Self-Supervised Small Soccer Player Detection and Tracking

Samuel Hurault
samuel.hurault@ens-paris-saclay.fr
ENS Paris-Saclay
Cachan, France

Coloma Ballester*
coloma.ballester@upf.edu
Universitat Pompeu Fabra
Barcelona, Spain

Gloria Haro*
gloria.haro@upf.edu
Universitat Pompeu Fabra
Barcelona, Spain

Figure 1: Extracted frame from our proposed tracking method applied to a wide-angle view soccer video sequence.

ABSTRACT

In a soccer game, the information provided by detecting and tracking brings crucial clues to further analyze and understand some tactical aspects of the game, including individual and team actions. State-of-the-art tracking algorithms achieve impressive results in scenarios on which they have been trained for, but they fail in challenging ones such as soccer games. This is frequently due to the player small relative size and the similar appearance among players of the same team. Although a straightforward solution would be to retrain these models by using a more specific dataset, the lack of such publicly available annotated datasets entails searching for other effective solutions. In this work, we propose a self-supervised pipeline which is able to detect and track low-resolution soccer players under different recording conditions without any need of ground-truth data. Extensive quantitative and qualitative experimental results are presented evaluating its performance. We also present a comparison to several state-of-the-art methods showing that both the proposed detector and the proposed tracker achieve top-tier results, in particular in the presence of small players. *Code available at <https://github.com/samuro95/Self-Supervised-Small-Soccer-Player-Detection-Tracking>.*

CCS CONCEPTS

• **Computing methodologies** → **Object detection; Tracking.**

*co-supervision

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MMSports'20, October 16, 2020, Seattle, WA, USA

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8149-9/20/10...\$15.00

<https://doi.org/10.1145/3422844.3423054>

KEYWORDS

Player detection; multi-player tracking; soccer; small object detection; self-supervised; single camera; CNN; neural networks.

ACM Reference Format:

Samuel Hurault, Coloma Ballester, and Gloria Haro. 2020. Self-Supervised Small Soccer Player Detection and Tracking. In *3rd International Workshop on Multimedia Content Analysis in Sports (MMSports'20)*, October 16, 2020, Seattle, WA, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3422844.3423054>

1 INTRODUCTION

Tracking data is nowadays a crucial information for professional sport teams. There is a growing demand for automatic statistic collection and analysis of soccer matches. Along with event data, player tracking data is a fundamental information to build automatic action understanding.

In this work, we aim to construct a simple and efficient soccer player tracking system that can be quickly and easily used after capturing a game. It is able to detect and track even in challenging scenarios such as the frame of the video shown in Figure 1. In such a camera setting, where a single wide-angle fixed camera is placed roughly at the center of the field, all players are visible in the video, which makes possible to track all of them during the whole game with just a single camera. However, this kind of view renders very small and low-resolution players on the image, making both their detection and their individual identification for tracking highly challenging. An alternative could be to use 4K images, as in [51], in order to obtain higher-resolution players, but this solution is not always available and it leads to very time-consuming algorithms.

In order to realize tracking in videos, recent proposals use a frame-wise detection in a first step followed by an association step that links detections in consecutive frames. The effectiveness of those methods heavily relies on the accuracy of the object detector. Our proposed method follows this same strategy while being robust

to the aforementioned drawback: in a first stage we train, without any manual annotations, a soccer player detector capable of recognising small players and then we establish associations among the previous detections, again following a self-supervised approach.

Numerous Deep Neural Networks (DNN) architectures [32, 41, 42] for object detection have been widely studied in the literature and achieve remarkable results. A first straightforward approach to detect players would be to use a pretrained object detection model, trained on a general dataset, like COCO [31], and to extract the human-class detections. We show in the supplementary material an application of a Faster R-CNN ResNet-50 FPN model pretrained on COCO [31], applied to a typical image of a soccer game. When these models are applied to detect players in usual soccer games videos, their performance is strongly affected. The reasons are numerous and include soccer player singular appearances, small relative size of far away players, presence of numerous non-player humans (like spectators, coaches), motion blur or strong occlusions.

For all these reasons, DNN models for soccer player detection are usually supervised, i.e., trained on ground-truth player positions. However, the manual annotation of bounding boxes is costly and time-consuming. To the best of our knowledge, no dataset of annotated soccer games, in various illumination conditions and scenarios has been made publicly available. Thus, we opt for a self-supervised strategy.

Our goal is to automatically achieve a knowledge transfer from a pretrained general human detector to a soccer player detector, without any manual annotation on the soccer images. Our final model is able to detect very small players, avoiding other humans present on the image (like spectators or coaches). Our approach is end-to-end as it takes as input the image without any pre-processing step and outputs the player positions. Figure 1 illustrates a detection result using our proposal. As it can be observed, small-size players are successfully detected. Additionally, in the same figure we illustrate the tracking result using different colors for the video frame bounding boxes corresponding to different tracking identities. To summarize, the contributions of our work are :

- A soccer player detector and tracker for wide-angle video games.
- A training method that does not require ground-truth data and learns a detector and a tracker that are accurate and robust to small players.
- An ablation study with a thorough quantitative analysis of all the steps of the proposed training procedure and its benefits for the final proposed detector and tracker, including a comparison with state-of-the-art methods.

The outline of the paper is as follows. Section 2 reviews the related work. Section 3 presents the proposed method for detection and tracking, with the main details on the model and architecture (more detailed information can be found in the supplementary material). An experimental evaluation of the proposed method together with a comparison to previous works is given in Section 4. Finally, conclusions are drawn in Section 5.

2 RELATED WORK

In this section, we survey the main proposed methods and deep architectures for object detection and tracking, in particular for small-size objects, and their applications to sport analysis.

Object detection with deep learning. CNNs have considerably improved the accuracy on object detection. Two families of detectors are currently popular: First, region proposal based detectors R-CNN [17], Fast R-CNN [16], Faster R-CNN [42] and, second, detectors that directly predict boxes for an image in one step such as YOLO [41] and SSD [32]. The goal of region proposal based detectors is to find regions of interest in the image that might contain an object and extract CNN features to classify them. Faster R-CNN [42] merges region proposal and classification phases into a single and end-to-end model that is faster and more accurate than the previous one. Feature pyramid networks (FPN) [30] were built on top of R-CNN systems to enhance their capacity to detect objects of varying sizes. They exploit the inherent multi-scale, pyramidal hierarchy of the CNN backbone of these networks to construct a pyramid of features. R-CNN sub-networks are then built on these semantically-strong multi-scale backbones. However, FPN models still show limited performances for detecting small objects.

Self-supervised domain adaptation for Object Detection. Some works have developed strategies to tackle the absence of annotated data for object detection in a particular domain. For example, some approaches use synthetically generated training images [20, 21], which are not competitive in the context of soccer. Similar to us, [9] uses a distillation framework to transfer the knowledge of a teacher network pretrained on the source domain to a student on the target domain. This simple approach has the advantage of not requiring any annotated data in the target domain. [45] proposes to improve the distillation by adding additional temporal consistency information obtained via object tracking.

Small object detection. Several methods based on generating new and challenging data have recently been proposed to enhance the capacity to detect small objects. Some works propose small object data augmentation: [24] copy-pastes small objects and [32] directly performs a zoom out operation on training images. Other methods leverage generative adversarial strategies to generate either super-resolved feature representations [29] or super-resolved image patches [1]. Besides, several works (e.g. [18, 43]) use context information surrounding the object proposal to improve their understanding.

Soccer player detection. Many non deep learning methods have been developed to detect soccer players [36, 46]. Classical approaches can be divided into two groups: blob-based approaches relying on background subtraction [12, 35] and feature-based classification methods [22, 34]. Beyond, deep learning approaches have exhibited great capacity in render more robust methods in many contexts. Probably due to the lack of publicly available annotated data, most of the sport player detection algorithms use the mentioned pretrained deep learning human detector models without fine-tuning on the specific sport scenario [5, 48, 51]. Some methods like [25, 33] devise specific CNN architectures for player detection. To train their models, [33] manually annotates images from two soccer games and [25] uses these annotations as well as the annotated game given by [11]. Their models are trained and tested, respectively, on two and three games only. This is insufficient in order to capture the visual variability of soccer games with e.g. various team color, player sizes, different stadiums or weather conditions. A method for detecting small players in soccer applied on multi-view videos is proposed by

[51]. Notice, that we focus on a more challenging scenario for tracking with single-camera views. As mentioned, [9] performs online semantic segmentation without requiring any manual annotation. On each specific game, they train a fast student network with data annotated by a slower pretrained teacher. This strategy can achieve good results but it adds complexity and the need of GPU resources to fine-tune on each specific game. Also, they evaluate their model on a standard TV broadcast game without focusing on possible small player failure cases.

Multi-object tracking. Multiple Object Tracking (MOT) refers to the task of estimating trajectories over time of multiple objects (see [8, 28, 36] for a detailed review). As previously mentioned, one of the most popular strategies for MOT is tracking by detection which mainly consists of a first step where the objects are detected followed by an identification step where the same object is matched through time and joined in individual trajectories. For instance, [2] proposes to adapt a state-of-the-art detector network into a tracker under the assumption that the target objects slightly move between consecutive frames. Then, a re-identification step trained on ground-truth tracking data is added to improve the results. Other approaches use graph-based algorithms for the tracking step (e.g., [4, 27, 52]). The so-called single-shot methods aiming to jointly detect and identify the targets have recently gained a lot of attention [4, 13, 49, 54]. In [49] a single network integrates both the detector and an embedding model for identification. Similarly, [13] processes two consecutive frames and generates two-frame tracklets that are afterwards combined into multi-frame tracks. A group of MOT methods leverage pose estimation of humans for tracking purposes [15, 37, 40, 50]. However, tracking strategies based on human pose exhibit difficulties as they struggle in correctly estimating pose for low-resolution persons such as the small players that usually appear in real wide-angle soccer videos, as tackled in this work.

3 METHOD

We propose a self-supervised soccer player detector and tracker that can be applied to wide-angle video games since it is robust to small players. Our method is specialized in soccer and does not require any labeled data. Our pipeline follows two steps : first a soccer player detection and then a tracking based on the detection results. In Section 3.1, we describe the detection framework which consists of the successive training of a teacher and a student for domain adaptation and knowledge distillation. Then, in Section 3.2, we describe the tracking framework which uses the very same detector to build player trajectories using spatial and visual consistency.

3.1 Soccer player detection

Our first goal is to train an accurate player detector. To do so, we specialize an object detection model, trained on a general dataset, in the detection of soccer players, without any further manual annotations. From a set of non-annotated training soccer images, denoted by $\mathcal{X}$, we adopt the following three-step training process also illustrated in Figure 2 :

- (1) Label the set of training images $\mathcal{X}$ with a pretrained human detector denoted as the initial teacher network $\mathcal{T}^{init}$. This set of initial noisy annotations $\mathcal{T}^{init}(\mathcal{X})$, are automatically

Figure 2: Illustration of our soccer player detection method. Each column corresponds to one step of the proposed method explained in Section 3.1.

corrected by suppressing False Positives and adding False Negatives resulting in a new set of labels $\mathcal{Y}^c$.

- (2) Fine-tune the initial teacher network on these corrected annotations $\mathcal{Y}^c$ in order to get the final teacher network $\mathcal{T}^f$ and its corresponding set of annotations $\mathcal{T}^f(\mathcal{X})$.
- (3) Distill the knowledge of the teacher $\mathcal{T}^f$ to a smaller and faster student network $\mathcal{S}$ which is our final accurate detector.

The baseline object detector we use is the well-known Faster-RCNN [42] network with Feature Pyramid Network (FPN) [30]. It is widely used in the literature due to its robustness to various object sizes while keeping low inference time. We use, respectively, ResNet50 and ResNet18 backbones for the teacher and student networks. We want our whole domain adaptation process to stay general so that it can be easily applied to any other object detection model. Therefore, we do not apply hard changes to the original FPN Faster-RCNN architecture. We adapt some parameters to make the model more sensitive to small players while keeping a reasonably low inference time. More details about the parameter choices can be found in Section 4.4 and in the supplementary material.

3.1.1 Automatic self-labeling of the training images. Our training images $\mathcal{X}$ are unlabeled TV broadcast soccer videos taken from a subset of the SoccerNet dataset [14]. More details about the dataset can be found in Section 4.1. We extract from $\mathcal{T}^{init}(\mathcal{X})$ the human annotations with a confidence higher than $\sigma = 0.8$. These annotations are then post-processed by the following two steps.

Field detection. Most of the False Positives consist of detections of supporters and coaches. We use a combination of field and line detection to avoid them. A first field mask is computed through green filtering and contour detection. To detect the line we use the pix2pix line extraction model of [6]. An example can be found in the bottom-left corner of Figure 2. Additional details and results can be found in the supplementary material. Finally, all the detected

bounding boxes with a proportion of the area in the field mask lower than 0.3 are discarded. Note that in practice this post-processing is also used to correct the $\mathcal{T}^f(X)$ annotations.

Adding missed players to the training data. We also add missed detections to the training data through a blob detection strategy. This is realized via green filtering, contour detection and human detection in the regions of the contours with appropriate size. Detailed information can be found in the supplementary material.

3.1.2 Fine-tuning of the teacher network on the corrected annotations. We perform a fine-tuning of the teacher $\mathcal{T}^{init}$ network on the corrected annotations $\mathcal{Y}^c$. For training, we use the same loss as in [42]. Let us denote this detection loss by $\mathcal{L}_{detection}$. This first retraining step will enable the big teacher network to first adapt to the new domain before transmitting the knowledge to a smaller student network. As shown in Section 4, this procedure improves the final student accuracy.

3.1.3 Training of the student network.

Student $\mathcal{S}$ architecture. We propose one small modification of the student $\mathcal{S}$ architecture to make the model consider the context surrounding a detection before assigning it a bounding box. More precisely, we follow the idea proposed by [43] and incorporate the context region enclosing each proposal in the decision process. This is done by concatenating to the feature maps of an extracted region of interest bounding box, the feature maps of this same box enlarged by a factor of 4. As explained in [43], adding the contextual regions is likely to help the detection of small objects.

Down-scaling of the training images. As explained in the introduction, the goal of our object detector is to be robust to small players like the ones that appear in Figure 1. On this kind of views, the entire field is constantly visible in the image. However, our training data (e.g. Figure 2) is composed of TV broadcast images that are often zoomed on the current game action. The players thus appear substantively bigger on the screen. We also want our detector to be robust to varying player sizes. To do so, we perform data augmentation by re-scaling each training image by a random factor chosen between a given value and 1. Let us denote this value by min_scale , which verifies $0 < min_scale \leq 1$. We add zero-padding to the resulting images to make them match their original size.

To train the final student network $\mathcal{S}$, we extract from $\mathcal{T}^f(X)$ the annotations with a confidence higher than $\sigma = 0.8$, post-process them by field detection and finally randomly down-scale the input image with min_scale equal to 0.1.

3.2 Soccer player tracking framework

We propose a tracking approach to associate the previous detections between consecutive frames. Our method is unsupervised, accurate for small players, simple and fast. It is only based on our player detector model. Given a video sequence $\{u_t\}_t$ where t denotes time, we define a player trajectory as a finite list of ordered player bounding boxes $T_k = \{b_{t_1}^k, b_{t_2}^k, \dots\}$. Here we assume that $0 \leq t_1 < t_2 < \dots$. At each frame u_t , we compute the set of potential bounding boxes $\mathcal{B}_t = \{b_t^{k_1}, b_t^{k_2}, \dots\}$ by running our player detector $\mathcal{S}$ and gathering all the annotations $\mathcal{S}(u_t)$ with a confidence higher than $\sigma_{track} > 0$. To extract the player trajectories from these bounding

boxes, we propose to consider two successive steps, explained in the following subsections.

3.2.1 Spatial consistency association. We first rely on spatial consistency, i.e., if two bounding boxes from the current and the previous frames are adjacent we consider them to belong to the same target. Formally, at $t > 0$, for each player bounding box b_t^i in $\mathcal{B}_t$ if there is a unique bounding box b_{t-1}^j in $\mathcal{B}_{t-1}$ such that the intersection-over-union between these two bounding boxes satisfies $IOU(b_t^i, b_{t-1}^j) > \tau_{iou}$, these detections are associated to the same track. In case b_{t-1}^j is associated to two different players in $\mathcal{B}_{t-1}$, these associations are discarded. Once a bounding box at frame t is associated to a track, it is removed from $\mathcal{B}_t$.

However, as explained in [37] and [2], the above criterion is based on the assumption that the tracked target has significant overlap between two successive frames. This assumption is not satisfied in some cases such as in low frame rate, fast player movements, or intense camera shifts. Furthermore, in the presence of crowded scenarios with partial occlusions between players, the IOU measure is not satisfactory. For these reasons, we consider a complementary visual consistency measure.

3.2.2 Re-identification and visual consistency. The players from frame u_{t-1} non directly associated with those from frame u_t with the previous criterion are first considered to be lost. We store killed (deactivated) tracks for a fixed number of N_{reID} frames. We then use a visual embedding extracted from the player detector network to compare the visual appearances of the remaining bounding boxes in $\mathcal{B}_t$ with the deactivated tracks. Using an additional ResNet50 like [2] would be time consuming, so we prefer to use the player detector model to extract visual information. The advantage is that this information is already freely available for each detection.

We propose to keep the original player detector $\mathcal{S}$ architecture and to directly extract the embedding features given by ROI pooling of the region proposals. However, a fine-tuning of the network is needed to make the individual features discriminative enough between players. Let us remark that we found that modifying the architecture by adding in parallel to the ROI head, a Siamese "Tracking head" ([47]) did not improve the quality of the final extracted embedding. Our proposed method remains simple and enables to unify detection and association, without any change in the architecture of the detector.

In principle, this fine-tuning would require tracking ground-truth data. Due to the absence of such data in soccer, we propose to use the following unsupervised approach:

- (1) Use annotations $\mathcal{B}_t = \{b_t^{k_1}, b_t^{k_2}, \dots\}$ given by the teacher $\mathcal{T}^f$.
- (2) Generate trajectories $\{T_k\}$ only with the spatial consistency association method explained in Section 3.2.1.
- (3) Keep only long enough trajectories, i.e. those where $|T_k| > 5$ frames.

We fine-tune the student network $\mathcal{S}$ on this data with a triplet loss, using the *batch hard* strategy introduced in [19]. Given a small training sub-sequence of N consecutive frames, we randomly sample K player tracks, and then randomly sample T detections for each track. This ensemble of detections forms a batch of KT

samples. We define the triplet loss over this batch as :

$$\mathcal{L}_{triplet}(\theta) = \sum_{k=1}^K \sum_{t=1}^T \left[m + \max_{u=1 \dots T} \|f_{\theta}(b_t^k) - f_{\theta}(b_u^k)\|_2 - \min_{i=1 \dots K, u=1 \dots T, i \neq k} \|f_{\theta}(b_t^k) - f_{\theta}(b_u^i)\|_2 \right]_+ \quad (1)$$

where $m > 0$ denotes the margin (in practice, we use $m = 2$), $\|\tilde{f}\|_2 = \langle \tilde{f}, \tilde{f} \rangle^{1/2}$, for $\tilde{f} \in \mathbb{R}^M$, and $[x]_+$ denotes, as usual, the maximum between a real value x and 0. f_{θ} takes a bounding box as input and outputs the resulting identification vector. θ denotes the weights of the backbone and ROI pooling layers.

In parallel, we need to back-propagate the detection loss in order not to lose the ability of the network to detect players. The detection loss, $\mathcal{L}_{detection}$, is computed with the teacher network annotations (see Section 3.1), with respect to the whole network parameters $\tilde{\theta}$ and it is summed over the N frames of the working sub-sequence. The final training loss is $\mathcal{L}(\tilde{\theta}) = \mathcal{L}_{detection}(\tilde{\theta}) + \mathcal{L}_{triplet}(\theta)$.

Tracking inference. For inference, we match the remaining elements in $\mathcal{B}_t$ with the deactivated tracks by the Hungarian graph bipartite matching algorithm [26]. The distance between b_t^i and a killed track $T_k = \{b_{t_1}^k, b_{t_2}^k, \dots\}$ is averaged over the last elements of the track. We define the single association cost between b_t^i and one player $b_{t_j}^k$ of the track as a combination of visual and spatial distances :

$$C(b_t^i, b_{t_j}^k) = \alpha D_{visual}(b_t^i, b_{t_j}^k) + (1 - \alpha) D_{spatial}(b_t^i, b_{t_j}^k) \quad (2)$$

where $0 \leq \alpha \leq 1$,

$$D_{visual}(b_t^i, b_{t_j}^k) = \|f_{\theta}(b_t^i) - f_{\theta}(b_{t_j}^k)\|_2, \quad (3)$$

$$D_{spatial}(b_t^i, b_{t_j}^k) = \|x_c(b_t^i) - x_c(b_{t_j}^k)\|_2, \quad (4)$$

and $x_c(b)$ denotes the center point of a bounding box b . A player-track association is established under two additional conditions on the visual and spatial distances : $D_{spatial} < D_{spatial_max}$ and $D_{visual} < D_{visual_max}$.

4 EXPERIMENTS

4.1 Training data

We train our method on a subset of SoccerNet dataset [14]. In particular, we extract only wide-angle sequences where the game is filmed with a wide angle of a lateral view from a fixed camera placed, roughly, at the center of the lateral tribune. Examples of these images are shown Figure 2 and in the supplementary material. It includes 2321 randomly extracted 5 seconds sequences from 50 unlabeled TV broadcast soccer games of resolution 720×1280 and frame rate 25 fps. In total, it comprises 290125 images. Different stadiums, illumination conditions, teams, and viewpoints are included.

4.2 Testing data

In order to quantitatively evaluate our detection and tracking results, we need to assess them on ground-truth annotated data. For detection evaluation, the position bounding boxes of all the players (and referees) in the field are needed. We use three different player position datasets for evaluation:

Figure 3: Histograms of the relative bounding box areas of players for each dataset. The SoccerNet statistics are computed with annotations given by the teacher network.

Table 1: Statistics on different detection datasets used. "bbox area" refers to the average area of the ground-truth player bounding boxes. "relat. area" is the average bounding box area divided by the whole image area. In SoccerNet we provide estimations since there is no ground truth annotation.

Dataset	# images	image size	bbox area	relat. area ($\times 10^6$)
SoccerNet	290125	720×1280	≈ 2200	≈ 2300
TVsoccer	250	720×1280	2580	2300
ISSIA	2600	1080×1920	4420	2140
SPD	550	720×1280	921	950

- "TVsoccer" : we manually annotated 250 images with wide views extracted from different TV broadcasted games of [53].
- "ISSIA" : 2600 frames from ISSIA-CNR (annotated) dataset [11]. We only perform our evaluation on camera 3 because it is the one placed on the center of the field and thus its view contains most of the players.
- "SPD" : 550 wide view annotated frames extracted from [33].

Various example images from these datasets are shown in the supplementary material. Additionally, statistics about these datasets can be found in Table 1, in particular comparing the different sizes of the players. The fourth column shows the average player bounding box sizes. Moreover, as the first pre-processing step of the player detector is to rescale the input image to a fixed image size, we also display in the fifth column the relative area, i.e., the average bounding box area divided by the whole image area. Figure 3 also plots the histograms of the relative bounding box areas in each dataset. Let us notice that the "TVsoccer" and "ISSIA" have similar player sizes to the SoccerNet training data, and that "SPD" contains smaller players than the other datasets. In order to fit the average size of training data with the one of "SPD", a rescaling of the training images with a scaling factor of approximately 0.6 would be needed.

To evaluate the tracking performance, an individual identity annotation is needed. We use the ISSIA-CNR dataset [11] and test on 5 different sequences of 10 seconds at 25 fps from camera 3.

We also qualitatively evaluate our proposed detection and tracking methods on sequences extracted from [39] (an example can be found in Figure 1) where really small players usually appear.

4.3 Evaluation metrics

For evaluation of the player detection model, we use the standard Average Precision (AP) measure defined by the PASCAL VOC object detection challenge [38], at IOU threshold 0.5. This measure evaluates the area under the Precision-Recall curve. Therefore, it determines the capacity of the detector at any score confidence threshold. However, while annotating data, we also need to evaluate the performance of the annotation at a fixed score detection threshold. In this case we will use the F1 score. It is a metric that combines recall and precision by taking their harmonic mean, $F1 = \frac{precision \cdot recall}{precision + recall}$.

To measure the performance of a tracker, we focus on the Multiple Object Tracking Accuracy (MOTA) [23] and ID F1 Score (IDF1) [44]. MOTA and IDF1 are meaningful combinations of the number of false positives and negatives of the detector, the total number of identity switches, the mostly tracked and mostly lost metrics [44].

4.4 Implementation details

We conduct all our experiments on a GPU NVIDIA TITAN X. To train and evaluate the FPN-FasterRCNN models, we use as baseline the native PyTorch implementation¹. Most of the native parameters are unchanged from the implementation of [30] and [42]. The images are scaled to a fixed input size of 768×1344 as a first preprocessing step. The native NMS is replaced by Soft-NMS, as proposed by [3]. We use pretrained models provided by torchvision, specifically $\mathcal{T}^{init}$ on COCO [31] and initialization of $\mathcal{S}$ on ImageNet [10]. The rest of the parameters used for detection and tracking are given in the supplementary material.

4.5 Ablation study

4.5.1 Player detection results.

Automatic labeling of the training images. We first evaluate the proposed pre-labeling of the SoccerNet training images, described in Section 3.1.1. Let us remind that we first keep the pretrained network annotations $\mathcal{T}^{init}(X)$ with confidence bigger than 0.8. We then automatically correct them with field segmentation and player blob detection to get $\mathcal{Y}^c$. To understand its effect, we provide the statistics of the generated labels with the different methods in Table 2. With field segmentation, approximately 5% of the original annotations are removed and, with blob detection, another 5% are added. It mainly corresponds to, respectively, a reduction in false positives and an increase in true positives. As it can be observed in the second column of Table 2, the added players are, in average, smaller than the original ones. This observation supports our claim that a pretrained detector shows poor performance for small objects. We also show in Table 3 the F1-score performances, on the three testing datasets, of the models and techniques used to automatically annotate the training set.

Retraining of the teacher network on the corrected annotations. We use these corrected annotations $\mathcal{Y}^c$ to retrain the initial teacher detection model $\mathcal{T}^{init}$ and get the final teacher $\mathcal{T}^f$. We have tried to fine-tune the pretrained model or to train it from scratch, with only the backbone weights pretrained on ImageNet and obtained very similar results in both cases. Therefore, in order to keep consistency

Table 2: Statistics of the first automatic annotation process

Method	# detections	avg bbox area
Pretrained object detector	14.2/image	2390
+ Field segmentation	-5.2%	-
+ Blob detection	+5.1%	1770

Table 3: F1-score of automatic annotation processes.

Dataset	TVsoccer	ISSIA	SPD
Pretrained detection model $\sigma > 0.8$	89.0	91.4	79.0
+ Field segmentation + blob detection	90.1	91.8	82.0
Retrained detection model $\sigma > 0.7$	93.1	93.0	83.5

Table 4: AP performances of the retrained teacher network $\mathcal{T}^f$ compared to the COCO pretrained version $\mathcal{T}^{init}$. The different columns indicate the results for the downscaled testing dataset.

Dataset	TVsoccer			ISSIA			SPD
Rescale factor	$\times 1$	$\times 0.75$	$\times 0.5$	$\times 1$	$\times 0.75$	$\times 0.5$	$\times 1$
$\mathcal{T}^{init}$	89.3	78.2	78.0	90.3	86.6	77.8	82.2
$\mathcal{T}^f, min_{scale} = 1$	95.7	93.9	84.3	97.6	94.9	86.8	77.1
$\mathcal{T}^f, min_{scale} = 0.5$	95.4	96.5	94.0	97.5	97.7	96.0	92.0

with the training of other models, in all the following experiments, the teacher $\mathcal{T}^f$ and the student $\mathcal{S}$ have been trained from scratch, with ResNet50 and ResNet18 backbones pretrained on ImageNet. The training process is described in Section 4.4.

The goal of this experiment is to retrain the teacher network $\mathcal{T}^{init}$ so that it annotates as accurately as possible our original SoccerNet training images. A first model is basically retrained on the given corrected annotations. We have previously seen with Table 2 that the missed detections of the pretrained model were due to the small size of the targets. Therefore, we conduct a second training using the re-scaling data augmentation method (see Section 3.1.3) with $min_scale = 0.5$. Table 4 evaluates the retrained models on the three evaluation datasets. Let us remind that SPD contains smaller players than TVsoccer and ISSIA. In order to simulate smaller players in these two last datasets, we also show in the columns $\times 0.75$ and $\times 0.5$, respectively, the results of the model applied to down-scaled images at fixed scales of 0.75 and 0.5, respectively. Smaller scales are not explored since, as it can be observed in the histogram of Figure 3, at scale 0.5, we already match the smallest player size present in the training images.

Results in Table 4 show that the retraining on self-annotated soccer data enables to enhance, by a large margin, the capacity of our given network to detect soccer players. We successfully perform domain adaptation of the human detector to the specific context of a soccer game. Furthermore, we show the necessity to perform down-scaling data augmentation on our TV broadcast training images in order for the detector to be robust to small players : using random re-scaling of the training images with $min_scale = 0.5$ enables to almost close the gap of the detector performance between the

¹https://github.com/pytorch/vision/blob/master/torchvision/models/detection/faster_rcnn.py

Table 5: AP performances of the student network training. The different columns indicate the results for the down-scaled testing dataset.

Dataset	TVsoccer			
Rescale factor	$\times 1$	$\times 0.75$	$\times 0.5$	$\times 0.3$
$\mathcal{T}^f$	95.4	96.5	94.0	83.12
$\mathcal{S}$ trained with $\mathcal{Y}^c$, $\min_{scale} = 0.5$	95.9	96.3	94.3	83.8
$\mathcal{S}$ trained with $\mathcal{Y}^c$, $\min_{scale} = 0.1$	94.9	96.3	95.6	93.5
$\mathcal{S}$, $\min_{scale} = 0.1$	95.1	96.2	96.1	93.6
$\mathcal{S}$, $\min_{scale} = 0.1$ w. context	95.2	96.5	96.4	93.5

Dataset	ISSIA			
Rescale factor	$\times 1$	$\times 0.75$	$\times 0.5$	$\times 0.3$
$\mathcal{T}^f$	97.5	97.7	96.0	87.2
$\mathcal{S}$ trained with $\mathcal{Y}^c$, $\min_{scale} = 0.5$	97.6	97.3	96.0	88.3
$\mathcal{S}$ trained with $\mathcal{Y}^c$, $\min_{scale} = 0.1$	96.8	96.7	97.1	94.3
$\mathcal{S}$, $\min_{scale} = 0.1$	96.6	96.9	97.4	95.7
$\mathcal{S}$, $\min_{scale} = 0.1$ w. context	97.3	97.7	97.9	95.9

Dataset	SPD			
Rescale factor	$\times 1$	$\times 0.75$	$\times 0.5$	
$\mathcal{T}^f$	92.0	88.7	83.5	
$\mathcal{S}$ trained with $\mathcal{Y}^c$, $\min_{scale} = 0.5$	96.0	95.1	91.3	
$\mathcal{S}$ trained with $\mathcal{Y}^c$, $\min_{scale} = 0.1$	96.6	96.2	96.2	
$\mathcal{S}$, $\min_{scale} = 0.1$	97.7	97.7	97.5	
$\mathcal{S}$, $\min_{scale} = 0.1$ w. context	98.0	97.8	97.9	

smallest (scale 0.5) and the biggest (scale 1) players without losing performance in bigger scales.

Knowledge distillation to the student network. Assessment on our final detector is now presented. We use the annotations generated by the previously retrained teacher $\mathcal{T}^f$ to train the student network $\mathcal{S}$. Once again, we train it from scratch, using the ResNet18 backbone pretrained on ImageNet. We follow the same training process described in Section 4.4. We train with random down-scaling of the training images, but this time with $\min_{scale} = 0.1$ in order to get a detector as robust as possible to small players. We also show in Tables 5 the results using the context feature map concatenation method introduced in Section 3.1.3. The AP performances of these trainings are reported, for each evaluation dataset, and for down-scaling factors $\{1, 0.75, 0.5, 0.3\}$ in the third block of lines in each subtable of Table 5. *SPD* is not evaluated at scale 0.3 due to the already small size of the players on the original images. For comparison, we also train $\mathcal{S}$ like the previous teacher on the original corrected annotations $\mathcal{Y}^c$ (results in the second block of lines in each subtable of Table 5). For the sake of comparison, we also report the performance of the teacher in the first blocks of Table 5.

The following observations can be extracted :

i) The first two lines of results in each subtable show that ResNet18 and ResNet50 backbones bring already similar results on *ISSIA* and *TVsoccer* when retrained on the original corrected annotations, while the ResNet18 backbone outperforms ResNet50 by a large margin on the *SPD* dataset. We hypothesize this difference by the fact that there exists a difference in the appearance of small players

present in the *SPD* dataset and small players "simulated" by down-scaling the SoccerNet image : as players in *SPD* are filmed from a higher point of view than most of the shots in SoccerNet, *TVsoccer* and *ISSIA*. It is very likely that a Resnet50 backbone, with higher complexity, slightly over-fits on the SoccerNet player appearances when the Resnet18 version, with less parameters, better generalizes to new soccer player appearances. We add additional details about this observation in the supplementary material.

ii) Using $\min_{scale} = 0.1$ enables a substantial boost of the performance for small players. This is at the cost of a little loss of accuracy for the biggest players (scale 1) of *ISSIA* and *TVsoccer*. As we want our final detector to target small players like in *SPD*, we keep $\min_{scale} = 0.1$.

iii) As expected, training the student $\mathcal{S}$ on the retrained teacher annotations $\mathcal{T}^f(X)$ rather than on the corrected labels $\mathcal{Y}^c$ obtained with the pretrained one improves the results, especially on small players. The intermediate retraining of the teacher permits a first domain adaptation and enables to train the final student $\mathcal{S}$ on more accurate labels.

iv) Finally, using context information improves the detection performances at all scales.

Figures 1 and 4 show visual results of the student detection, in particular with very small players. Additional examples are shown in the supplementary material.

4.5.2 Player tracking results.

Fine-tuning of the visual embedding. We first associate all the detections to tracks using only the spatial consistency IOU criteria. In this particular setting, we find on the validation sequences that the optimal IOU threshold value is $\tau_{IOU} = 0.7$. We use this value to create tracks on our unlabeled SoccerNet training sequences. As a post-processing step, we also remove all the tracks with length smaller than 5 detections since such small tracks are not reliable. The generated trajectories are used to fine-tune the object detector network so as to generate inter-player discriminative visual representations. The training is conducted using the triplet loss (1) with a margin value $m = 2$, $K = 5$ and $P = 10$.

Inference of the final tracker. We report the final results of the tracker in Table 6. In order to simulate the performances with smaller players, we evaluate the tracking performances at different fixed down-scaling factors. For a deeper analysis, we provide three blocks of results. The first one ("No triplet loss FT") displays the tracking performances before triplet loss fine-tuning of the network, and the second one ("Final Tracker") after this retraining step. Notice that it gives a small improvement in MOTA possibly confirming that the performance of the tracker mostly relies on the accuracy of the player bounding box positions. To confirm it, the third block, denoted as "Oracle tracker", shows the performance of the proposed tracking framework applied to ground-truth player positions. We observe that with ideal bounding boxes, our tracking method follows all the players almost perfectly. The performance gap between the final tracker and the oracle one is mainly explained by the limits of our detection method in occlusion situations. We illustrate in Figure 4 the results of our tracker on two sequences with small players. Additional examples are shown in the supplementary material, including failure cases and a discussion.

Figure 4: Visual performances of the proposed tracker on two challenging sequences (1 second apart frames).

4.6 Comparison with other methods

4.6.1 Soccer player detection methods. Our whole framework can be easily adapted for any given object detection model. The improvements showed in Section 4.5 between the pretrained model and our proposed final one prove the efficiency of the proposed method in the case of FPN-FasterRCNN. Due to unavailability of data or code, it is hard to compare our method with most of the detection and tracking frameworks specialized in soccer. However, we can still compare with [25] and [33] using the results given in their respective papers. Both methods proposed CNN architectures trained in the very same dataset on which they perform evaluation. On the opposite, our method is robust to evaluation on various game conditions. [33] does not use standard object detection metrics. We can still compare the precision of their model : on the SPD dataset they announce a precision value of 0.962 while our model achieves a final precision of 0.97. Table 7 shows that we also improve by large margin the AP performance of [25].

4.6.2 Multi-object tracking methods. We compare in Table 8 our tracker with the state-of-the art Multi-object tracking (MOT) methods [49] and [2]. Like us, [49] uses a triplet loss for learning appearance embedding and a network architecture based on Feature Pyramid Networks. [2] is also based on FasterRCNN and associates the detector bounding boxes using spatial information and visual embeddings. As we are interested in the tracking of small players, we report in the first two rows of results of Table 8 the results of their whole frameworks on the ISSIA dataset with $\times 0.5$ down-scaled images. The third column "FPS" indicates the computational time for each method in frame per second. Our proposal outperforms by very large margin these methods. However, even if our method is not trained on any ground-truth data it has been adapted to small soccer player detection. For a fairer comparison, we also compute the tracking results of [2] but from the detections given by our detector $\mathcal{S}$. These results, shown in the third row of Table 8, indicated by "tracking_wbw with $\mathcal{S}$ ", display a clear boost on their performance. Our proposed tracking method and [2] obtain really close results on the detections given by our $\mathcal{S}$ detector, corroborating the benefits of our detector. Our approach is however simpler, as we do not use any additional re-id network and is computationally two times faster. Moreover, our method is unsupervised while the tracking extension (re-id module) in [2] has been trained with ground truth tracking data. Finally, let us remark that the combination of our detector with the architecture of [49] is not possible as their model integrates detection and tracking in a joint estimation.

Table 6: MOTA and IDF1 performance of our final MOT on ISSIA dataset, for different down-scaling factors. Extra blocks of results provide a further analysis (details in the text).

Method	Scaling factor	MOTA	IDF1
No triplet loss FT	1.	87.3	74.6
	0.75	92.3	82.8
	0.5	93.1	81.7
Final tracker	1.	87.7	76.6
	0.75	93.4	82.3
	0.5	92.9	80.9
Oracle tracker	1	99.9	99.3
	0.75	99.7	95.0
	0.5	100.	96.8

Table 7: AP comparison of the proposed detector with the soccer player detection method [25].

Dataset	ISSIA	SPD
FootAndBall [25]	92.1	88.5
Ours	97.3	98.0

Table 8: IDF1, MOTA and running time (frames per second, FPS) comparison of the proposed MOT with the methods [49] and [2]. Best results in bold and second best underlined.

Method	MOTA	IDF1	FPS
towards_rt_mot [49]	46.4	24.7	13.1
tracking_wbw [2]	23.5	27.8	3.4
tracking_wbw with our detections	<u>91.0</u>	85.7	3.4
Ours	92.9	<u>80.9</u>	<u>6.5</u>

5 CONCLUSION

This paper presents a deep-learning-based soccer player detection and tracking method robust to small players and assorted conditions (such as different stadiums, teams, viewpoints and illumination settings). As a key component, our method is trained in an unsupervised way, so no manual annotations are needed. We have performed a thorough quantitative analysis of the different components of the method while highlighting the key ones. The comparison with state-of-the-art methods, both in detection and tracking, shows that our method offers competitive results, especially in the detection stage. As future work, we plan to improve the tracking results in challenging cases such as long-time occlusions.

ACKNOWLEDGMENTS

SH acknowledges support by ENS Paris-Saclay. CB and GH acknowledge support by MICINN/FEDER UE project, ref. PGC2018-098625-B-I00, H2020-MSCA-RISE-2017 project, ref. 777826 NoMADS, and RED2018-102511-T. We thank Patricia Vitoria for her comments that greatly improved the manuscript.

REFERENCES

- [1] Yancheng Bai, Yongqiang Zhang, Mingli Ding, and Bernard Ghanem. 2018. Sodmtgan: Small object detection via multi-task generative adversarial network. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 206–221.
- [2] Philipp Bergmann, Tim Meinhardt, and Laura Leal-Taixe. 2019. Tracking without bells and whistles. In *Proceedings of the IEEE international conference on computer vision*. 941–951.
- [3] Navaneeth Bodla, Bharat Singh, Rama Chellappa, and Larry S Davis. 2017. Soft-NMS—improving object detection with one line of code. In *Proceedings of the IEEE international conference on computer vision*. 5561–5569.
- [4] Guillem Brasó and Laura Leal-Taixé. 2020. Learning a neural solver for multiple object tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6247–6257.
- [5] Matija Buric, Miran Pobar, and Marina Ivašić-Kos. 2018. Object detection in sports videos. 1034–1039. <https://doi.org/10.23919/MIPRO.2018.8400189>
- [6] Jianhui Chen and James J Little. 2019. Sports camera calibration via synthetic data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. 0–0.
- [7] Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, and Jian Sun. 2018. Cascaded pyramid network for multi-person pose estimation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 7103–7112.
- [8] Gioele Ciaparrone, Francisco Luque Sánchez, Siham Tabik, Luigi Troiano, Roberto Tagliaferri, and Francisco Herrera. 2020. Deep learning in video multi-object tracking: A survey. *Neurocomputing* 381 (2020), 61–88.
- [9] Anthony Cioppa, Adrien Deliege, Maxime Istasse, Christophe De Vleeschouwer, and Marc Van Droogenbroeck. 2019. ARTHuS: Adaptive Real-Time Human Segmentation in Sports Through Online Distillation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- [10] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. 2009. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*.
- [11] T. D’Orazio, M. Leo, N. Mosca, P. Spagnolo, and P. L. Mazzeo. 2009. A Semi-automatic System for Ground Truth Generation of Soccer Video Sequences. In *2009 Sixth IEEE International Conference on Advanced Video and Signal Based Surveillance*. 559–564.
- [12] T. D’Orazio, M. Leo, P. Spagnolo, P. L. Mazzeo, N. Mosca, M. Nitti, and A. Distant. 2009. An Investigation Into the Feasibility of Real-Time Soccer Offside Detection From a Multiple Camera System. *IEEE Transactions on Circuits and Systems for Video Technology* 19, 12 (2009), 1804–1818.
- [13] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. 2017. Detect to track and track to detect. In *Proceedings of the IEEE International Conference on Computer Vision*. 3038–3046.
- [14] Silvio Giancola, Mohieddine Amine, Tarek Dghaily, and Bernard Ghanem. 2018. SoccerNet: A Scalable Dataset for Action Spotting in Soccer Videos. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)* (2018), 1792–179210.
- [15] Rohit Girdhar, Georgia Gkioxari, Lorenzo Torresani, Manohar Paluri, and Du Tran. 2018. Detect-and-track: Efficient pose estimation in videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 350–359.
- [16] Ross Girshick. 2015. Fast R-CNN. *CoRR* abs/1504.08083 (2015). http://www.cv-foundation.org/openaccess/content_iccv_2015/papers/Girshick_Fast_R-CNN_ICCV_2015_paper.pdf
- [17] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. 2014. Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation. In *Proceedings of the 2014 IEEE Conference on Computer Vision and Pattern Recognition (CVPR ’14)*. IEEE Computer Society, USA, 580–587. <https://doi.org/10.1109/CVPR.2014.81>
- [18] Georgia Gkioxari, Ross Girshick, and Jitendra Malik. 2015. Contextual action recognition with r* cnn. In *Proceedings of the IEEE international conference on computer vision*. 1080–1088.
- [19] Alexander Hermans, Lucas Beyer, and Bastian Leibe. 2017. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737* (2017).
- [20] Han-Kai Hsu, Chun-Han Yao, Yi-Hsuan Tsai, Wei-Chih Hung, Hung-Yu Tseng, Maneesh Singh, and Ming-Hsuan Yang. 2020. Progressive domain adaptation for object detection. In *The IEEE Winter Conference on Applications of Computer Vision*. 749–757.
- [21] Naoto Inoue, Ryosuke Furuta, Toshihiko Yamasaki, and Kiyoharu Aizawa. 2018. Cross-domain weakly-supervised object detection through progressive domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 5001–5009.
- [22] Zdravko Ivankovic, Milos Rackovic, and Miodrag Ivkovic. 2013. Automatic player position detection in basketball games. *Multimedia Tools and Applications* 72 (2013), 2741–2767.
- [23] Rangachar Kasturi, Dmitry Goldgof, Padmanabhan Soundararajan, Vasant Manohar, John Garofolo, Rachel Bowers, Matthew Boonstra, Valentina Korzhova, and Jing Zhang. 2008. Framework for performance evaluation of force, text, and vehicle detection and tracking in video: Data, metrics, and protocol. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 31, 2 (2008), 319–336.
- [24] Mate Kisantal, Zbigniew Wojna, Jakub Murawski, Jacek Naruniec, and Kyunghyun Cho. 2019. Augmentation for small object detection. *arXiv preprint arXiv:1902.07296* (2019).
- [25] Jacek Komorowski, Grzegorz Kurzejamski, and Grzegorz Sarwas. 2020. FootAndBall: Integrated player and ball detector. In *VISIGRAPP*.
- [26] Harold W Kuhn. 1955. The Hungarian method for the assignment problem. *Naval research logistics quarterly* 2, 1-2 (1955), 83–97.
- [27] Laura Leal-Taixé, Cristian Canton-Ferrer, and Konrad Schindler. 2016. Learning by tracking: Siamese CNN for robust target association. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*. 33–40.
- [28] Laura Leal-Taixé, Anton Milan, Konrad Schindler, Daniel Cremers, Ian Reid, and Stefan Roth. 2017. Tracking the trackers: an analysis of the state of the art in multiple object tracking. *arXiv preprint arXiv:1704.02781* (2017).
- [29] Jianan Li, Xiaodan Liang, Yunchao Wei, Tingfa Xu, Jiashi Feng, and Shuicheng Yan. 2017. Perceptual generative adversarial networks for small object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1222–1230.
- [30] Tsung-Yi Lin, Piotr Dollar, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. 2017. Feature Pyramid Networks for Object Detection. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [31] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollar, and Larry Zitnick. 2014. Microsoft COCO: Common Objects in Context. In *ECCV (eccv ed.)*. European Conference on Computer Vision. <https://www.microsoft.com/en-us/research/publication/microsoft-coco-common-objects-in-context/>
- [32] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott E. Reed, Cheng-Yang Fu, and Alexander C. Berg. 2016. SSD: Single Shot MultiBox Detector. In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 9905)*, Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling (Eds.). Springer, 21–37. https://doi.org/10.1007/978-3-319-46448-0_2
- [33] Keyu Lu, Jianhui Chen, James J. Little, and Hangen He. 2017. Light Cascaded Convolutional Neural Networks for Accurate Player Detection. In *British Machine Vision Conference 2017, BMVC 2017, London, UK, September 4-7, 2017*. BMVA Press. <https://www.dropbox.com/s/ydsw8sfq8t3fa/0171.pdf?dl=1>
- [34] Sławomir Mackowiak. 2013. Segmentation of Football Video Broadcast. *International Journal of Electronics and Telecommunications* 59 (2013), 75–84.
- [35] M. Manafifard, H. Ebadi, and H. Abrishami Moghaddam. 2017. Multi-Player Detection in Soccer Broadcast Videos Using a Blob-Guided Particle Swarm Optimization Method. *Multimedia Tools Appl.* 76, 10 (May 2017), 12251–12280. <https://doi.org/10.1007/s11042-016-3625-6>
- [36] Mehrtaash Manafifard, Hamid Ebadi, and H Abrishami Moghaddam. 2017. A survey on player tracking in soccer videos. *Computer Vision and Image Understanding* 159 (2017), 19–46.
- [37] Guanghan Ning, Jian Pei, and Heng Huang. 2020. Lighttrack: A generic framework for online top-down human pose tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. 1034–1035.
- [38] Rafael Padilla, Sergio L Netto, and Eduardo AB da Silva. 2020. A Survey on Performance Metrics for Object-Detection Algorithms. In *2020 International Conference on Systems, Signals and Image Processing (IWSSIP)*. IEEE, 237–242.
- [39] Svein Arne Pettersen, Dag Johansen, Håvard Johansen, Vegard Berg-Johansen, Vamsidhar Reddy Gaddam, Asgeir Mortensen, Ragnar Langseth, Carsten Griwodz, Håkon Kvale Stensland, and Pål Halvorsen. 2014. Soccer Video and Player Position Dataset. In *Proceedings of the 5th ACM Multimedia Systems Conference* (Singapore, Singapore) (MMSys ’14). Association for Computing Machinery, New York, NY, USA, 18–23.
- [40] Nan Ran, Longteng Kong, Yunhong Wang, and Qingjie Liu. 2019. A robust multi-athlete tracking algorithm by exploiting discriminant features and long-term dependencies. In *International Conference on Multimedia Modeling*. Springer, 411–423.
- [41] Joseph Redmon, Santosh Kumar Divvala, Ross B. Girshick, and Ali Farhadi. 2015. You Only Look Once: Unified, Real-Time Object Detection. *CoRR* abs/1506.02640 (2015). [arXiv:1506.02640](http://arxiv.org/abs/1506.02640) <http://arxiv.org/abs/1506.02640>
- [42] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. In *Advances in Neural Information Processing Systems* 28, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (Eds.). Curran Associates, Inc., 91–99. <http://papers.nips.cc/paper/5638-faster-r-cnn-towards-real-time-object-detection-with-region-proposal-networks.pdf>
- [43] Yun Ren, Changren Zhu, and Shunping Xiao. 2018. Small object detection in optical remote sensing images via modified faster R-CNN. *Applied Sciences* 8, 5 (2018), 813.
- [44] Ergys Ristani, Francesco Solera, Roger Zou, Rita Cucchiara, and Carlo Tomasi. 2016. Performance measures and a data set for multi-target, multi-camera tracking. In *European Conference on Computer Vision*. Springer, 17–35.

- [45] Aruni RoyChowdhury, Prithvijit Chakrabarty, Ashish Singh, SouYoung Jin, Huaizu Jiang, Liangliang Cao, and Erik Learned-Miller. 2019. Automatic adaptation of object detectors to new domains using self-training. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 780–790.
- [46] H. Shih. 2018. A Survey of Content-Aware Video Analysis for Sports. *IEEE Transactions on Circuits and Systems for Video Technology* 28, 5 (2018), 1212–1231.
- [47] Bing Shuai, Andrew G Berneshawi, Davide Modolo, and Joseph Tighe. 2020. Multi-Object Tracking with Siamese Track-RCNN. *arXiv preprint arXiv:2004.07786* (2020).
- [48] R. Theagarajan, F. Pala, X. Zhang, and B. Bhanu. 2018. Soccer: Who Has the Ball? Generating Visual Analytics and Player Statistics. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. 1830–18308.
- [49] Zhongdao Wang, Liang Zheng, Yixuan Liu, and Shengjin Wang. 2019. Towards real-time multi-object tracking. *arXiv preprint arXiv:1909.12605* (2019).
- [50] Bin Xiao, Haiping Wu, and Yichen Wei. 2018. Simple baselines for human pose estimation and tracking. In *Proceedings of the European conference on computer vision (ECCV)*. 466–481.
- [51] Jianfeng Xu, Lertniphonphan Kanokphan, and Kazuyuki Tasaka. 2018. Fast and Accurate Object Detection Using Image Cropping/Resizing in Multi-View 4K Sports Videos. In *Proceedings of the 1st International Workshop on Multimedia Content Analysis in Sports (Seoul, Republic of Korea) (MMSports'18)*. Association for Computing Machinery, New York, NY, USA, 97–103. <https://doi.org/10.1145/3265845.3265852>
- [52] Bo Yang and Ram Nevatia. 2012. An online learned CRF model for multi-target tracking. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2034–2041.
- [53] Junqing Yu, Aiping Lei, Zikai Song, Tingting Wang, Hengyou Cai, and Na Feng. 2018. Comprehensive Dataset of Broadcast Soccer Videos. *2018 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)* (2018), 418–423.
- [54] Yifu Zhan, Chunyu Wang, Xinggang Wang, Wenjun Zeng, and Wenyu Liu. 2020. A Simple Baseline for Multi-Object Tracking. *arXiv preprint arXiv:2004.01888* (2020).

A TRAINING DATA

Our training images are unlabeled TV broadcast soccer videos taken from a subset of the SoccerNet dataset, where we extract only wide views. Figure 5 and the left column of Figure 6 show some examples of these training images.

Figure 5: Examples of SoccerNet training images.

B DETAILS OF THE METHOD

In this section, we detail the algorithms used to remove false positives and add false negatives detections to the teacher $\mathcal{T}^{init}$ annotations. These methods have to be as accurate as possible and can be, to a certain extent, computationally expensive as they are not used for the final inference of the student $\mathcal{S}$.

B.1 Field detection

As a first step for detecting the field we find a mask that contains the potential pixels of the field. Then, it is successively refined. As observed in Figure 5, coaches are often present in the green bottom part of the field, and those should not be included in the mask. The field mask is computed thanks to following successive processing steps :

- Application of a green filter to create a first binary image. We use as upper and lower color threshold the values (15,50,50) and (70,255,255).
- Selection of the contour with the biggest connected component.
- Approximation of the contour by a polygon.
- Detection of the bottom line of the field and removal of the mask pixels below it.

To detect the bottom line we use the [6] pix2pix model to get a binary image approximating the lines of the field. We then run a Hough line transform on this binary image and use thresholds on the length, angle and position of the detected lines. Some examples of the final estimated field mask are shown in Figure 6.

B.2 Add missed players to the training data

We detail here the blob detection strategy used to add players to the training data. Given an input image u , we work on the pixels belonging to the previously extracted field mask. Note that the teacher $\mathcal{T}^{init}$ has already detected some of the players $\mathcal{T}^{init}(u)$.

We first compute blobs through green filtering and contour selection. We select the contours for which the enclosing bounding box verifies : i) it has a similar area as the bounding boxes of the players already detected by the teacher $\mathcal{T}^{init}$ and ii) it does not exist any detection in $\mathcal{T}^{init}(u)$ with intersection-over-union $IOU < 0.2$.

We then run a human pose estimation algorithm [7] on the corresponding enlarged extracted regions. If the confidence of the

Figure 6: Three results obtained by the proposed field detection approach. Left: original image. Right: Output of our field detector.

pose detection is higher than 0.8, it is considered as a player and it is added to the annotations. We have observed in our experiments that the pretrained version of the human pose estimator [7] performs better on small humans than the multi-person FPN-FasterRCNN model, $\mathcal{T}^{init}$.

We display in Figure 7 some examples of the training images with corrected annotations $\mathcal{Y}^c$.

Figure 7: Examples of training images with the annotations of $\mathcal{Y}^c$. In white the annotations of the pretrained teacher $\mathcal{T}^{init}$ and in green the added players detected with the strategy detailed in Section B.2 to include missing players.

C IMPLEMENTATION DETAILS

We indicate here additional parameter values involved in the proposed framework :

- i) We train the teacher and student detectors during 20 epochs using the SGD optimizer with an initial learning rate of 10^{-3} decayed with a factor 0.1 at epochs 8, 14 and 17.
- ii) The fine-tuning with triplet loss is done on 5 epochs with initial learning rate of 10^{-4} .
- iii) For tracking inference we find the optimal parameters to be :

Table 9: AP on SPD of fine-tuned $\mathcal{T}^f$ and S before and after fine-tuning on SPD images.

Model	$\mathcal{T}^f$	S
Before fine-tuning	92.0	96.0
After fine-tuning	99.3	98.9

$N_{reID} = 10$; $\alpha = 0.03$; $D_{visual_max} = 4$ and $D_{spatial_max} = (1/16) * image_width$.

D QUANTITATIVE RESULTS

In Section 4.5.1 of our paper, Table 5 shows that the Resnet18 backbone outperforms Resnet50 on the SPD dataset. We made the hypothesis that it is due to the fact that there exists a small difference between the appearance of the players in SPD and SoccerNet training images.

In Table 9 we propose to fine-tune $\mathcal{T}^f$ (Resnet50 backbone) and S (Resnet18 backbone) on 80% of the SPD dataset. We then evaluate on the remaining 20%. With this fine-tuning it is now the Resnet50 backbone that outperforms its Resnet18 counterpart. This test confirms that the Resnet50 backbone slightly overfits on the typical TV broadcast player appearances of the training data. Besides providing a faster inference, using a smaller backbone allows the student S to better generalize.

E VISUAL RESULTS

E.1 Player detection

We show in Figure 9 additional player detection results on the SPD (row 1), $ISSIA$ (row 2) and TV_soccer (row 3) datasets, as well as on the *panorama* images of [39] (row 4). In order to test the detection on very small players, we down-scale and pad the input images. We show that our network is able to detect almost all the players in very challenging images, even when a player contains only a few pixels.

To show the strength of our network we present in Figure 8 results with the pretrained FPN-FasterRCNN, $\mathcal{T}^{init}$, on the same challenging images. It fails at finding small players and detects a lot of non-player humans.

Failure cases. Figure 9 includes some failure cases. For example, in the second image of SPD (row 1), non-players are detected on the side of the field. In this particular case, the distinction between players and non-player is challenging because the latter appear on the field. The network can also fail when the background around a player is not green field, like in the third image of $ISSIA$ (row 2).

E.2 Player tracking

Figure 10 shows two tracking results on the most challenging SPD and *panorama* sequences. We present four images of each sequence, one every 5 frames, from top to bottom.

We observe that almost all the players are correctly followed. One failure case appears in the first column, in a typical crowded scene. Due to occlusion, the white player 20 is lost between the frames 2 and 3. The re-ID module fails to re-identify him at frame 4 and it is assigned a new id 21. The reason is that the appearance of the

Figure 8: Player detections results with the human annotations of FPN-FasterRCNN pretrained on COCO ($\mathcal{T}^{init}$).

player between frame 2 and 4 has changed substantially. Moreover, the bounding box given by the tracker is not really precise for a small player, it can involve only one part of the player, or its close neighbours. In these cases, the visual embedding extracted from this bounding box is not good enough to perform an association.

(a) *SPD dataset*

(b) *ISSIA dataset with down-scaling $\times 0.5$*

(c) *TV_soccer dataset with down-scaling $\times 0.5$*

(d) *images from dataset [39]*

Figure 9: Player detection results on different datasets, for different down-scaling levels.

Figure 10: Player tracking results on a *SPD* sequence and a *panorama* sequence.