Object Detection for Ice Surface Localisation in Youth Hockey

Parisa Dehghan, Amirhossein Azadi, Allison Clouthier, T. Blaine Hoshizaki

Abstract Sports-related brain injury is a pressing issue, particularly in high-impact sports like ice hockey where impact velocity plays an important role in determining the magnitude of head impacts and subsequent risk of injury. However, existing methods for measuring impact velocity, such as GPS tracking and manual video analysis, are costly making them inaccessible, especially for youth leagues. This study introduces an automated, cost-effective method using computer vision to determine player velocity from 2D video. The initial step involves localising the field, achieved through a novel approach employing YOLOv5 to detect specific landmarks on the ice surface. With a dataset of over 9,900 annotated images, YOLOv5 demonstrates exceptional performance, achieving an F1 score and precision-recall of 0.99 at an 80% confidence level, and mAP scores of 98.5% and 64.5% at IoU thresholds of 0.5 and 0.5:0.95, respectively. By detecting at least four landmarks per frame, homography matrices were calculated to obtain a top-down view, completing the localisation process. This approach achieved an average IoU of 0.96, validating its accuracy in field localisation and demonstrating its potential for improving accessibility and cost-efficiency in measuring impact velocity in ice hockey.

Keywords Head impacts, homography matrix, ice hockey, localisation, object detection.

I. INTRODUCTION

For the most part, the majority of sports-related traumatic brain injuries (TBI) that contribute to mental health deficits and increased risk for neurodegenerative disease do not result in loss of consciousness or serious symptoms, especially in youth sport and remain under reported. Effective diagnosis, management and prevention requires accurate and reliable injury data. Automated data capture technology using video provides an opportunity to obtain biomechanical measures of head impact characteristics to measure brain tissue damage in sport. In ice hockey, the incidence of concussion is notably high, reaching up to 8.8 concussions per 100 athletic exposures [3]. Despite this, the actual incidence among youth hockey players remains uncertain due to ineffective monitoring of this demographic [4]. Hockey is recognised as one of the sports with the highest rates of head injury among youth. Athletes aged 5 to 18 years account for 65% of all sports and recreation-related head injuries treated in US emergency departments [5]. Repetitive head impacts (RHI), both high and low magnitude, continue to be associated with long-term brain health issues including depression, early cognitive decline, and chronic brain disease and chronic traumatic encephalopathy (CTE) [6-7].

Biomechanical measurements including impact velocity, location, mass, compliance, and direction can be used to obtain the dynamic response of the head and subsequently calculate potential brain tissue damage [9-11]. The magnitude and interaction of these parameters collectively influences brain motion during impact. Impact velocity is an important variable used to calculate the amount of energy transferred to the brain during impact and resulting injury severity [12]. Various methods, including GPS tracking devices [13-14], video analysis [15-16], time gates [17], model-based image matching (MBIM) [18], and Doppler effect devices [19], have been employed to measure velocity in sports. However, these methods are expensive and necessitate specialised technology, and few sports groups, especially youth, have access to such technology [20]. Manual video analysis is time-consuming and labor-intensive, making them impractical for use by sport organizations to monitor head trauma or creating the large databases needed for research. An objective, low-cost method for accurately documenting head impact

P. D. is a PhD student in Human Kinetics (pdehg058@uottawa.ca, +1 (343) 988 8234), A. A. is a PhD student in Human Kinetics, A.C. is an Assistant Professor of Biomedical Engineering, and T. B. H. is a Full Professor of Biomechanics in the Department of Human Kinetics, at the University of Ottawa in Ottawa, Canada.

characteristics during games, especially at the youth level is necessary. Automating data collection systems, facilitated by the application of computer vision techniques using 2D video, presents an opportunity to document head impact characteristics for a large number of sports and athletes. A challenging but important contributor to brain trauma magnitude involves the calculation of player velocity at the time of impact. Accurate estimation of player velocities requires detection, identification, and tracking of players to monitor their trajectories across the field [21]. Players' absolute velocities can be obtained using a top-down view of the ice surface to establish location. The vast majority of game video used in youth sport do not have a top-down view of the playing surface. However, a top-down view can be achieved using homography, a transformation matrix that maps each frame onto a reference plane. Localising the ice rink in each frame is essential for this calculation.

To calculate head impact trauma in sport the accuracy of a novel method for ice rink localisation using various angles from 2-d video to obtain player head impact velocities was investigated. Calculating player velocities is critical to automating head impact data capture in sports like ice hockey. Automating head trauma data collection from video provides for economical head impact documentation and the creation of a comprehensive datasets. Big data supports research investigating the associations between impact characteristics and clinical outcomes contributing to our understanding of brain trauma and injury.

This paper utilised game footage from the primary broadcast camera of youth games (peewee, midget, and Atom), positioned in the arena stands above the centre ice line. A novel method has been proposed in this paper for localising the ice surface, which uses an object detection algorithm to detect landmarks on the ice surface. Specifically, we employed YOLOv5[23] for landmark detection on the ice rink, aiming to identify at least four points to localise the rink in the top view frame[22]. This paper includes: a) Proposing a novel localisation method utilising YOLOv5 [23] for field landmark detection, b) introducing an ice rink dataset comprising over 9,000 images containing landmarks of the playing surface and c) leveraging computer vision techniques in developing an objective tool for measuring brain trauma in ice hockey.

The prevention of sports-related TBI is an important public health concern. Technologies such as deep neural networks and computer vision can facilitate the creation of extensive datasets to investigate the complex interplay between head trauma and mental health. Furthermore, it only requires a 2D video of the game and does not necessitate specialised equipment, making it accessible to youth participants who make up the majority of ice hockey players [5].

Using broadcast footage for computer vision tasks eliminates the need for additional specialised hardware to produce analytics. However, utilising broadcast footage poses its own unique challenges. Cameras used in broadcasts frequently pan, tilt, and zoom to follow the action, necessitating compensation for this movement in the analytics process. This becomes simpler when videos are captured from stationary cameras. Since hockey broadcast videos often lack camera parameters, determining location information must be automated from the video feed itself. Consequently, rink registration is conducted to determine how pixels in video frames correspond to the top-down view of the rink. The sports field localisation task involves determining the planar transform between the view of the hockey rink captured in the broadcast video frame and an overhead view of a hockey rink template [24-25]. This transform is defined by a homography matrix [22] between the two views of the plane of the ice surface. The homography matrix, H, is a 3 × 3 matrix with eight degrees of freedom, representing the transformation between two planes, typically scaled by a factor s. Equation 1 provides an example of a standard homography matrix, which establishes the correspondence between a point [x' y' 1]T on the rink model and a point [x y 1]T on the ice surface in the broadcast frame. The elements of H incorporate translation, rotation, and scale factors [26]. Our method dynamically derives the homography matrix for each video frame based on visible landmarks, allowing for accurate transformation to a top-down view. This frame-specific computation ensures precise localisation and alignment, which are essential for tasks like velocity calculation and player tracking.

$$S\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = H\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & 1 \end{bmatrix}\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \qquad \text{(Eq. 1)}$$

A standard method for computing the transformation between the plane of the ice surface in the broadcast video and the rink model involves detecting field markings in the frame, including points, lines, and line intersections, and matching them with their counterparts in the model. After the detection of multiple key points or landmarks in the frame, the RANSAC algorithm is employed. This involves the selection of a random subset of four points from these landmarks and the estimation of the homography matrix using them, typically through the Direct Linear Transform (DLT) algorithm [26]. The fit of all detected landmarks to this estimated homography is then evaluated, with well-fitting points being classified as inliers. This iterative process ensures robustness, with repetitions made multiple times using different sets of points. Ultimately, the homography with the highest number of inliers is chosen. Through the utilisation of RANSAC, outliers are effectively managed, resulting in the identification of the most accurate set of points for the homography and thereby rectifying perspective distortion [26]. An illustration of this process is depicted in Fig. 1, demonstrating how an image appears when transformed to the top-down view, as well as how the top-down template appears when transformed onto the image.

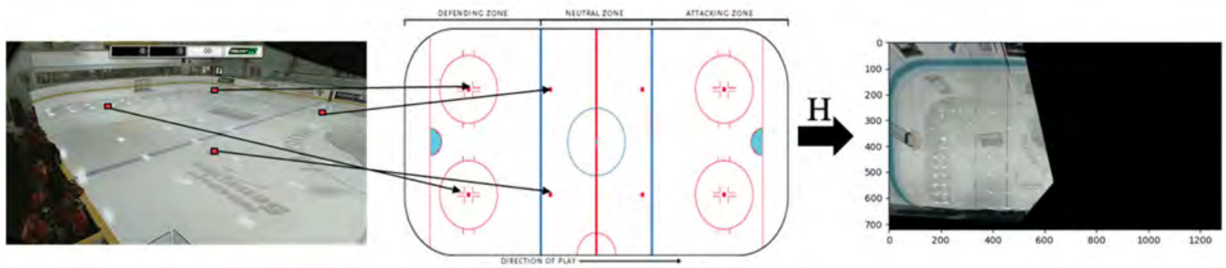

Fig. 1. Example of warping a video frame onto the overhead rink template (and vice versa) using homography.

Several computer vision techniques have proposed various methods for the localisation of ice surface. A method that combining point, line, and ellipse matches was used to estimate homography in ice hockey rinks, building upon the DLT algorithm [27]. A VGG16 semantic segmentation network to detect lines and estimate the camera position using branch and bound was also used [24]. Researchers have partitioned ice hockey footage into video shots and utilised a ResNet18-based regressor to determine homography between each frame and the ice-rink model [28]. Researchers have also employed the pix2pix network to extract lines from soccer broadcast video and compared them with a database of synthetic edge images for field localisation [29-30]. [31] Lines were identified and categorised on a soccer field [31], while a line segment detector to identify intersections and match them to a template was utilised [32]. Zones from soccer and basketball datasets were detected initialising camera pose estimation through a dictionary lookup and refining it with a spatial transformer network [33]. Reference [34-36] utilised Semantic segmentation or keypoint segmentation was used for sports field localisation [34-36], while a two-step method for refining homography estimates by combining the warped template and frame and minimising estimation error has been proposed [37].

The diversity of approaches published in sports field localisation underscores the absence of consensus on an optimal solution. This is compounded by the limited accessibility to algorithms and datasets from other research groups, resulting in each new contribution introducing a novel approach instead of building upon existing methods. Many projects achieve high accuracy using proprietary datasets, hindering meaningful comparisons. Despite introducing various methods, recent advancements suggest that deep network architectures offer superior performance with faster computation [38]. This study proposed a novel localisation method utilising YOLOv5 [23] as an object detection algorithm to identify landmarks on the ice surface, as detailed in the methodology section. This presents an important step for accurately determining player velocity.

II. METHODS

Data Collection

A dataset was created to support ice rink localisation techniques. This dataset comprised 9963 images capturing youth ice hockey games at various age levels, including Novice, Peewee, and Atom matches. The images were captured by a stationary camera positioned at the centre of the rink, which panned to capture all angles of the play area during recording. The dataset offered comprehensive views of the ice surface and was diversified to include frames where lines and pitches did not have perfect shapes. This diverse training data helped the model generalise better and handle variations effectively. Each image was manually annotated using annotation software named Roboflow [39] to highlight eight distinct landmarks within the ice rink, treating each as a separate object of interest. These landmarks encompass critical points such as the intersection of blue lines, goal lines, and centre lines with the boards, the intersection of goal posts with goal lines, the intersection of centre face-off circles with centre lines, and all face-off spots, as depicted in Fig.2. Well-defined guidelines were implemented to ensure accuracy, and training was provided for the labelers, significantly minimizing potential errors. Once the initial annotation was completed, another researcher reviewed a random sample of the annotations to check for inconsistencies. During the annotation process, markings with similar shapes were grouped together. In the programming section, it was described how each member of the same group was specified and categorised into individual classes based on their relative positions on the ice surface. Bounding boxes were drawn around each object during annotation to ensure accurate representation. The original images maintained an average size of 720 x 1280 pixels. The training process validation for this dataset is observed in Fig.3. This figure presents annotated image samples, describing the classes defined in Figure 2 as they appear in each frame.

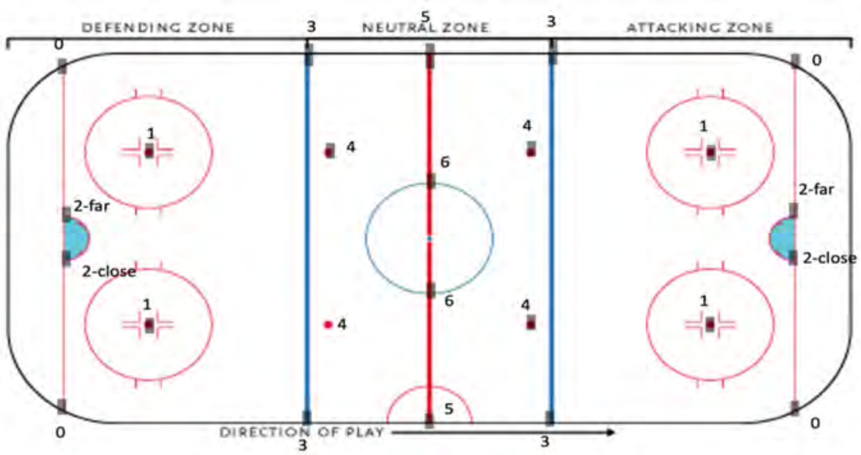

Fig.2. Top view of an ice rink with annotated landmarks labelled with their corresponding class names.

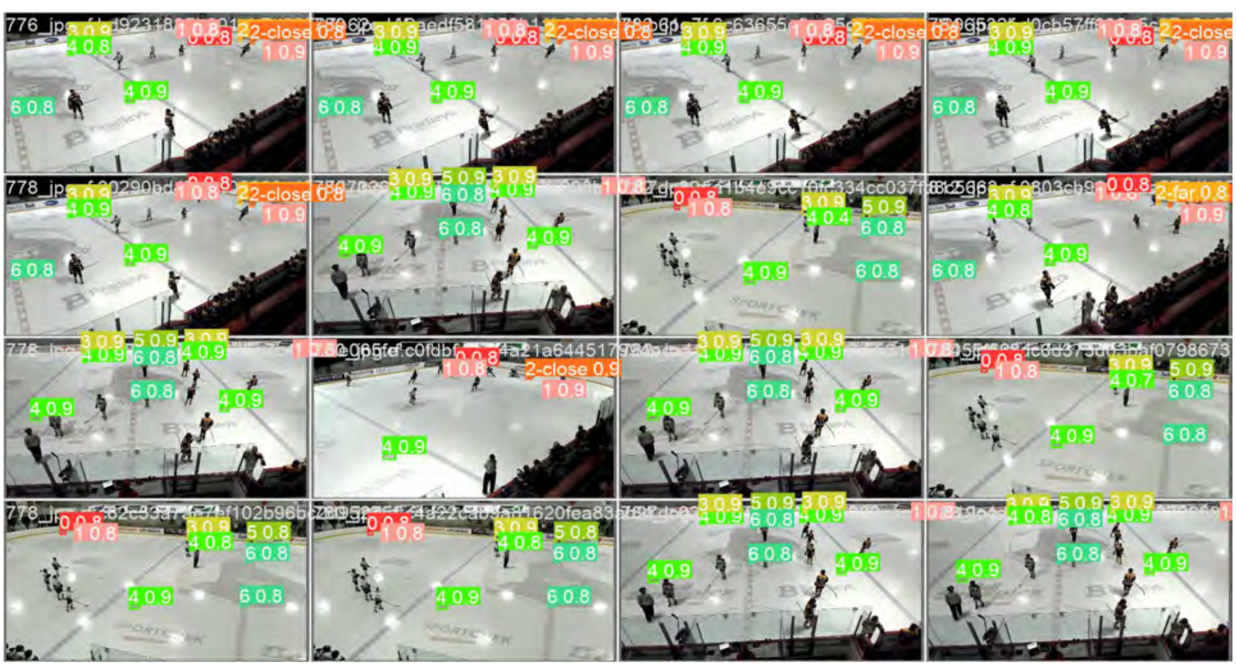

Fig.3. Ice surface's landmark dataset.

Data Preprocessing

The hockey images underwent preprocessing to prepare them for use in the YOLOv5 object detection algorithm. This process included resizing the images to 650 x 650 and applying filters to enhance contrast and remove noise. Subsequently, the dataset was augmented by rotating and flipping images, increasing the number of images to 14650. The images were randomly divided into two sets: training data (80%), and validation data (20%). The training data were used to train the machine learning model. The validation data were withheld during training and employed to provide an unbiased assessment of the model's performance throughout the training process. This information was used to fine-tune the model's hyperparameters.

Training the YOLOv5 model

The subsequent stage involved using a dataset to train the YOLOv5 model for ice hockey object detection. YOLO (You Only Look Once) is a prevalent deep learning-based algorithm for object detection known for its wide global receptive field, grid division, anchor frame matching, and multi-semantic fusion detection mechanism. Unlike traditional methods, which typically involve a multi-stage process of feature extraction followed by object identification and classification using separate algorithms like sliding window approaches or region-based methods such as R-CNN [40], YOLO directly predicts bounding boxes and object probabilities using Convolutional Neural Networks (CNNs) [41], thereby significantly improving detection accuracy. YOLOv5, built on the PyTorch [42] framework, boasts rapid detection speeds, reaching up to 140 frames per second. It incorporates the Mosaic data augmentation technique from YOLOv4, enhancing the detection of small targets by combining input images through random scaling, cropping, and arrangement. YOLOv5 has four variants, YOLOv5s, YOLOv5m, YOLOv5l, and YOLOv5x, each differing in network depth. Deeper networks yield more feature maps but also involve more complex computations. In this study, YOLOv5s pretraining weights were employed for training. The training process involved feeding images into the algorithm and adjusting neural network weights to improve its capability in detecting and classifying landmarks on the ice rink.

Experimental Environment

The system used Windows 10 Pro with PyTorch (1.8.0) with an AMD Ryzen 9 3900X 12-core processor@3.8 GHz with 32G memory CPU and a NVIDIA GeForce RTX 3090 GPU with CUDA (11.7).

Evaluation of the Object Detection Model

Once the YOLOv5 model had completed its training phase, it underwent evaluation using various metrics to assess its performance, including precision, recall, and F1-score. Additionally, the loss of validation and loss of training, representing the error between predicted and actual values during model training, were monitored, with lower values indicating better performance. These metrics served as indicators of the accuracy and efficiency of the model in detecting and classifying landmarks within the ice hockey images.

Precision, one of the essential components of the F1 score, measures the accuracy of positive predictions made by the model, indicating how many of the predicted landmarks are correct. Mean Average Precision (mAP), calculated at Intersection over Union (IOU) thresholds of 0.5 and 0.5-0.95, provided additional insight into the accuracy of predicted bounding boxes. Recall evaluates the model's ability to find all the relevant landmarks, indicating how many were successfully detected. The F1-score is a combined measure of precision and recall, providing a balanced assessment of the model's overall performance in detecting landmarks on the ice surface. The precision-recall curve (PR) is another metric that evaluates the trade-off between precision and recall at different confidence levels. These metrics collectively provided a comprehensive evaluation of the model's effectiveness in detecting landmarks on the ice surface.

In this paper the following terms were defined as: True positive (TP) refers to instances where the model correctly predicts the presence and location of a landmark within an image, with an IOU score exceeding 0.5. False positive (FP) occurs when the model incorrectly predicts the presence of a landmark instance within an image. This may happen if the model predicts a landmark where none exists or assigns a different label to the predicted landmark compared to the ground truth. False negative (FN) arises when the model erroneously predicts the absence of a landmark instance within an image with an IOU score below 0.5. This could occur if the model overlooks a landmark instance present in the ground truth or assigns a smaller bounding box to the predicted landmark compared to the ground truth.

$$\text{Recall} = \frac{TP}{TP+FN}, \qquad \text{(Eq. 2)}$$

$$\text{Precision} = \frac{TP}{TP+FP}, \qquad \text{(Eq. 3)}$$

$$\text{F1-score} = \frac{2 * \text{Precision} * \text{Recall}}{\text{Precision} + \text{Recall}}, \qquad \text{(Eq. 4)}$$

By meticulously examining the true positive, false positive, and false negative parameters and their associated scores, we can gain valuable insights into the model's performance. This process not only helped us identify areas for improvement but also allowed us to refine the model's capabilities for more accurate landmark detection and classification.

Programming

Once objects were detected across the frames, those containing at least four detected objects were selected. The centre point of each bounding box enclosing these objects was identified to represent the coordinates of the respective landmarks. In alignment with Fig. 2, landmarks with similar shapes were grouped into classes. A Python script was developed to distinguish between different objects within the same class. This differentiation was based on various factors, including the objects' relative coordinates to each other, their positions along the centre line, and their relative placements within the image boundaries, such as mid-top or mid-bottom. For each landmark on the ice surface, a subclass name was assigned, as illustrated in Fig. 4. After initial detection of each object by YOLOv5 in a frame, our programme specified the corresponding label for each subclass. Therefore, if we have class four detected in a frame four times, our programme determined which one belongs to 4-right-top, which one is 4-left-top, which one is 4-right-bottom, and which one is 4-left-bottom. It should be noted that several landmarks are the same in Figure 2 and 4 however in Figure 4 the positions of each landmark is identified.

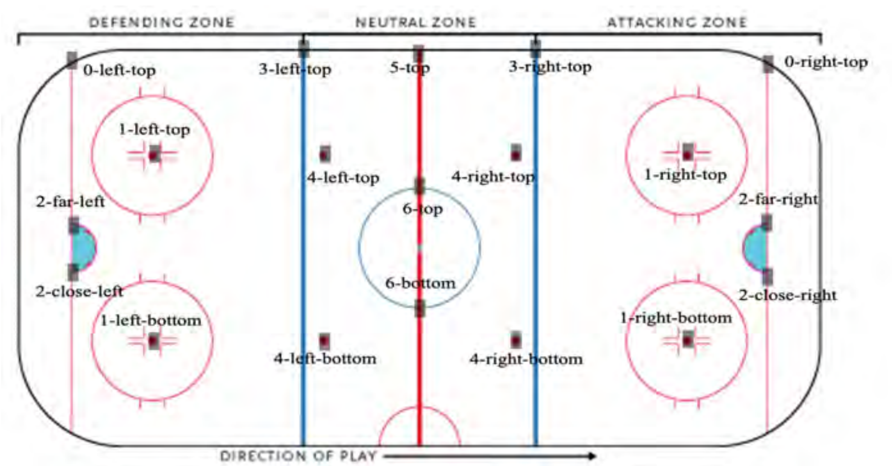

Fig .4. Reference frame including the landmarks with their true label according to their position.

A homography matrix was derived by accurately identifying correspondence between at least four points on a top-view reference image and the game video. This matrix determined the transformation to map the current view to the reference view. The top-view reference image had a size of 612 x 258 pixels. Application of this transformation matrix to each game video frame effectively transformed them into a top-view perspective, thus achieving localisation. Metrics that are used to evaluate sports field localisation are IOU_{part} and IOU_{whole} (Fig.5). IOU_{part} measures the overlap between the ground truth part (red line) and the predicted transformed region (yellow line) of the visible part of the field, while IOU_{whole} assesses the overlap for the entire field area after transformation. It evaluates the overall accuracy of the entire scene's transformation from the original perspective to the top-down view. By illustrating these metrics, Figure 5 helps clarify how effectively the transformation process preserves the spatial relationships and alignment of landmarks in the top-down view. In this paper, IOU_{part} was used to evaluate our localisation method.

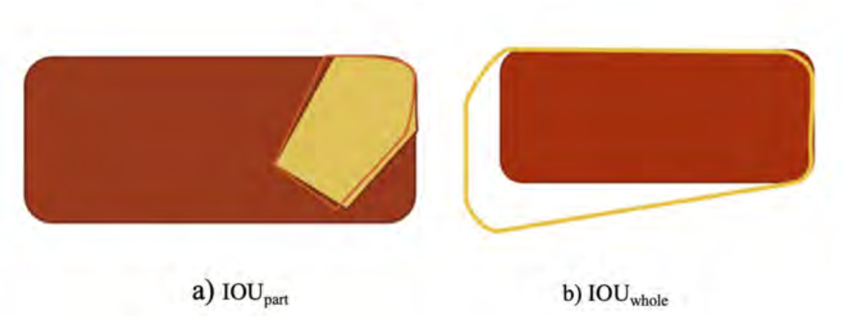

Fig. 5. Illustration of IOU Part versus Whole, highlighting the distinction in bounding box overlap assessment [43].

III. RESULTS

The number of epochs was set to 300. In Figure 6, the performance of YOLOv5 on landmark's detection is presented through various metrics during training and validation. The curves depict the loss associated with bounding box predictions (Train/Box Loss, Val/Box Loss), object presence predictions (Train/Obj Loss, Val/Obj Loss), and class predictions (Train/Class Loss, Val/Class Loss), reflecting the model's ability to accurately predict these parameters. Notably, in both training and validation, as epochs increased, the loss consistently decreased, highlighting the network's robust performance and effective learning process. Upon evaluation, the YOLOv5 model achieved a validation precision score of 0.98 and a recall score of 0.97. The mAP_0.5 (mean Average Precision at IoU 50%) and mAP_0.5:0.95 (Mean Average Precision at IoU 50% to 95% in increments of 0.05) of the landmark classification model were 98.5% and 64.5%, respectively. As illustrated in Fig. 6, as the number of epochs increased, the loss curve gradually stabilised, indicating improved model training effectiveness. These results confirmed the effectiveness of our approach in accurately predicting landmarks from frames taken at various camera angles. F1 and precision-recall curves, evaluating integrated Precision and Recall (PR) indices, are depicted in Fig. 7. It is evident that almost all eight classes performed very well, with F1 and PR evaluations achieving 0.99 at a confidence level of 80%. The spike observed in the Precision-Confidence Curve for the "2-close" class can occur due to fluctuations in data distribution or the model's adjustments, and it usually smooths out as training progresses. Overall accuracy and performance across all classes are the main concerns, and as long as the trend shows improvement and stability, occasional spikes were considered acceptable.

Homography matrices were calculated and applied to each frame by correlating the detected landmarks with their corresponding coordinates on the reference frame (see Fig. 2), enabling localisation. Subsequently, the IOU score was computed, yielding an average of 96.1. This high IOU score indicates the accuracy of the localisation process, affirming the method's effectiveness in precisely mapping objects to their respective landmarks on the ice surface. Figures 8 and 9 depict test images from two video frames of youth games and their corresponding warped frames in the top view, achieved through our network. Figures 8 and 9 demonstrate the accuracy of our localisation and transformation process by comparing the detected landmarks in the original image (left) with the transformed top-down view (right). Focusing on specific landmarks in the left image, they align precisely between the original and transformed images. For instance, objects in class 6 in the left image correspond with their locations in the top-down view on the right, overlaying correctly on the points at the top and bottom of the centre circle, indicating high precision. It is noteworthy that this transformation provides a top-down view of the frame, allowing us to calculate horizontal plane velocity, and does not include the vertical component.

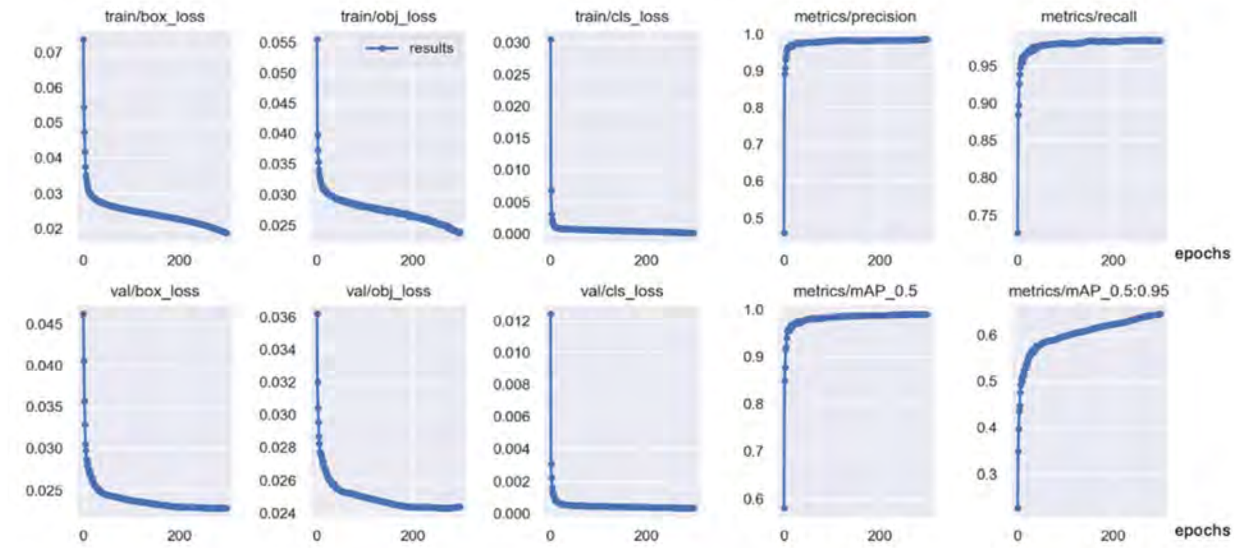

Fig. 6. Train and validation loss curves, mAP_0.5 and mAP_0.5:0.95 curves, precision curve, recall curve.

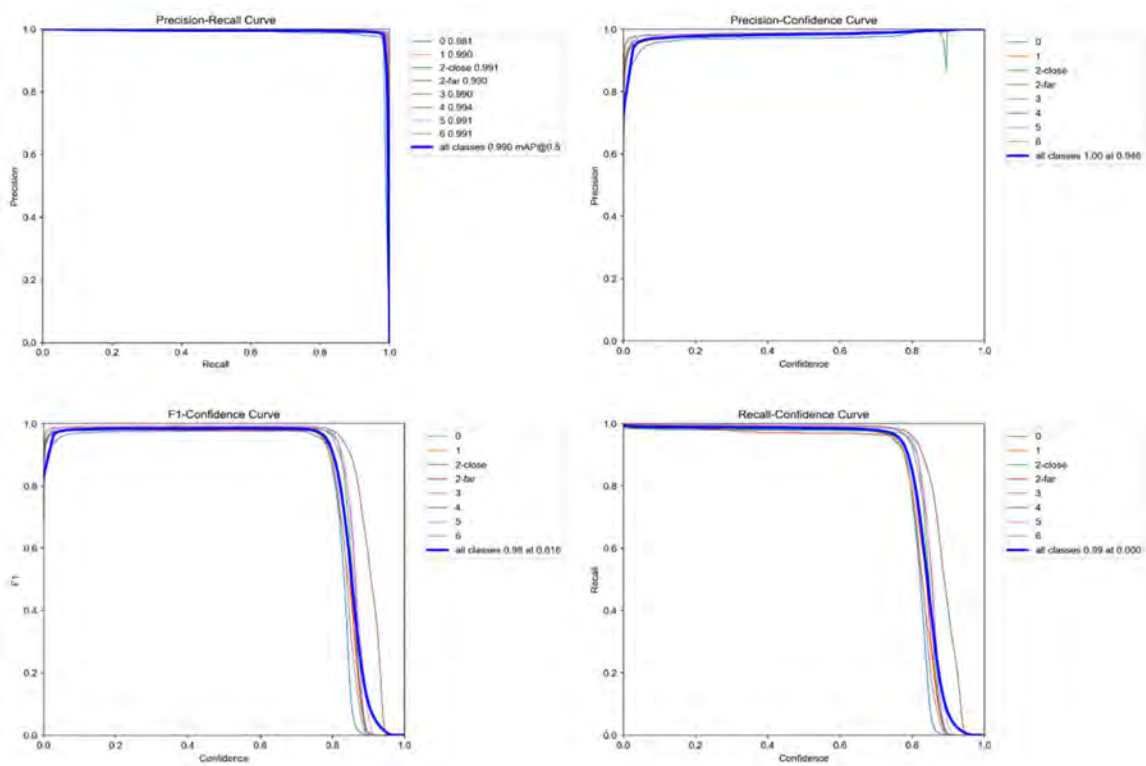

Fig. 7. Precision-Recall Curve, Precision-Confidence Curve, F1-Confidence Curve, Recall-Confidence Curve.

Fig. 8. Test image of a sample of a video frame including centre line.

Fig. 9. Test image of a sample of a video frame recording on right side of the rink.

IV. DISCUSSION

Medical science professionals [44-46] have identified head trauma especially in youth contact sports as a contributor to increases in mental health problems and risk of neurodegenerative disease. Youth, who represent the vast majority of athletes playing sporting activities have few financial or clinical resources to effectively monitor and manage head trauma. Impact velocity is an important characteristic of head impact severity with accurate impact velocities necessary to obtain accurate head impact severity measures. A top-view frame of players is essential for calculating accurate player plane velocity. This research investigated a novel method employing object detection techniques from various angles using 2d video for localising the ice surface in youth ice hockey games. This localisation enabled an accurate measurement of player velocity during head impacts, a critical factor in determining the magnitude of impact and assessing brain trauma. The entire rink was mapped by matching the visible landmarks from the current frame with those on the top-down view. Applying the homography matrix ensureed the whole rink was accurately aligned, as shown in Figures 8 and 9. This study yielded promising results, achieving an IoU score of 0.96 and localising youth ice hockey players without needing multiple cameras or specialised setups, which is described as excellent when compared to research in the field [24,29,37]. It is important to note that the homography transformation is generally more accurate near the reference points and might exhibit some error as we move further from the points of reference as a result the extrapolation needed in those areas, The high IoU score indicates that our overall transformation was accurate across the entire rink. The accuracy of the transformation process will be accessed using the player's coordinates from the top-down view to calculate their horizontal velocities and compared direct measures of player horizontal velocities.

The proposed method involves the application of biomechanical measures from 2d video to create large data sets supporting the investigation of the relationship between head trauma and neurodegenerative disease. The method employed to calculate impact velocities using 2d video support further development of an economic and efficient method to identify and document head trauma, creating large data sets to contribute to research designed to increase safety in sports medicine.

While our research has yielded promising results, it is essential to acknowledge its limitations. One such limitation is the calculation of homography, which requires a minimum of four points in the frame. Frames where fewer than four objects are detected are not applicable. However, in this study the cameras that record youth ice hockey games typically provide a long-shot view, capturing at least four landmarks in most frames. Our model was also trained on a sufficient number of images that accommodate most of the scenes and can capture enough landmarks. To enhance the accuracy of our method, more landmarks could be added to the object detection process, particularly in zoomed-in shots. This would ensure that the necessary points for homography calculation are generally available and could be a potential area for future research.

.

V. CONCLUSION

Accurate measures of impact velocity in sport head injuries is fundamental to establishing associated risk of brain injury. This research presents a novel method for localising the ice surface in youth ice hockey games employing object detection techniques from various angles using 2 d video. By leveraging computer vision technology, specifically YOLOv5 [23], we successfully detected landmarks on the ice surface, enabling accurate localisation without requiring specialised setups or multiple camera arrays. Upon evaluation, the YOLOv5 model achieved a validation precision score of 0.98 and a recall score of 0.97. The mAP (IOU[0.5]) and mAP (IOU [0.5: 0.95]) of the landmark classification model were 98.5% and 64.5%, respectively. As the number of epochs increased, the loss curve gradually stabilised, indicating improved model training effectiveness. These results support the effectiveness of our approach in accurately predicting landmarks from frames taken at various camera angles. F1 and precision-recall curves, evaluating integrated Precision and Recall (PR) indices, confirmed that almost all eight classes perform exceptionally well, with F1 and PR evaluations achieving 0.99 at a confidence level of 80%. These results validate the effectiveness of our approach in accurately predicting landmarks from frames taken at various camera angles.

Subsequently, the IOU score was computed, yielding an average of 96.1. This high IOU score indicates the accuracy of the localisation process, affirming the method's effectiveness in precisely mapping objects to their respective landmarks on the ice surface. The study achieved promising results, with a high IOU score of 0.96, indicating the accuracy of the localisation process.

The proposed method for obtaining impact velocities from various angles using 2 d video provides a promising approach to documenting, managing and studying head impact velocity and the risk of brain injury in youth sport. This study contributes to sports medicine by facilitating the measurement of player velocity during head impacts, crucial for assessing brain trauma opening avenues for further analysis of player performance, injury risk assessment, and trauma management in ice hockey.

VI. REFERENCES

[1] Gavett, B. E., Stern, R. A., & McKee, A. C. (2011). Chronic traumatic encephalopathy: a potential late effect of sport-related concussive and subconcussive head trauma. *Clinics in sports medicine*, *30*(1), 179-188.

[2] Daneshvar, D. H., Nowinski, C. J., McKee, A. C., & Cantu, R. C. (2011). The epidemiology of sport-related concussion. *Clinics in sports medicine*, *30*(1), 1-17.

[3] Donaldson, L., Asbridge, M., & Cusimano, M. D. (2013). Bodychecking rules and concussion in elite hockey. *PloS one*, *8*(7), e69122.

[4] Williamson, I. J. S., & Goodman, D. (2006). Converging evidence for the under-reporting of concussions in youth ice hockey. *British journal of sports medicine*, *40*(2), 128-132.

[5] Kontos, A. P., Elbin, R. J., Sufrinko, A., Dakan, S., Bookwalter, K., Price, A., ... & Collins, M. W. (2016). Incidence of concussion in youth ice hockey players. *Pediatrics*, *137*(2).

[6] Cubon, V. A., Putukian, M., Boyer, C., & Dettwiler, A. (2011). A diffusion tensor imaging study on the white matter skeleton in individuals with sports-related concussion. *Journal of neurotrauma*, *28*(2), 189-201.

[7] Bazarian, J. J., Zhu, T., Zhong, J., Janigro, D., Rozen, E., Roberts, A., ... & Blackman, E. G. (2014). Persistent, long-term cerebral white matter changes after sports-related repetitive head impacts. *PloS one*, *9*(4), e94734.

[8] Hoshizaki, B., Post, A., Kendall, M., Karton, C., & Brien, S. (2013). The relationship between head impact characteristics and brain trauma. *J Neurol Neurophysiol*, *5*(1), 1-8.

[9] Gennarelli, T. A., Thibault, L. E., Adams, J. H., Graham, D. I., Thompson, C. J., & Marcincin, R. P. (1982). Diffuse axonal injury and traumatic coma in the primate. *Annals of Neurology: Official Journal of the American Neurological Association and the Child Neurology Society*, *12*(6), 564-574.

[10] Kleiven, S. (2003). Influence of impact direction on the human head in prediction of subdural hematoma. *Journal of neurotrauma*, *20*(4), 365-379.

[11] Post, A., Hoshizaki, T. B., Gilchrist, M. D., Brien, S., Cusimano, M. D., & Marshall, S. (2014). The influence of dynamic response and brain deformation metrics on the occurrence of subdural hematoma in different regions of the brain. *Journal of neurosurgery*, *120*(2), 453-461.

[12] Ignacy, T., Post, A., Gardner, A. J., Gilchrist, M. D., & Hoshizaki, T. B. (2022). Comparison of dynamic response and maximum principal strain of diagnosed concussion in professional men's rugby league. *Proceedings of the Institution of Mechanical Engineers, Part P: Journal of Sports Engineering and Technology*, *236*(4), 266-276.

[13] Duffield, R., Reid, M., Baker, J., & Spratford, W. (2010). Accuracy and reliability of GPS devices for measurement of movement patterns in confined spaces for court-based sports. *Journal of science and medicine in sport*, *13*(5), 523-525.

[14] Aughey, R. J. (2011). Applications of GPS technologies to field sports. *International journal of sports physiology and performance*, *6*(3), 295-310.

[15] Castagna, C., Varley, M., Póvoas, S. C., & D'Ottavio, S. (2017). Evaluation of the match external load in soccer: Methods comparison. *International journal of sports physiology and performance*, *12*(4), 490-495.

[16] D'Apuzzo, N. (2002). Surface measurement and tracking of human body parts from multi-image video sequences. *ISPRS journal of Photogrammetry and Remote Sensing*, *56*(5-6), 360-375.

[17] Alexander, J. P., Hopkinson, T. L., Wundersitz, D. W., Serpell, B. G., Mara, J. K., & Ball, N. B. (2016). Validity of a wearable accelerometer device to measure average acceleration values during high-speed running. *The Journal of Strength & Conditioning Research*, *30*(11), 3007-3013.

[18] Tierney, G. J., Joodaki, H., Krosshaug, T., Forman, J. L., Crandall, J. R., & Simms, C. K. (2018). Assessment of model-based image-matching for future reconstruction of unhelmeted sport head impact kinematics. *Sports biomechanics*, *17*(1), 33-47.

[19] Hader, K., Mendez-Villanueva, A., Palazzi, D., Ahmaidi, S., & Buchheit, M. (2016). Metabolic power requirement of change of direction speed in young soccer players: not all is what it seems. *PloS one*, *11*(3), e0149839.

[20] Naik, B. T., Hashmi, M. F., & Bokde, N. D. (2022). A comprehensive review of computer vision in sports: Open issues, future trends and research directions. *Applied Sciences*, *12*(9), 4429.

[21] Barris, S., & Button, C. (2008). A review of vision-based motion analysis in sport. *Sports medicine*, *38*, 1025-1043.

[22] Dubrofsky, E. (2009). Homography estimation. *Diplomová práce. Vancouver: Univerzita Britské Kolumbie*, *5*.

[23] Redmon, J., Divvala, S., Girshick, R., & Farhadi, A. (2016). You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 779-788).

[24] jiangyounfar, N., Fidler, S., & Urtasun, R. (2017). Sports field localization via deep structured models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 5212-5220).

[25] Vats, K., Fani, M., Clausi, D. A., & Zelek, J. (2021). Puck localization and multi-task event recognition in broadcast hockey videos. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 4567-4575).

[26] Hartley, R., & Zisserman, A. (2003). *Multiple view geometry in computer vision*. Cambridge university press.

[27] Gupta, A. (2010). *Using line and ellipse features for rectification of broadcast hockey video* (Doctoral dissertation, University of British Columbia).

[28] Fani, M., Walters, P. B., Clausi, D. A., Zelek, J., & Wong, A. (2021). Localization of ice-rink for broadcast hockey videos. *arXiv preprint arXiv:2104.10847*.

[29] Sharma, R. A., Bhat, B., Gandhi, V., & Jawahar, C. V. (2018, March). Automated top view registration of broadcast football videos. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)* (pp. 305-313). IEEE.

[30] Chen, J., & Little, J. J. (2019). Sports camera calibration via synthetic data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops* (pp. 0-0).

[31] Cuevas, C., Quilon, D., & García, N. (2020). Automatic soccer field of play registration. *Pattern Recognition*, *103*, 107278.

[32] Tsurusaki, H., Nonaka, K., Watanabe, R., Konno, T., & Naito, S. (2021). Sports camera calibration using flexible intersection selection and refinement. *ITE Transactions on Media Technology and Applications*, *9*(1), 95-104.

[33] Sha, L., Hobbs, J., Felsen, P., Wei, X., Lucey, P., & Ganguly, S. (2020). End-to-end camera calibration for broadcast videos. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (pp. 13627-13636).

[34] Tarashima, S. (2020). SFLNet: direct sports field localization via CNN-based regression. In *Pattern Recognition: 5th Asian Conference, ACPR 2019, Auckland, New Zealand, November 26–29, 2019, Revised Selected Papers, Part I 5* (pp. 677-690). Springer International Publishing.

[35] Leonardo, C., Pablo, M. N., Stefano, S., Vivek, J., Dubout, C., Félix, R., ... & Pascal, F. (2020). Real-time camera pose estimation for sports fields. *Machine Vision and Applications*, *31*(3).

[36] Nie, X., Chen, S., & Hamid, R. (2021). A robust and efficient framework for sports-field registration. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision* (pp. 1936-1944).

[37] Jiang, W., Higuera, J. C. G., Angles, B., Sun, W., Javan, M., & Yi, K. M. (2020). Optimizing through learned errors for accurate sports field registration. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision* (pp. 201-210).

[38] Chen, J., Zhu, F., & Little, J. J. (2018, March). A two-point method for PTZ camera calibration in sports. In *2018 IEEE winter conference on applications of computer vision (WACV)*(pp. 287-295). IEEE.

[39] Alexandrova, S., Tatlock, Z., & Cakmak, M. (2015, May). RoboFlow: A flow-based visual programming language for mobile manipulation tasks. In *2015 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 5537-5544). IEEE.

[40] Ren, S., He, K., Girshick, R., & Sun, J. (2016). Faster R-CNN: Towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence*, *39*(6), 1137-1149.

[41] Chua, L. O. (1998). *CNN: A paradigm for complexity* (Vol. 31). World Scientific.

[42] Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., ... & Chintala, S. (2019). Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, *32*.

[43] Walters, P. (2021). *Towards Efficient Ice Surface Localization From Hockey Broadcast Video* (Master's thesis, University of Waterloo).

[44] McAllister, T., & McCrea, M. (2017). Long-term cognitive and neuropsychiatric consequences of repetitive concussion and head-impact exposure. *Journal of athletic training*, *52*(3), 309-317.

[45] Zetterberg, H., Winblad, B., Bernick, C., Yaffe, K., Majdan, M., Johansson, G., ... & Blennow, K. (2019). Head trauma in sports–clinical characteristics, epidemiology and biomarkers. *Journal of internal medicine*, *285*(6), 624-634.

[46] McKee, A. C., Mez, J., Abdolmohammadi, B., Butler, M., Huber, B. R., Uretsky, M., ... & Alosco, M. L. (2023). Neuropathologic and clinical findings in young contact sport athletes exposed to repetitive head impacts. *JAMA neurology*, *80*(10), 1037-1050.