

# Evaluating deep tracking models for player tracking in broadcast ice hockey video

Kanav Vats, Mehrnaz Fani, David A. Clausi, and John S. Zelek

Department of Systems Design Engineering, University of Waterloo  
{k2vats, mfani, dclausi, jzelek}@uwaterloo.ca

**Abstract.** Tracking and identifying players is an important problem in computer vision based ice hockey analytics. Player tracking is a challenging problem since the motion of players in hockey is fast-paced and non-linear. There is also significant player-player and player-board occlusion, camera panning and zooming in hockey broadcast video. Prior published research perform player tracking with the help of handcrafted features for player detection and re-identification. Although commercial solutions for hockey player tracking exist, to the best of our knowledge, no network architectures used, training data or performance metrics are publicly reported. There is currently no published work for hockey player tracking making use of the recent advancements in deep learning while also reporting the current accuracy metrics used in literature. Therefore, in this paper we compare and contrast several state-of-the-art tracking algorithms and analyze their performance and failure modes in ice hockey.

**Keywords:** ice hockey · deep learning · tracking.

## 1 Introduction

Ice hockey is played by an estimated 1.8 million people worldwide [10]. As a team sport, the positioning of the players and puck on the ice are critical to team offensive and defensive strategy [22]. The location of players on the ice is essential for hockey analysts for determining the location of play and analyzing game strategy and events. In ice hockey, prior published research [15, 5] perform player tracking with the help of handcrafted features for player detection and re-identification. Okuma *et al.* [15] track hockey players by introducing a particle filter combined with mixture particle filter (MPF) framework [23], along with an Adaboost [24] player detector. The MPF framework [23] allows the particle filter framework to handle multi-modality by modelling the posterior state distributions of  $M$  objects as an  $M$  component mixture. A disadvantage of the MPF framework is that the particles merge and split in the process and leads to loss of identities. Moreover, the algorithm does not have any mechanism to prevent identity switches and lost identities of players after occlusions. Cai *et al.* [5] improve upon [15] by using a bipartite matching for associating observations with targets instead of using the mixture particle filter framework. However, the

algorithm is not trained or tested on broadcast videos, but performs tracking in the rink coordinate system after a manual homography calculation.

Remarking that there is a lack of publicly available research for tracking ice hockey players making use of recent advancements in deep learning, in this paper we track and identify hockey players in broadcast NHL videos and analyze performance of several state-of-the-art deep tracking models on the ice hockey dataset. We also annotate and introduce a new hockey player tracking dataset on which the deep tracking models are tested.

## 2 Related work

There are a number of recent studies dealing with player tracking in basketball [19, 13, 27] and soccer [20, 9, 21, 7]. For basketball player tracking, Sangüesa *et al.* [19] demonstrated that deep features perform better than classical handcrafted features for basketball player tracking. Lu *et al.* [13] perform player tracking in basketball using a Kalman filter by making the assumption that the relationship between time and player's locations is approximately linear in a short time interval. Zhang *et al.* [27] perform basketball player tracking in a multi camera setting.

In soccer, Theagarajan *et al.* [20] track players using the deep SORT algorithm [26] for generating tactical analysis and ball possession statistics. Hurault *et al.* [9] introduce a self-supervised detection algorithm to detect small soccer players and track players in non-broadcast settings using a triplet loss trained re-identification mechanism, with embeddings obtained from the detector itself. Theiner *et al.* [21] present a pipeline to extract player position data on the soccer field from video. The player tracking was performed with the help of CenterTrack [29]. However, the major focus of the work was on detection accuracy rather than tracking and identification. Gadde *et al.* [7] use a weakly supervised transductive approach for player detection in soccer broadcast videos by treating player detection as a domain adaptation problem. The dataset used is generated with the help of the deep SORT algorithm [26].

## 3 Methodology

We experimented with five state-of-the-art tracking algorithms [3, 26, 28, 1, 4] on the hockey player tracking dataset. The algorithms include four online tracking algorithms [3, 26, 28, 1] and one offline tracking algorithm [4]. SORT [3], deep SORT [26] and MOT Neural Solver [4] are tracking by detection (TBD) algorithms. Tractor [1] and FairMOT [28] are joint detection and tracking (JDT) algorithms.

Tracking by detection (TBD) is a widely used approach for multi-object tracking. TBD consists of three steps: (1) detecting objects (hockey players in our case) frame-by-frame in the video (2) calculating affinity between detected objects (3) inference - linking player detections using calculated affinities to produce tracks. Concretely, in TBD, the input is a set of object detections

**Table 1.** Tracking algorithms compared for hockey player tracking.

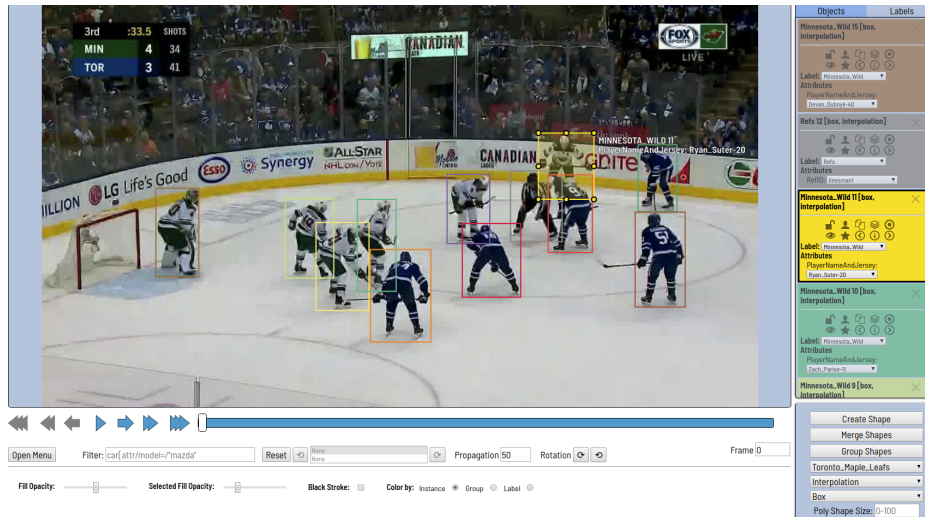
Algorithm	Description
SORT [3]	Kalman filter with simple IOU based re-id.
Deep SORT [26]	Kalman filter with deep CNN based re-id.
Tracktor [1]	JDT algorithm with separate detection and re-id networks.
FairMOT [28]	JDT algorithm with combined object detection and re-id network.
MOT Neural Solver [4]	Tracking using graph message passing with edge classification.

$O = \{o_1, \dots, o_n\}$ , where  $n$  denotes the total number of detections in all video frames. A detection  $o_i$  is represented by  $\{x_i, y_i, w_i, h_i, I_i, t_i\}$ , where  $x_i, y_i, w_i, h_i$  denotes the coordinates, width, and height of the detection bounding box.  $I_i$  and  $t_i$  represent the image pixels and timestamp corresponding to the detection. Affinity calculation consists of calculating affinity between detections  $o_i$  by obtaining appropriate features. The features can be simple intersection over union (IOU) based [3] or using deep networks [25]. After affinity calculation, a set of trajectories  $T = \{T_1, T_2 \dots T_m\}$  is found that best explains  $O$  where each  $T_i$  is a time-ordered set of observations. This is done through an appropriate inference technique. Two widely used inference techniques are filtering [3, 25] and graphical formulation [4]. As an example of graphical formulation, the MOT Neural Solver [4] models the tracking problem as an undirected graph  $G = (V, E)$ , where  $V = \{1, 2, \dots, n\}$  is the set of  $n$  nodes for  $n$  player detections for all video frames. In the edge set  $E$ , every pair of detections is connected so that trajectories with missed detections can be recovered. The problem of tracking is posed as splitting the graph into disconnected components where each component is a trajectory  $T_i$ . After computing each node embedding and edge embedding using a CNN (affinity calculation), the model then solves a graph message passing problem. The message passing algorithm classifies whether an edge between two nodes in the graph belongs to the same player trajectory.

Joint detection and tracking (JDT) [1, 28] is the latest trend in multi-object tracking research. These methods either (1) Convert an object detector to a tracker by estimating the location of a bounding box in the adjacent frames [1] or (2) Perform detection and re-identification using a single network [28]. Bergmann *et al.* [1] use the bounding box regressor of a Faster RCNN [16] detector to regress the position of a person in the next frame. The re-identification is performed using a separate siamese network. Zhang *et al.* [28] perform object detection and re-identification with the same network using separate detection and re-identification branches. The differences and similarities between the five tracking algorithms are summarized in Table 1. We refer the readers to the publications of the respective tracking papers [3, 26, 28, 1, 4] for more detail.

## 4 Dataset

The player tracking dataset consists of a total of 84 broadcast NHL game clips with a frame rate of 30 frames per second (fps) and resolution of  $1280 \times 720$  pixels. The average clip duration is 36 seconds. The 84 video clips in the dataset

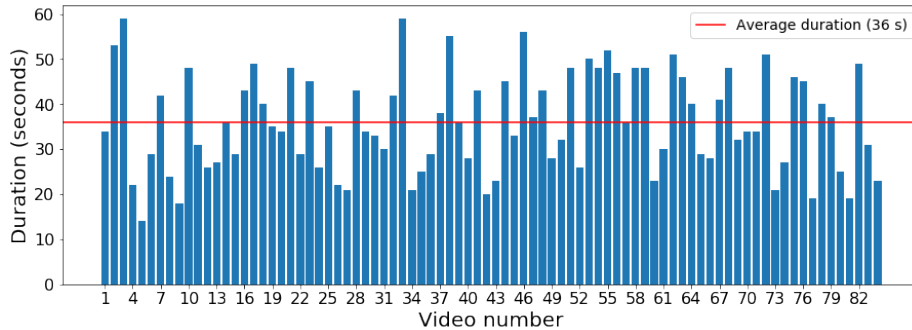


**Fig. 1.** CVAT tool used for tracking annotations. The tool offers the ability to annotate bounding boxes with each box having one label - home or away team. Each player bounding box has player name and jersey number as attributes. CVAT also offers an interpolation mode which alleviates the need to draw bounding boxes multiple times for adjacent frames.

are extracted from 25 NHL games. The duration of the clips is shown in Fig. 2. Each video frame in a clip is annotated with player and referee bounding boxes and player identity consisting of player name and jersey number. The annotation is carried out with the help of the open source computer vision annotation tool (CVAT) <sup>1</sup>. An illustration of an annotation job using the CVAT tool is shown in Fig. 1. The dataset is split such that 58 clips are used for training, 13 clips for validation, and 13 clips for testing. To prevent any game-level bias affecting the results, the split is made at the game level, such that the training clips are obtained from 17 games, validation clips from 4 games and test split from 4 games respectively.

Table 2 compares the size of the dataset with other tracking datasets in literature. The hockey player tracking dataset is comparable in size with other tracking datasets used in literature. As compared to pedestrian datasets (MOT 16 [14] and MOT20 [6]), the bounding boxes per frame is less in our dataset since the maximum number of players on the screen can be 12, with usually less than 12 players actually in broadcast camera field of view (FOV). The NHL game videos used to create this dataset have been obtained from Stathletes Inc. with permission.

<sup>1</sup> Found online at: <https://github.com/openvinotoolkit/cvat>



**Fig. 2.** Duration of videos in the player tracking dataset. The average clip duration is 36 seconds. The red horizontal line represents the average clip duration.

**Table 2.** Comparison of hockey tracking dataset with other tracking datasets in literature. Our hockey player tracking dataset is comparable to other multi-object tracking datasets commonly used in literature.

Dataset	Videos/sequences	Frames	Bounding boxes	Domain
MOT16 [14]	14	11, 235	292, 733	Pedestrians
MOT20 [6]	8	13, 410	2, 102, 385	Crowded pedestrian scenes
KITTI-T [8]	50	10, 870	65, 213	Autonomous driving
Ours	84	91, 807	773, 545	Ice hockey players

#### 4.1 Annotation process

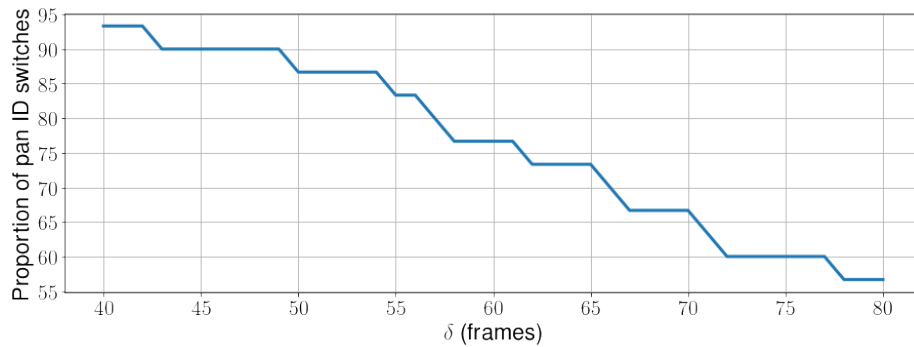
15 annotators annotated the whole dataset using the CVAT tool. The average time taken to annotate one minute of video is 10.45 minutes. The total time taken to annotate all 84 videos is 527 minutes. The manual annotation was done such that a bounding box as tight as possible was drawn around a player/referee. Linear interpolation was used to interpolate bounding box positions. Additionally, unlike other tracking datasets such as MOT16 [14] and MOT20 [6], the same ground truth identity was assigned to a player leaving camera FOV at a particular frame and re-entering after some time. If a player was occluded by board or another player, the bounding box was annotated based on the best guess of the tightest box enclosing the full body of the player. For quality control, all bounding boxes were checked to make sure each box has label-name(name of the player). When a player enters/exits the scene, his bounding box was labeled even if he was partially in camera FOV. Whenever players were occluded by other players, revision of annotations was performed to ensure high quality.

## 5 Results

Player detection is performed using a Faster-RCNN network [17] with a ResNet50 based Feature Pyramid Network (FPN) backbone [11] pre-trained on the COCO

**Table 3.** Comparison of the overall tracking performance on test videos of the hockey player tracking dataset. ( $\downarrow$  means lower is better,  $\uparrow$  mean higher is better)

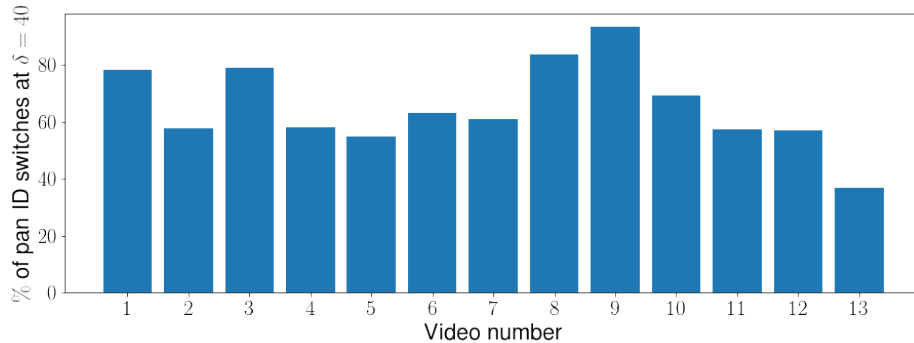
Method	IDF1 $\uparrow$	MOTA $\uparrow$	ID-switches $\downarrow$	False positives (FP) $\downarrow$	False negatives (FN) $\downarrow$
SORT [3]	53.7	92.4	673	2403	5826
Deep SORT [26]	59.3	94.2	528	1881	4334
Tractor [1]	56.5	94.4	687	1706	<b>4216</b>
FairMOT [28]	61.5	91.9	768	<b>1179</b>	7568
MOT Neural Solver [4]	<b>62.9</b>	<b>94.5</b>	<b>431</b>	1653	4394

**Fig. 3.** Proportion of pan identity switches vs.  $\delta$  plot for video number 9. Majority of the identity switches (90% at a threshold of  $\delta = 40$  frames) occur due to camera panning, which is the main cause of error.

dataset - a large scale object detection, segmentation, and captioning dataset, popular in computer vision [12] and fine tuned on the hockey tracking dataset. The object detector obtains an average precision (AP) of 70.2 on the test videos (Table 4). The accuracy metrics for tracking used are the CLEAR MOT metrics [2] and Identification F1 score (IDF1) [18]. A ground truth object missed by the trackers is called a false negative (FN) whereas a false alarm is called a false positive (FP). For any tracker, a low number of false positives (FP) and false negatives (FN) are favoured. An important metric is the number of identity switches (IDSW), which occurs when a ground truth ID  $i$  is assigned a tracked ID  $j$  when the last known assignment ID was  $k \neq j$ . A low number of identity switches is an indicator of accurate tracking performance. For sports player tracking, the IDF1 is considered a better accuracy measure than Multi Object Tracking accuracy (MOTA) since it measures how consistently the identity of a tracked object is preserved with respect to the ground truth identity. The overall results are shown in Table 3. The best tracking performance is achieved using the MOT Neural Solver tracking model [4] re-trained on the hockey dataset. The MOT Neural Solver model obtains the highest MOTA score of 94.5 and IDF1 score of 62.9 on the test videos.

**Table 4.** Player detection results on the test videos.  $AP$  stands for Average Precision.  $AP_{50}$  and  $AP_{75}$  are the average precision at an IOU of 0.5 and 0.75 respectively.

$AP$	$AP_{50}$	$AP_{75}$
70.2	95.9	87.5



**Fig. 4.** Proportion of pan-identity switches for all videos at a threshold of  $\delta = 40$  frames. On average, pan-identity switches account for 65% of identity switches.

## 6 Discussion

From Table 3 it can be seen that the MOTA score of all methods is above 90%. This is because MOTA is calculated as

$$MOTA = 1 - \frac{\sum_t (FN_t + FP_t + IDSW_t)}{\sum_t GT_t} \quad (1)$$

where  $t$  is the frame index and  $GT$  is the number of ground truth objects. MOTA metric counts detection errors through the sum  $FP + FN$  and association errors through  $IDSWs$ . Since false positives (FP) and false negatives (FN) heavily rely on the performance of the player detector, the MOTA metric highly depends on the performance of the detector. For hockey player tracking, the player detection accuracy is high because of the sufficiently large size of players in broadcast video and a reasonable number of players and referees (with a fixed upper limit) to detect in the frame. Therefore, the MOTA score for all methods is high.

The SORT [3] algorithm obtains the least IDF1 score and the highest number of identity switches. This is due to the linear motion model assumption and simple IOU score for re-identification. Deep SORT [25], on the other hand uses features obtained from deep network for re-identification resulting in better IDF1 score and lower identity switches. For JDT based networks, performing detection and re-identification with a single network using a multi-task loss performs better than having separate networks for detection and re-id tasks, evident by better performance of FairMOT [28] compared to Tractor [1]. JDT tracking

**Table 5.** Tracking performance of MOT Neural Solver model for the 13 test videos ( $\downarrow$  means lower is better,  $\uparrow$  means higher is better).

Video #	IDF1 $\uparrow$	MOTA $\uparrow$	ID-switches $\downarrow$	False positives (FP) $\downarrow$	False negatives (FN) $\downarrow$	Duration (sec.)
1	78.53	94.95	23	100	269	36
2	61.49	93.29	26	48	519	29
3	55.83	95.85	43	197	189	43
4	67.22	95.50	31	77	501	49
5	72.60	91.42	40	222	510	40
6	66.66	90.93	38	301	419	35
7	49.02	94.89	59	125	465	48
8	50.06	92.02	31	267	220	34
9	53.33	96.67	30	48	128	29
10	55.91	95.30	26	65	193	26
11	56.52	96.03	40	31	477	45
12	87.41	94.98	14	141	252	35
13	62.98	94.77	30	31	252	22

algorithms, however, [28, 1] do not show any significant improvement over deep SORT evident by lower identity switches of deep SORT in comparison. The MOT Neural Solver method achieves the highest IDF1 score of 62.9 and significantly lower identity switches than the other methods. This is because the other trackers use a linear motion model assumption which does not perform well with the motion of hockey players. Sharp changes in player motion often leads to identity switches. The MOT Neural Solver model, in contrast, has no such assumptions since it poses tracking as a graph edge classification problem.

Table 5 shows the performance of the MOT Neural solver for each of the 13 test videos. We do a failure analysis to determine the cause of identity switches and low IDF1 score in some videos. The major sources of identity switches are severe occlusions and players going out of the camera FOV (due to camera panning and/or player movement). We define a pan-identity switch as an identity switch resulting from a player leaving and re-entering camera FOV due to camera panning. It is very difficult for the tracking model to maintain identity in these situations since players of the same team look identical with features such as, jersey color, helmet model, visor model, stick model, glove model, skate model, tape color etc unidentifiable from bounding boxes cropped from 720p broadcast clips. During a pan-identity switch, a player going out of the camera FOV at a particular point in screen coordinates can re-enter at any other point. We estimate the proportion of pan-identity switches to determine the contribution of panning to total identity switches.

To estimate the number of pan-identity switches, since we have quality annotations, we make the assumption that the ground truth annotations are accurate and there are no missing annotations in the ground truth. Based on this assumption, there is a significant time gap between two consecutive annotated detections of a player only when the player leaves the camera FOV and comes back again. Let  $T_{gt} = \{o_1, o_2, \dots, o_n\}$  represent a ground truth tracklet, where  $o_i = \{x_i, y_i, w_i, h_i, I_i, t_i\}$  represents a ground truth detection. A pan-identity switch is expected to occur during tracking when the difference between timestamps (in frames) of two consecutive ground truth detections  $i$  and  $j$  is greater



than a sufficiently large threshold  $\delta$ . That is

$$(t_i - t_j) > \delta \quad (2)$$

Therefore, the total number of pan-identity switches in a video is approximately calculated as

$$\sum_G \mathbb{1}(t_i - t_j > \delta) \quad (3)$$

where the summation is carried out over all ground truth trajectories and  $\mathbb{1}$  is an indicator function. Consider the video number 9 in Table 5 having 30 identity switches and a low IDF1 of 53.33. We plot the proportion of pan identity switches, that is

$$= \frac{\sum_G \mathbb{1}(t_i - t_j > \delta)}{IDSW_s} \quad (4)$$

against  $\delta$ , where  $\delta$  varies between 40 and 80 frames. From Fig. 3 it can be seen that majority of the identity switches ( 90% at a threshold of  $\delta = 40$  frames) occur due to camera panning. Visually investigating the video confirmed the statement. Fig. 4 shows the proportion of pan-identity switches for all videos at a threshold of  $\delta = 40$  frames. On average, pan identity switches account for 65% of identity switches in the videos. This shows that the tracking model is able to tackle a majority of other sources of errors which include minor occlusions and lack of detections. The primary source of errors are pan-identity switches and extremely cluttered scenes.

## 7 Conclusion

In this paper, we test five state-of-the-art tracking algorithms on the ice hockey dataset and analyzed their performance. From the performance of trackers we infer that trackers with a linear motion model do not perform well on hockey dataset, evident by the high number of identity switches occurring in models with linear motion assumption. The best performance is obtained by the MOT neural solver model [4], that uses a graph based approach towards tracking without any linear motion model assumption. Also, the IDF1 metric is a better metric for hockey player tracking since the MOTA metric is heavily influenced by player detection accuracy. We find that the main source of error in hockey player tracking in broadcast video are pan-identity switches - identity switches results due to players going outside the broadcast camera FOV.

## 8 Acknowledgments

This work was supported by Stathletes through the Mitacs Accelerate Program and the Natural Sciences and Engineering Research Council of Canada (NSERC). We also acknowledge Compute Canada for hardware support.

## References

1. Bergmann, P., Meinhardt, T., Leal-Taixé, L.: Tracking without bells and whistles. In: The IEEE International Conference on Computer Vision (ICCV) (October 2019)
2. Bernardin, K., Stiefelhagen, R.: Evaluating multiple object tracking performance: The clear mot metrics. *EURASIP Journal on Image and Video Processing* **2008** (01 2008). <https://doi.org/10.1155/2008/246309>
3. Bewley, A., Ge, Z., Ott, L., Ramos, F., Upcroft, B.: Simple online and realtime tracking. In: 2016 IEEE International Conference on Image Processing (ICIP). pp. 3464–3468 (2016). <https://doi.org/10.1109/ICIP.2016.7533003>
4. Braso, G., Leal-Taixe, L.: Learning a neural solver for multiple object tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
5. Cai, Y., de Freitas, N., Little, J.J.: Robust visual tracking for multiple targets. In: Leonardis, A., Bischof, H., Pinz, A. (eds.) *Computer Vision – ECCV 2006*. pp. 107–118. Springer Berlin Heidelberg, Berlin, Heidelberg (2006)
6. Dendorfer, P., Rezatofighi, H., Milan, A., Shi, J.Q., Cremers, D., Reid, I.D., Roth, S., Schindler, K., Leal-Taix'e, L.: Mot20: A benchmark for multi object tracking in crowded scenes. *ArXiv abs/2003.09003* (2020)
7. Gadde, C.A., Jawahar, C.: Transductive weakly-supervised player detection using soccer broadcast videos. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 965–974 (January 2022)
8. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: 2012 IEEE Conference on Computer Vision and Pattern Recognition. pp. 3354–3361 (2012). <https://doi.org/10.1109/CVPR.2012.6248074>
9. Hurault, S., Ballester, C., Haro, G.: Self-supervised small soccer player detection and tracking. In: Proceedings of the 3rd International Workshop on Multimedia Content Analysis in Sports. p. 9–18. *MMSports '20*, Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3422844.3423054>, <https://doi.org/10.1145/3422844.3423054>
10. IIHF: Survey of Players (2018), available online: <https://www.iihf.com/en/static/5324/survey-of-players>
11. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 936–944 (2017). <https://doi.org/10.1109/CVPR.2017.106>
12. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft *coco*: Common objects in context. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) *Computer Vision – ECCV 2014*. pp. 740–755. Springer International Publishing, Cham (2014)
13. Lu, W.L., Ting, J.A., Little, J.J., Murphy, K.P.: Learning to track and identify players from broadcast sports videos. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **35**(7), 1704–1716 (2013). <https://doi.org/10.1109/TPAMI.2012.242>
14. Milan, A., Leal-Taixé, L., Reid, I.D., Roth, S., Schindler, K.: Mot16: A benchmark for multi-object tracking. *ArXiv abs/1603.00831* (2016)
15. Okuma, K., Taleghani, A., de Freitas, N., Little, J.J., Lowe, D.G.: A boosted particle filter: Multitarget detection and tracking. In: Pajdla, T., Matas, J. (eds.) *Computer Vision - ECCV 2004*. pp. 28–39. Springer Berlin Heidelberg, Berlin, Heidelberg (2004)

16. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1. p. 91–99. NIPS'15, MIT Press, Cambridge, MA, USA (2015)
17. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. In: Cortes, C., Lawrence, N., Lee, D., Sugiyama, M., Garnett, R. (eds.) Advances in Neural Information Processing Systems. vol. 28. Curran Associates, Inc. (2015)
18. Ristani, E., Solera, F., Zou, R., Cucchiara, R., Tomasi, C.: Performance measures and a data set for multi-target, multi-camera tracking. In: Hua, G., Jégou, H. (eds.) Computer Vision – ECCV 2016 Workshops. pp. 17–35. Springer International Publishing, Cham (2016)
19. Sangüesa, A.A., Ballester, C., Haro, G.: Single-camera basketball tracker through pose and semantic feature fusion. ArXiv **abs/1906.02042** (2019)
20. Theagarajan, R., Bhanu, B.: An automated system for generating tactical performance statistics for individual soccer players from videos. IEEE Transactions on Circuits and Systems for Video Technology **31**(2), 632–646 (2021). <https://doi.org/10.1109/TCSVT.2020.2982580>
21. Theiner, J., Gritz, W., Müller-Budack, E., Rein, R., Memmert, D., Ewerth, R.: Extraction of positional player data from broadcast soccer videos. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 823–833 (January 2022)
22. Thomas, A.C.: The impact of puck possession and location on ice hockey strategy. Journal of Quantitative Analysis in Sports **2**(1) (2006)
23. Vermaak, Doucet, Perez: Maintaining multimodality through mixture tracking. In: Proceedings Ninth IEEE International Conference on Computer Vision. pp. 1110–1116 vol.2 (Oct 2003). <https://doi.org/10.1109/ICCV.2003.1238473>
24. Viola, P., Jones, M.: Rapid object detection using a boosted cascade of simple features. In: Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001. vol. 1, pp. I–I (Dec 2001). <https://doi.org/10.1109/CVPR.2001.990517>
25. Wojke, N., Bewley, A.: Deep cosine metric learning for person re-identification. In: 2018 IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 748–756. IEEE (2018). <https://doi.org/10.1109/WACV.2018.00087>
26. Wojke, N., Bewley, A., Paulus, D.: Simple online and realtime tracking with a deep association metric. In: 2017 IEEE International Conference on Image Processing (ICIP). pp. 3645–3649. IEEE (2017). <https://doi.org/10.1109/ICIP.2017.8296962>
27. Zhang, R., Wu, L., Yang, Y., Wu, W., Chen, Y., Xu, M.: Multi-camera multi-player tracking with deep player identification in sports video. Pattern Recognition **102**, 107260 (2020). <https://doi.org/https://doi.org/10.1016/j.patcog.2020.107260>, <https://www.sciencedirect.com/science/article/pii/S0031320320300650>
28. Zhang, Y., Wang, C., Wang, X., Zeng, W., Liu, W.: Fairmot: On the fairness of detection and re-identification in multiple object tracking. International Journal of Computer Vision pp. 1–19 (2021)
29. Zhou, X., Koltun, V., Krähenbühl, P.: Tracking objects as points. European Conference on Computer Vision (ECCV) (2020)