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.

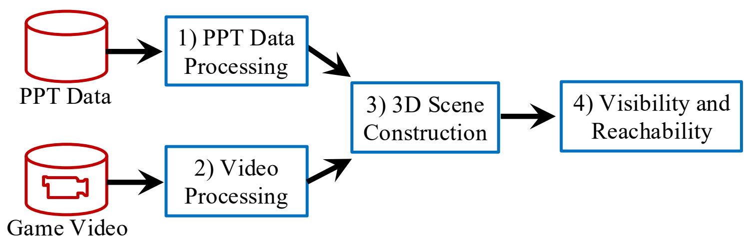

Fig. 1. The framework for computing net visibility and reachability.

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 856$^{\text{th}}$ 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.

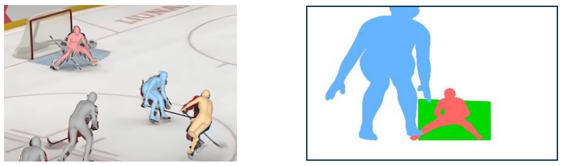

(a) (b)

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.

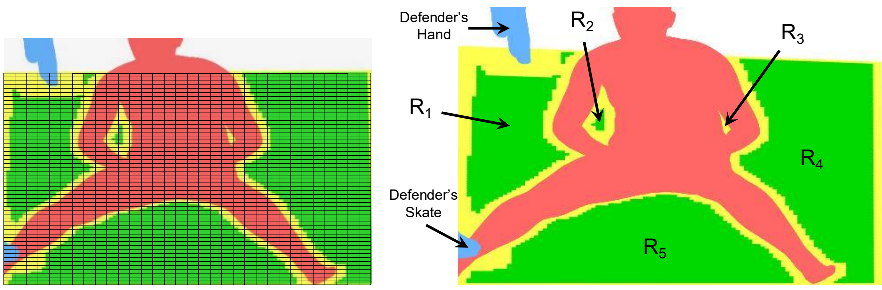

(a) Finding reachable portions (b) Visible and reachable regions

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.

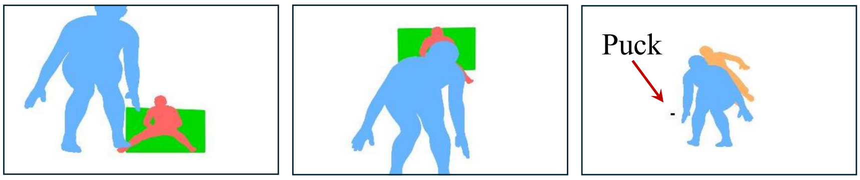

(a) Puck's view (b) Shooter's view (c) Goaltender's view

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.

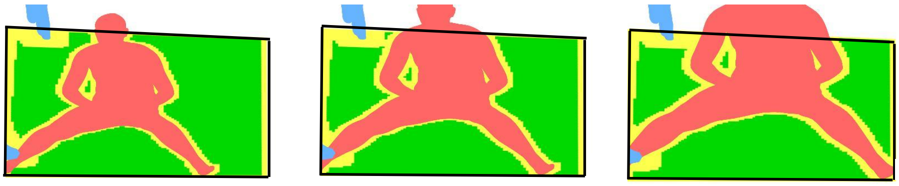

-4 feet	Original	+4 feet
V = 71 %	V = 66 %	V = 62 %
R = 57 %	R = 50 %	R = 45 %
D_{puck} = 34.3 ft	D_{puck} = 30.0 ft	D_{puck} = 25.7 ft
D_{net} = 2.4 ft	D_{net} = 6.5 ft	D_{net} = 10.9 ft

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.

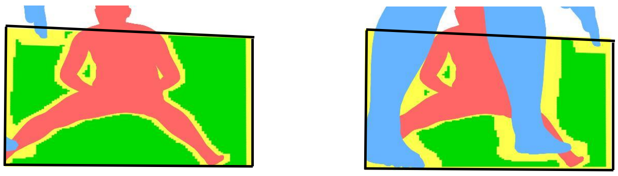

Original	D (#55) in shooting lane
V = 66 %	V = 46 %
R = 50 %	R = 26 %

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.