

This CVPR 2021 workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version;

the final published version of the proceedings is available on IEEE Xplore.

Detecting and Matching Related Objects with One Proposal Multiple Predictions

Yang Liu

Luiz G. Hafemann Michael Jamieson Sportlogiq, Montreal, Canada Mehrsan Javan

{yang, luiz, mikej, mehrsan}@sportlogig.com

Abstract

Tracking players in sports videos is commonly done in a tracking-by-detection framework, first detecting players in each frame, and then performing association over time. While for some sports tracking players is sufficient for game analysis, sports like hockey, tennis and polo may require additional detections, that include the object the player is holding (e.g. racket, stick). The baseline solution for this problem involves detecting these objects as separate classes, and matching them to player detections based on the intersection over union (IoU). This approach, however, leads to poor matching performance in crowded situations, as it does not model the relationship between players and objects. In this paper, we propose a simple yet efficient way to detect and match players and related objects at once without extra cost, by considering an implicit association for prediction of multiple objects through the same proposal box. We evaluate the method on a dataset of broadcast ice hockey videos, and also a new public dataset we introduce called COCO +Torso. On the ice hockey dataset, the proposed method boosts matching performance from 57.1% to 81.4%, while also improving the meanAP of player+stick detections from 68.4% to 88.3%. On the COCO + Torso dataset, we see matching improving from 47.9% to 65.2%. The COCO + Torso dataset, code and pre-trained models will be released at https: //github.com/foreverYoungGitHub/detectand-match-related-objects.

1. Introduction

Detecting players is the foundation of automated player tracking in sport games [18, 12, 4], in a tracking-by-detection framework [2, 21, 20]. The detection task normally focuses on finding bounding boxes around players and, for some applications, a box around the ball. For sports like soccer or basketball, these detections can be sufficient for game analysis and estimation of players' physical metrics [15]. In sports like hockey, polo and tennis, where players use a tool (e.g. a racket or a hockey stick), these detections can be sufficient for game analysis and estimation of players' physical metrics [15].

tions are not sufficient for all applications. For these sports, it is also important to detect additional bounding boxes, either around the tools, or a larger bounding box around the player with the tool, and make the association to the player wielding it.

As an example, we consider a tracking system for hockey. For the application of classifying the type of a shot, the hockey stick needs to be tracked, and associated with the player holding it. In order to estimate the player's pose and stick location, a bounding box containing both the player and the stick can be used. On the other hand, for player identification (recognizing jersey numbers, or extracting Re-ID features), a bounding box only detecting the player is more appropriate. During crowded scenes (e.g. players fighting over a puck), players are close to each other, and the extended bounding box can contain part of another player, degrading performance for player identification.

The naive solution for this problem is to detect these boxes independently. This requires associating the regular and extended bounding boxes for the player, for instance using the Hungarian method. In crowded scenes, heuristics such maximizing intersection-over-union (IoU) or minimizing visual feature distance are not reliable for matching. Therefore, it is appealing to solve the problem of detecting the related boxes and matching them together at the same time in an end-to-end framework.

In this work, we propose a robust yet efficient method to extract both types of bounding boxes at the same time, where the association between them is implicit. This implicit association comes at no extra cost, making it appealing for real-time tracking systems. Furthermore, the model can be trained without annotating the extended boxes in every image, thereby reducing the overall annotation effort.

The main contributions for this paper are:

 We propose a method to simultaneously detect and match related bounding boxes. The related objects are naturally grouped by implicit association without extra cost. This includes the introduction of SetNMS to suppress duplicate sets of detections and to keep more possible detection sets in crowded scenes. We validate our proposed method on a broadcast ice hockey dataset and a public dataset we introduce called COCO +Torso dataset. Compared with the baseline model, the proposed methods performs better in both detection and matching for these datasets. Based on the ablation study, the proposed method is particularly suited for applications with highly overlapped associations.

2. Related work

2.1. Object Detection in Sports

Object detection has been widely used in automated tracking systems for sports [18, 12, 4, 19]. The task poses additional challenges to generic object detection, including camera distortion and fast movement (for broadcast videos), as well as frequent crowded scenes. Thaler *et al.* [17] and Faulkner *et al.* [7] adopt HOG and Haar features to detect and classify players in a sliding window for soccer and football. Acuna [1] introduces an end-to-end CNN-based object detector in the basketball. In order to improve the detection result for fisheye cameras in soccer, Cioppa *et al.* [4] applies a single stage object detector on cropped and rotated image patches.

2.2. Object Detection with Multiple Predictions

Object detection with multiple predictions is a new approach for end-to-end CNN-based object detectors. The methods proposed in the literature are applied to two stage object detectors, with the goal of improving detector performance in crowded scenes. Chu *et al.* [3] introduces multiple instance detection head to improve the object detection in the extremely high overlap scenario by predicting multiple unmatched objects from single proposal box. Zhang *et al.* [22] and Huang *et al.* [11] propose paired-RPN to predict the matched human visible part and human body with two proposal boxes to improve the pedestrian detection in crowd scenes. It is worth noting that these methods were designed to improve human detection and matching problems.

3. Framework

In order to detect and match related objects, we design the detector to output multiple predictions from a single proposal box, as shown in Figure 1.

During training, we introduce a specific loss function to jointly optimize the prediction of multiple associated objects from a single prediction box. For inference, we introduce the SetNMS procedure to suppress duplicated sets of detections.



Figure 1: Illustration of the proposed network architecture for one proposal multiple prediction, applied to the ice hockey application. The arrows represent the data flow in the network: the red arrows indicate the loss functions applied to the networks during training, while the blue arrows shows the post-processing operations in inference.

3.1. Detection Head Design for One Proposal Multiple Predictions

The loss functions in modern end-to-end detectors are based on proposal boxes: given a set of ground truth bounding boxes for an image, these are first matched to proposed boxes and then each proposal box is refined to its box location and target label. Besides not modeling relationships between boxes that belong to the same group, this creates another problem: each proposed box is forced to correspond to a single ground truth box, which creates an issue when multiple ground truth boxes are too close to each other. This is a common situation in the problems at hand. For example, the box for a player and for the player holding the stick normally have high overlap.

We formulate the detection problem to jointly predict a set of related objects, such that the association between the objects is implicit, addressing both the issue of associating bounding boxes, and the problem of training a detector when multiple ground-truth boxes overlap. We consider $b_{\text{base}} = \{x_{\text{base}}, y_{\text{base}}, w_{\text{base}}\}$ the *base* bounding box (e.g. for the player), and a sequence of N associated bounding boxes $B_{\text{extra}} = \{b_i | 1 \leq i \leq N\}$. We model the associated bounding boxes as offsets from the base box. That is, for the extra box i, the model estimates $\{x_{\text{offset}}^i, y_{\text{offset}}^i, w_{\text{offset}}^i, h_{\text{offset}}^i\}$, and the bounding box is specified as $b_i = \{x_{\text{base}} + x_{\text{offset}}^i, y_{\text{base}} + y_{\text{offset}}^i, w_{\text{base}} + w_{\text{offset}}^i\}$. Finally, given an input image X, the detector estimates a set of paired (matched) boxes $\{b_{\text{base}}\} \cup B_{\text{extra}}$, together with a class c (for the whole set) and and an objectness score s.

This formulation assumes that the related objects can be predicted from the base object. Based on this assumption, the proposal boxes are first matched to the objects in the base class, and the related objects are matched to the same proposal boxes. Therefore, the same proposal box is used for the group of related objects. Since the predicted objects are already matched, there is only one predicted score and class for each matched group. Additionally, the problem for high overlapped ground truth bounding boxes (for objects on the same group) is solved, as the related objects share the same proposal box, instead of competing for it. It should be noticed that, even though the detection head we propose is designed on top of single-stage detectors like Yolo/FPN, it can be easily extended to other two stage detector frameworks like Faster R-CNN and Mask R-CNN.

3.2. Loss Functions

We consider the following loss function:

$$L_{\text{total}} = L_{\text{cls}} + L_{\text{loc}}^{\text{base}} + \sum_{i=1}^{n} (L_{\text{loc}}^{i} + L_{\text{constraint}}^{i})$$
(1)

Where L_{cls} is the classification loss for the set. In this work we use the focal loss [13]; L_{loc}^{base} and L_{loc}^{i} are the localization losses for the base box and extra boxes, respectively, that are defined below. Finally, $L_{constraint}^{i}$ is a term to represent the relationship between boxes from the base classes and the related classes *i*. This is an optional term, that we can use to model domain knowledge for a particular application, as described below and in Equation 3.

For the localization losses, we use the Distance-IoU Loss [23], defined as follows:

$$L_{\rm loc} = 1 - {\rm IoU} + \frac{||c - c_{\rm gt}||^2}{d^2}$$
(2)

Where IoU indicate the intersection-over-union of the predicted and ground-truth boxes, c and c_{gt} are the center points of the predicted and GT boxes (respectively), d is the diagonal length of the smallest enclosing box that covers the two boxes.

We also model an optional constraint loss, that can be used to model relationships between base and extra bounding boxes. In particular, we consider a term to encourage that an extra bounding box b_i be strictly larger than the base bounding b_{base} . This term is appropriate for situations where the extra bounding boxes are intended to always enclose the base bounding box. This loss term is defined by adding a penalty if the extra bounding box does not cover the base bounding box. We use the following loss for the x component:

$$L_{\text{constraint}_x}^i = \begin{cases} \text{smooth}_{L1}(x_{\text{base}} - x_i), & \text{if } x_i > x_{\text{base}} \\ 0, & \text{otherwise} \end{cases} (3)$$

Where smooth_{L1} is the Smooth L1 loss defined in [8]. We use equivalent losses for y, x + w and y + h, penalizing the

following conditions: $y_i > y_{\text{base}}$, $x_i + w_i < x_{\text{base}} + w_{\text{base}}$, and $y_i + h_i < y_{\text{base}} + h_{\text{base}}$.

Algorithm 1 SetNMS

	Input									
	$B = \{\{b_1^1,, b_1^C\},, \{b_N^1,, b_N^C\}\}$: boxes.									
	$S = \{s_1,, s_N\}$: scores.									
	$C = \{c_1,, c_N\}$: classes.									
	Ω_{IoU} : IoU threshold for SetNMS.									
1:	procedure SetNMS (B, S, C, Ω_{IoU})									
2:	$R \leftarrow \{\}$									
3:	while $B eq \emptyset$ do									
4:	Let T be the highest scored group									
5:	Remove T from B, S and C , add it to R									
6:	for $\left\{b_{i}^{1},,b_{i}^{C} ight\}\in B$ do									
7:	if $c_i \neq c_T$ then Continue									
8:	$IoU_{set} \leftarrow \left\{ IoU(b_T^1, b_i^1),, IoU(b_T^C, b_i^C) \right\}$									
9:	if $\min(IoU_{set}) > \Omega_{IoU}$ then									
0:	Remove <i>i</i> -th element from B, S, C									
	return R									

3.3. SetNMS

1

During inference time, detectors normally use a nonmaximum suppression (NMS) algorithm to avoid duplicated predictions to the same object. In the method above, we predict a set of matched boxes for each proposal, and therefore we need to adapt this procedure to work with a set of predictions - we name this process SetNMS.

The SetNMS algorithm is described in Algorithm 1. Since the predicted matched group shares the same classification result and confidence score, SetNMS suppresses false positive predictions by taking all the matched boxes in the group into consideration rather than taking one predicted box. Only all the matched boxes in the group that have high overlap with the boxes in the kept group will be considered as false positive predictions and be removed. Since the SetNMS compares the ensemble group of the matched boxes rather than a single box, it also keeps more reliable boxes in the crowd scene and more robust for the whole ensemble group compared to the original NMS.

4. Experiments

In this section, we discuss our evaluation protocol, and discuss the results of our approach on our ice hockey dataset and the COCO +Torso dataset.

4.1. Datasets and Evaluation Metric

Ice hockey dataset is a dataset for ice hockey player detection in broadcast videos. The images are sampled from multiple ice hockey leagues including the National Hockey League (NHL) and minor leagues. Figure 2 shows two images of this dataset. Two classes are considered: "player" -



Figure 2: Samples from the Ice hockey, COCO +Torso and COCO +Torso_S datasets. The blue boxes indicate the detections for the base class, which is the player class in ice hockey dataset and person class in COCO +Torso dataset, while the green boxes are the bounding boxes for the related classes, which is the player+stick class in ice hockey dataset and torso class in COCO +Torso dataset.

a tight bounding box around the player, and "player+stick" - a bounding box that includes the player and the hockey stick. Only part of the images contains the "player+stick" label. This dataset contains 15095 images in the training set and 1886 images in the validation set. Of these, only 3286 images in the training set and 1634 images in the validation set have the annotation for the "player+stick" class. For the samples where the player+stick annotation is not available, we simply remove the localization loss for the related classes in Equation 1. Note that we always consider the constraint loss, that drives the model to predict player+stick bounding boxes that are always larger than the player bounding boxes, even when we do not have annotations for player+stick. We benchmark the proposed method to predict both classes.

COCO +Torso dataset is a public dataset introduced in this work, based on the COCO[14] and DensePose[9] datasets. The DensePose dataset is mainly used to predict dense human body joints and contains 3 labels: human bounding boxes, human body joints and human densepose. Based on these three labels, we generated the COCO +Torso dataset featuring detections for two classes: a "person" class with their matched "torso". The human body joints are used to filter the human bounding boxes which does not contains the any torso joints annotation, while the DensePose annotation is used to generate the torso bounding boxes. Two sample images in this dataset are shown in Figure 2 (b). The training and validation sets for COCO +Torso dataset are split based on its original split set in the DensePose dataset. The final COCO +Torso dataset contains 26437 and 5984 images in training and validation set respectively.

COCO +Torso_S dataset In order to simplify the COCO person-torso detection and matching task, we also sample an easier dataset from COCO +Torso dataset and named it as COCO +Torso_S dataset. This simplified dataset only contains the images which all the persons in the image have the matched torso bounding boxes. Two sample images in this dataset is shown in Figure 2 (c). After sampling, the training and validation sets for the COCO +Torso_S dataset contain 13483 and 2215 images separately.

Evaluation Metric. Since we consider a joint detection and matching task, the standard log-average miss rate (MR) [5] and average precision $(AP^{.50})$ [6] are adopted as detection evaluation metrics for each class. The MRis computed in the false positive per image (FPPI) with a range of $[10^{-2}, 100](MR^{-2})$, while the $AP^{0.5}$ calculates the average precision with IoU threshold 0.5 with a range of [0,1]. In order to evaluate model performance in object matching, we also extend the log-average miss rate and average precision and introduce the matching log-average miss rate (MR_{match}) and average precision $(AP_{match}^{.50})$. Compared with the standard MR and $AP^{.50}$, in MR_{match} and $AP^{.50}_{match}$, the true positive (TP) is counted only if the IoU score for all the objects in the group is higher than IoU threshold. The range of $AP_{match}^{.50}$ is in $[0, \min(AP_1^{.50}, ..., AP_C^{.50})]$, while the range of MR_{match} is in $[\max(MR_1, ..., MR_C), 100]$, where C is the classes for the grouped objects.

Dataset	Method	$AP^{0.5}(\uparrow)$			$MR^{-2}(\downarrow)$		
		Player	Player+Stick	Match	Player	Player+Stick	Match
	Yolo Player Only	98.9			1.6		
Ice Hockey	Yolo	98.5	68.4	57.1	1.9	52	63.1
	Yolo+MP	99	88.3	81.4	1.5	20.8	33.2
		Person	Torso	Match	Person	Torso	Match
	Yolo	72.7	67.6	47.9	34.2	45.6	71.8
COCO L Torso	Yolo+MP	74.7	66.1	65.2	29.2	48.5	52.1
0000 + 10180	FPN	71	67.7	45.9	35.7	46.1	73.3
	FPN+MP	72.7	64.4	63.3	30.7	48.9	52.3
	Yolo	97.1	90.5	69.3	4.6	15	51.1
	Yolo+MP	98.1	87.3	86.1	2.7	21.5	24.1
$COCO + 10180_S$	FPN	96.3	91.2	68.3	5.5	13.8	51.2
	FPN+MP	97.7	87.7	87.1	3.1	21.5	22.8

Table 1: Detection and matching results on the three datasets. +**MP** indicates using the detection head and postprocessing methods proposed in this paper.

4.2. Experimental protocol

In these experiments, we adopt the state-of-the-art single stage object detector pretrained at COCO dataset [14] as our baseline. The feature extractor in the baseline models is ResNet-18[10], while YOLOv3[16] and FPN[13] are used as the detection head in the baseline models. It should be noted that there is only 1 anchor box for each proposal region in the anchor-based single stage detector. The model input size for ice hockey dataset is 1280x720, while the input size for COCO +Torso datasets is 320x320.

The same training strategy is applied for all experiments: we first train only the detection head for 60 epochs, using Adam with learning rate 4×10^{-4} and then fine-tune the whole network for 250 epochs using Adam with learning rate 10^{-4} . During the training, we adopt center sampling and scale matching to match the proposal boxes with ground truth bounding box. We also use standard data augmentation methods: random resizing, cropping, flipping, random contrast and saturation. We do not employ multi-scale predictions during training or testing.

4.3. Results

Ice hockey dataset: We consider two baselines for this dataset: a detector trained only for the player class and a detector trained on the full dataset, but treating the two classes as separate. For the latter, we match the player and player+stick predictions at inference time using the Hungarian algorithm, based on the IoU and confidence score between boxes of the two classes. We compare these baselines to the proposed method that outputs detections for both classes from the same proposal, which performs the matching implicitly. The results are shown in Table 1. We can see that the baseline detector achieves 98.9 AP in player class when it is only trained for this class. When the ex-

tra player+stick class is introduced to the baseline detector, the AP_{player} drops 0.4 and the $AP_{player+stick}$ only achieves 68.4. We hypothesize that this worse performance is caused by the the two classes competing for the same proposal boxes in situations with high overlap between the two classes. The proposed detector solves this issue by design, and achieve similar accuracy for the Player class, while significantly boosting performance for the Player+stick class to 88.3%.

We notice a more significant difference on the the matching performance (AP_{match} and MR_{match}). The AP_{match} for the baseline method is 57.1 which is 83.4% of its upper bound (AP for the player+stick class), while the AP_{match} for the proposed detector achieves 81.4%, which is the 92.2% of its upper bound. The MR_{match} for the proposed detector is also 1.90 times lower than the baseline detector as well, showing that the proposed detector the best result in all detection and matching evaluation metrics.

COCO +Torso dataset: In order to check the generalization of our proposed method on non-sport tasks, we trained and evaluated the proposed framework in the COCO +Torso dataset to detect the bounding boxes for the person and torso classes and matched the corresponding boxes for each person. The result for this experiment is shown in Table 1. Compared with the baseline models, the proposed detector with one proposal multiple prediction head improves AP_{person} by 1-2 percent points, but decreases AP_{torso} by 1.5-3.5 percent points. We notice, however, a significant improvement in matching performance: AP_{match} has 1.36 and 1.24 improvement for COCO +Torso and COCO +Torso_S dataset respectively. This result shows that the proposed framework is suitable for applications where the *matching* task is important.

Ablation Study on Proposal Boxes Matching: The re-



Figure 3: Density distribution for IoU scores between the base objects and related objects among ice hockey, COCO +Torso and COCO +Torso $_S$ dataset.

sults from Table 1 show that the proposed method achieves better results on base class detection and group matching in all experiments. However, while the proposed model trained with ice hockey dataset has significant improvement in the related classes, the same model trained on the COCO +Torso dataset performs worse than the regular detector for the extra "Torso" class. In order to obtain insights on which situation the proposed framework works better, we conducted two ablation studies: we first analyze the distribution of IoU scores between base and extra bounding boxes (Figure 3). In a second experiment, we compute the average number of proposal boxes matches per image, among all feature maps (Figure 4).

Figure 3 shows the distribution of the IoU scores for matched objects in the three datasets (e.g. the IoU between a player and the associated box of player+stick). We notice a very different distribution for the ice hockey dataset, where the overlap between the two classes is very high (mean of 0.94), whereas for the COCO +Torso datasets, we see a mean around 0.5 (0.46 for COCO +Torso_S). Since the IoU score of paired objects is only 0.5 in the COCO +Torso dataset, in the regular detector, the groundtruth boxes can



Figure 4: The number of proposal boxes matched per image in each feature map for the three datasets. The x axis is the stride between the feature map and original input size, while the y axis is the number of matched proposal boxes for each class.

Dataset	Method	$AP^{0.5}(\uparrow)$			
		Player	+Stick	Match	
	NMS	98.9	88.1	81.2	
Ice Hockey	JointNMS[22]	98.8	88	81	
	SetNMS	99	88.3	81.4	
		Person	Torso	Match	
	NMS	76.4	66.7	66.5	
COCO	JointNMS[22]	76.3	66.6	65.3	
+Torso	SetNMS	74.7	66.1	65.2	
	NMS	98.4	87	85.7	
COCO	JointNMS[22]	98.2	86.9	85.4	
$+Torso_S$	SetNMS	98.1	87.3	86.1	

Table 2: Ablation studies on NMS methods for the one proposal multiple prediction model.

be assigned with the suitable proposal boxes based on its center and size. However, for ice hockey dataset, since the matched objects are highly overlapped with each other, these objects almost always need to compete for the same proposal boxes in regular detector. Therefore, the proposed method is particularly suited for tasks with high-overlap between the base bounding box and extra boxes.

Figure 4 shows the number of proposal box matches, per feature map, in the three different resolutions considered by the detection heads (for stride s the feature map has size $\frac{h}{s} \times \frac{w}{s}$). The left plot (a) shows that the player and player+stick tend to match to the proposal boxes located in the same feature maps in ice hockey dataset. This indicates that the same feature maps should contain the rich information for both matched classes in this dataset. However, in COCO +Torso dataset, since the torso objects are much smaller than the person objects, the paired objects are matched to the proposal boxes in the different feature maps in the regular detectors (plots (b) and (c)). In the proposed method, the proposal boxes are matched to the bounding box of the *base* objects, and its feature map may not be suitable for the related matched objects. Therefore, the proposed detector performs slightly worse in the torso classes in COCO +Torso dataset. This reinforces the previous finding that the proposed method is particularly suited for applications where the base and extra classes have high overlap.

Ablation Study on NMS methods: In order to validate the proposed SetNMS, we conducted an ablation study with three different NMS methods: (i) the baseline (denoted simply by NMS) considers performing NMS only with the IoU score of the base class; (ii) JointNMS[22] removes group of objects as long as objects in *one class* are overlapped; (iii) the proposed SetNMS, which removes the group only if all of objects in the group are overlapped with other groups. The result of this comparison is shown in Table 2.

Based on its formulation, the proposed SetNMS should keep the highly overlapped groups in a crowded scene, but also introduce more false alarm groups with low confidence scores. In Table 2, we see that the results for different methods differ only by less than 1 percent point. Set-NMS achieves the best result in all detection and matching evaluation metrics for Ice Hockey dataset, while NMS achieves the best result for COCO +Torso dataset. For COCO +Torso_S dataset, NMS has better AP in person detection while the best torso detection and objects match AP is generated by SetNMS. This result indicates that when the model is not confident about the predicting boxes, the SetNMS actually performs worse than the normal NMS. Therefore, the choice of the best NMS procedure can be application-specific.

5. Conclusion

In this paper, we propose an efficient object detector that can simultaneously detect and match player bounding boxes and related bounding boxes (e.g. player holding a stick), that can be used for applications where detecting just the player is not sufficient. This method considers an implicit association for the multiple predictions through the same proposal box to detect and match the related objects without extra cost. Our experiments show that this method is particularly suited for applications that require high matching accuracy, and situations where the base and related bounding boxes have high overlap.

References

- David Acuna. Towards Real-time Detection and Tracking of Basketball Players using Deep Neural Networks. In Proceedings of the Conference on Neural Information Processing Systems (NIPS), pages 4–9, 2017. 2
- [2] Alex Bewley, Zongyuan Ge, Lionel Ott, Fabio Ramos, and Ben Upcroft. Simple Online and Realtime Tracking. In Proceedings of the IEEE International Conference on Image Processing (ICIP), pages 3464–3468. IEEE, 2016. 1
- [3] Xuangeng Chu, Anlin Zheng, Xiangyu Zhang, and Jian Sun. Detection in Crowded Scenes: One Proposal, Multiple Predictions. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 12214–12223, 2020. 2
- [4] Anthony Cioppa, Adrien Deliege, Noor Ul Huda, Rikke Gade, Marc Van Droogenbroeck, and Thomas B Moeslund. Multimodal and Multiview Distillation for Real-time Player Detection on a Football Field. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 880–881, 2020. 1, 2
- [5] Piotr Dollar, Christian Wojek, Bernt Schiele, and Pietro Perona. Pedestrian Detection: An Evaluation of the State of the Art. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(4):743–761, 2011. 4
- [6] Mark Everingham, SM Ali Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The

PASCAL Visual Object Classes Challenge: A Retrospective. International Journal of Computer Vision, 111(1):98–136, 2015. 4

- [7] H. Faulkner and A. Dick. AFL Player Detection and Tracking. In Proceedings of the International Conference on Digital Image Computing: Techniques and Applications (DICTA), pages 1–8, 2015. 2
- [8] R. Girshick. Fast R-CNN. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), pages 1440–1448, 2015. 3
- [9] Rıza Alp Güler, Natalia Neverova, and Iasonas Kokkinos. DensePose: Dense Human Pose Estimation in the Wild. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 7297–7306, 2018.
 4
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 770–778, 2016. 5
- [11] Xin Huang, Zheng Ge, Zequn Jie, and Osamu Yoshie. NMS by Representative Region: Towards Crowded Pedestrian Detection by Proposal Pairing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (CVPR), pages 10750–10759, 2020. 2
- [12] Jacek Komorowski, Grzegorz Kurzejamski, and Grzegorz Sarwas. FootAndBall: Integrated player and ball detector. In Proceedings of the International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP), pages 47–56, 2020. 1, 2
- [13] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal Loss for Dense Object Detection. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 2980–2988, 2017. 3, 5
- [14] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common Objects in Context. In Proceedings of the European Conference on Computer Vision (ECCV), pages 740–755. Springer, 2014. 4, 5
- [15] M. Manafifard, H. Ebadi, and H. Abrishami Moghaddam. A Survey on Player Tracking in Soccer Videos. *Computer Vision and Image Understanding*, 159:19–46, June 2017. 1
- [16] Joseph Redmon and Ali Farhadi. YoloV3: An Incremental Improvement. arXiv preprint arXiv:1804.02767, 2018. 5
- [17] M. Thaler and W. Bailer. Real-Time Person Detection and Tracking in Panoramic Video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1027–1032, 2013.
 2
- [18] Graham Thomas, Rikke Gade, Thomas B. Moeslund, Peter Carr, and Adrian Hilton. Computer Vision for Sports: Current Applications and Research Topics. *Computer Vision and Image Understanding*, 159:3–18, June 2017. 1, 2
- [19] Roman Voeikov, Nikolay Falaleev, and Ruslan Baikulov. TTNet: Real-time Temporal and Spatial Video Analysis of Table Tennis. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pages 884–885, 2020. 2

- [20] Zhongdao Wang, Liang Zheng, Yixuan Liu, Yali Li, and Shengjin Wang. Towards Real-Time Multi-Object Tracking. arXiv:1909.12605, July 2020. 1
- [21] Nicolai Wojke, Alex Bewley, and Dietrich Paulus. Simple Online and Realtime Tracking with a Deep Association Metric. In *Proceedings of the IEEE International Conference on Image Processing (ICIP)*, pages 3645–3649. IEEE, 2017. 1
- [22] Kevin Zhang, Feng Xiong, Peize Sun, Li Hu, Boxun Li, and Gang Yu. Double Anchor R-CNN for Human Detection in a Crowd. arXiv:1909.09998, 2019. 2, 6
- [23] Zhaohui Zheng, Ping Wang, Wei Liu, Jinze Li, Rongguang Ye, and Dongwei Ren. Distance-IoU Loss: Faster and Better Learning for Bounding Box Regression. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2020.
 3