarios.
3. **Data Limitations:** The analysis relies solely on event-level data. Incorporating player tracking data could capture additional context such as defensive pressure, player positioning, and ice surface congestion.
4. **Scope Limitations:** The current model focuses only on passes that directly precede shots. A more comprehensive approach might model the full sequence of events leading to scoring chances.

Future research should address these limitations through domain adaptation techniques to handle distribution shift, alternative feature representations that avoid information leakage, and incorporation of richer data sources including player tracking and video analysis.

5 Code Access Link

Code: github.com/KingKobra7899/KROYE_LINHAC_2025

References

1. MacKenzie, R., Cushion, C.: Passing and its role in ice hockey offense. International Journal of Sports Science & Coaching **15**(3), 345–357 (2020)
2. Schulte, O.: Offensive strategies in professional ice hockey. Journal of Quantitative Analysis in Sports **14**(4), 169–178 (2018)