

# Extraction of Positional Player Data from Broadcast Soccer Videos

Jonas Theiner<sup>1</sup>      Wolfgang Gritz<sup>1</sup>      Eric Müller-Budack<sup>2</sup>  
 Robert Rein<sup>3</sup>      Daniel Memmert<sup>3</sup>      Ralph Ewerth<sup>1,2</sup>

<sup>1</sup>L3S Research Center, Leibniz University Hannover, Hannover, Germany

<sup>2</sup>TIB – Leibniz Information Centre for Science and Technology, Hannover, Germany

<sup>3</sup>Institute of Exercise and Sport Informatics, German Sport University Cologne

{theiner, ewerth}@l3s.de      {r.rein, d.memmert}@dshs-koeln.de

## Abstract

*Computer-aided support and analysis are becoming increasingly important in the modern world of sports. The scouting of potential prospective players, performance as well as match analysis, and the monitoring of training programs rely more and more on data-driven technologies to ensure success. Therefore, many approaches require large amounts of data, which are, however, not easy to obtain in general. In this paper, we propose a pipeline for the fully-automated extraction of positional data from broadcast video recordings of soccer matches. In contrast to previous work, the system integrates all necessary sub-tasks like sports field registration, player detection, or team assignment that are crucial for player position estimation. The quality of the modules and the entire system is interdependent. A comprehensive experimental evaluation is presented for the individual modules as well as the entire pipeline to identify the influence of errors to subsequent modules and the overall result. In this context, we propose novel evaluation metrics to compare the output with ground-truth positional data.*

## 1. Introduction

Match analysis in soccer is very complex and many different factors can affect the outcome of a match. The question is which so-called key performance parameters allow for the characterization of successful teams [27, 34, 40, 44]. While team behavior can be differentiated into a hierarchical scheme consisting of individual, group, and team tactics, different metrics are necessary to capture behavior at each level [13, 40]. Researchers have recognized that game plays should be segmented into different phases since tactics vary greatly [31] across them. Performance in soccer is

also determined by physiological factors [10] such as running distance [8]. For this reason, it has been suggested to link such information to tactical parameters [3].

To carry out such analyses, the player positions on the field are required. Current tracking technologies allow the recording of several million data points representing player and ball positions during a match by using additional hardware, e.g., multiple static cameras or sensors on players. However, they are difficult to obtain, for instance, due to licensing, financial restrictions, or competitive concerns, i.e., a club normally does not want or disclose its own team's data. In contrast, broadcast video recordings of soccer matches can be accessed more easily. In this paper, we introduce a modular pipeline to extract the two-dimensional positions of the visible players from ordinary broadcast recordings. As illustrated in Figure 1, the system involves sports field registration, shot boundary detection, shot type classification, player detection, and team assignment.

**Application Novelty:** While commercial approaches like [51, 52, 56] primarily use multiple static cameras for position data generation from video data, the TV feed is concretely used by *SkillCorner* [49] and *Track160* [57]. However, only the final output of such systems is accessible [49, 57]. To the best of our knowledge, neither their quality, nor used architectures or even information about training data and applicability to own data is publicly reported. While individual sub-tasks were tackled in research, its combination for the joint real-world task of *player position estimation* has not been studied yet (also not beyond soccer). Even individual sub-modules have not been sufficiently evaluated in terms of applicability to real-world data. For the essential step of sports field registration, recent approaches [37, 46] are evaluated only on a single small-scale dataset [16]. Potential for generalization were mentioned [5, 37] with the use of many cost-intensive annota-

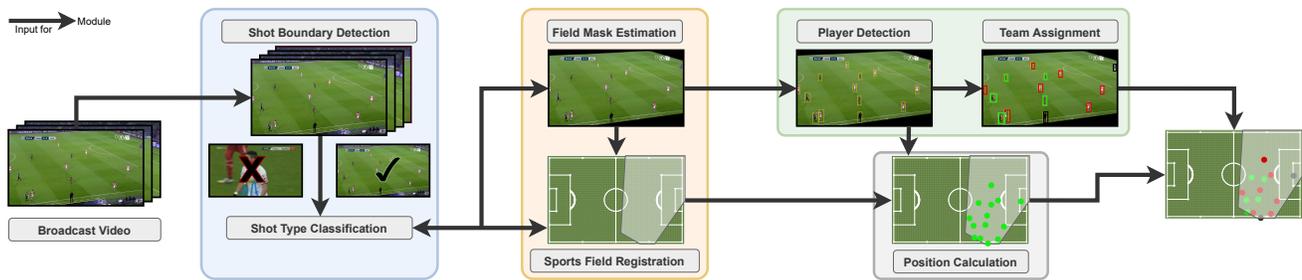


Figure 1. Proposed pipeline to extract positional data with team assignment from broadcast videos: The video is pre-processed to segment the field and detect shot boundaries. The camera type is estimated to extract shots from the main camera. Subsequently, the sports field is registered and the extracted homography matrix is used to transform the sport field and player detections in order to obtain two-dimensional coordinates for each player. Team assignment is performed by clustering the player’s bounding boxes.

tions from various data sources. Furthermore, the influence of errors in individual modules and their connections has not been explored. To tackle this demanding real-world task is of interest for the computer vision community as well for sports science, and has direct applications.

**Contributions:** In contrast to commercial systems and related work, we provide the first transparent baseline for *player position estimation* with interchangeable modules, that relies on state-of-the-art techniques and freely available data, while evaluating each module. We demonstrate the generalizability on multiple datasets where the applied models were not originally trained on. The proposed pipeline is also applicable to the so-called “tactic-cam” that is located next to the main camera. It usually covers the entire soccer field (without any cuts) and is consequently of interest for video analysts. To evaluate the global task, estimated positions are compared to ground-truth positional data. This comparison is not trivial due to non-visible players in the video and the influence of errors of individual modules. Therefore, we propose novel evaluation metrics and identify the impact of errors on the final system output.

The remainder of the paper is organized as follows. Section 2 gives a brief overview of related work. The pipeline itself is introduced in Section 3. In Section 4, the different system components and the accuracy of the extracted positional data are evaluated. Finally, Section 5 discusses the results and describes possible areas of future research.

## 2. Related Work

Since the global task of *player position estimation* has not yet been addressed, we briefly review related work for all individual sub-tasks in this section.

Great progress has been made in recent years for **sports field registration** with monocular non-static cameras. Cuevas et al. [7] trained a probabilistic decision tree to classify specific line segments as an intermediate step for homography estimation and integrated a self-verification step to judge whether a predicted homography matrix is cor-

rect. Homayounfar et al. [16] propose a solution that relies on field segmentation and Markov random fields. Sharma et al. [47] and Chen and Little [4] propose the nearest neighbor search from a synthetic dataset of pairs of edge images and camera images for fully-automated registration. Jiang et al. [19] present a two-step deep learning approach that initially estimates a homography and minimizes the error using another deep network instead of the Lucas-Kanade algorithm [1]. Citraro et al. [6] suggest an approach that also takes into account the position of players and is trained on a separate dataset for uncalibrated cameras. Sha et al. [46] propose an end-to-end approach for area-based field segmentation, camera pose estimation, and online homography refinement that allows end-to-end training and efficient inference. Nie et al. [37] tackle the challenge when no prior knowledge about the camera is available and propose a multi-task network to simultaneously detect a grid of key points and dense field features to estimate and refine a homography matrix end-to-end. This approach seems suitable since also temporal consistency is verified for successive frames. However, a very large number of training samples is required to achieve the desired accuracy and generalizability, but training data are not publicly available except for the *WorldCup2014* dataset (*WC14* [16]).

**Shot boundary detection** (e.g., [14, 26, 53, 58]) and **shot type classification** (e.g., [45, 54]) are necessary pre-processing steps for many tasks of video analysis. It enables the distinction between different camera shot types. Related work in the context of soccer distinguishes between three [48], four [35] or five [59] different camera shot types. For the extraction of positional data, the main camera (with the largest distance) offers the most useful information, because it normally covers a larger part of the field depicting several players.

There are several approaches for the **detection of players** in sports analysis [21, 28, 43, 60]. Although general-purpose approaches for object detection [25, 42] are also able to detect persons, sports offer specific challenges. For

example, the players are usually small, they differ in scale due to the distance from the camera, they can occlude one another, and there is blur caused by camera movement. Nevertheless, specialized approaches [21, 60] compare themselves to general-purpose detectors such as the *Single Shot Detector (SSD)* [25] or *Faster R-CNN* [42]. Komorowski et al. [21] have recently introduced a computationally much more efficient method with results similar to a fine-tuned *Faster R-CNN*. In team sports, the jerseys of the teams are designed so that they can be easily recognized by their color. Thus, for **team assignment** of the (detected) players, color information can be used as a discriminant feature. Hand-crafted (color) features ([9, 30, 55]) or features from convolutional neural networks (CNNs) ([18, 22, 29]) are exploited and clustered by these approaches for team assignment. An approach for player detection and team discrimination [18] addresses the problem of occlusions and errors in object detection [32].

### 3. Player Position Estimation in Soccer Videos

A frequent problem in the field of automatic sports analysis is the lack of publicly available datasets. Currently, there is no public dataset that provides positional data for given broadcast soccer videos. Besides, related work solely considered sub-problems of the overall task of *player position estimation*. This section describes a pipeline as well as the choice and modifications of individual components that solve all required sub-tasks for *player position estimation* to predict the two-dimensional player positions on the field given an input (broadcast) video (Figure 1).

After all relevant (main camera) shots are identified (Section 3.1), the step of sports field registration is essential to extract position data (Section 3.2). A homography matrix is determined and used to transform the positions of the players from the image plane into world coordinates (Section 3.3).

#### 3.1. Shot Boundary and Shot Type Detection

We aim at estimating player positions in frames recorded by the main camera since it is most frequently used and shows the area of the game that is relevant for tactical analysis, as shown in Figure 1. We first extract shots from the television (TV) broadcast using the widely applied *TransNet* [26, 50] for shot boundary detection. Since our objective is to gather only valuable positional data, we subsequently apply shot type classification to identify shots captured by the main camera. We exploit the homography matrices estimated by the sports field registration approach presented in Section 3.2. We found that the homography matrices do not change fundamentally in successive frames captured by the main camera. On the other hand, all other cameras that, for example, capture player close-ups or actions depict no or only small fractions of the sports field

causing large errors and consequently inconsistencies in the predicted homography matrices. For this reason, we calculate the average  $\bar{\mathcal{L}}_H$  of the homography changes for each shot. The homography change for two successive frames  $t$  is defined as  $\mathcal{L}_H(H_t, H_{t+1}) = \|H_t - H_{t+1}\|_2$  where each entry in  $H$  is (min-max) normalized for each shot. Finally, we classify each shot as the main camera shot if the condition  $\bar{\mathcal{L}}_H \leq \tau$  is fulfilled.

#### 3.2. Sports Field Registration

The task of sports field registration aims at determining a homography matrix  $H$  for the transformation of an image from the (main) camera into two-dimensional sports field coordinates. Formally, the matrix  $H$  defines a two-dimensional projective transformation and is defined by a  $3 \times 3$  matrix with eight degrees of freedom. We use Chen and Little’s approach [4] as the basis for sports field registration. The camera calibration is defined as the nearest neighbor search in a synthetic dataset of edge map camera pairs. We choose this approach for multiple reasons: (1) It obtains almost state-of-the-art performance on the only test set for soccer [16], (2) it does not rely on manual annotations to obtain training data [5, 19, 37], and (3) is adaptable to other environments (e.g., stadiums and camera parameters) by changing only a few hyper-parameters, as shown in our experiments (Section 4.4).

Chen and Little [4] adopt a *pix2pix* [17] model for field segmentation and the subsequent detection of the field markings. The edge images generated in this way are compared with a dataset of synthetic edge images for which the camera parameters are known ( $x, y, z$  position, focal length, pan, tilt). This comparison is based on a Siamese CNN [15], which takes two edge images as input. Feature vectors are used to construct the reference database. The nearest neighbor search on the feature vectors is then applied by computing the  $L_2$  distance over all pairs. The camera parameters of the nearest neighbor in the synthetic dataset are used to determine an initial homography matrix. This initial estimation is refined using the Lucas-Kanade algorithm [1].

#### 3.3. Player Detection and Position Estimation

Sports analysis offers some specific challenges for the task of object detection and tracking, e.g., the objects (like players) are often small because they are far away from the camera. Camera motion causes blur in the players’ silhouettes. But the movement with unpredictable changes of players’ direction and pace poses problems also for well-tested approaches. Therefore, some approaches address these problems in the architectural design [21, 28]. Zhou et al. [61] solves object detection and tracking based on the object center and should therefore be less susceptible to movements of the players. In Section 4.2, a comparison of three approaches is performed.

To determine the actual position of each player on the field, we can utilize the predicted homography matrix  $H$ , which maps pixel coordinates to sports field coordinates. We define the image position  $\tilde{p} \in \mathbb{R}^2$  of players as the center of the bottom of the detected bounding box, which usually corresponds to the feet of the player. The predicted position  $\hat{p} \in \mathbb{R}^2$  of the player on the field is then calculated with the inverse homography matrix and the detected image positions of the players:  $\hat{p} = H^{-1}\tilde{p}$ .

**Self-Verification (sv):** The predicted positions can be used to verify the homography matrix extracted by the sports field registration. Assuming that most player positions should be assigned to a coordinate within the sports field, the system can automatically discard individual frames where the sports field registration is obviously erroneous. If one of the projected player positions is far outside the dimensions of the field including a tolerance distance  $\rho$  in meter, then normally there is an error in the homography estimation. The smaller the value  $\rho$  is chosen, the more frames are discarded, because only smaller errors in the homography estimation are being tolerated. Intuitively, a tolerance distance between two and five meters seems reasonable which is proven experimentally (Section 4.5).

**Team Assignment:** Assuming that for some sports analytic tasks the position of the goalkeeper is of minor relevance (e.g., formation or movement analysis) and it is extremely rare that both goalkeepers are visible in the video at the same time, they are ignored in the team assignment step. Due to the different jersey type and color it requires context information (i.e. the location) to correctly assign the team.

Another problem is that coaches and attendants also protrude onto the sports field with their bodies due to the perspective of the camera so that the number of visible classes which appear in a frame cannot be predetermined. We present a simple approach that provides a differentiation between only two classes (team A and B) based on the object detection and assumes that the use of an unsupervised clustering method is more appropriate in this domain since it does not rely on any training data and the player detection results are already available with high quality. We apply *DBScan* [11] to determine two dominant clusters representing the field players of both teams. Any unassigned detection, which should include goalkeepers, referees, and other persons, is discarded. The feature vectors are formed based on the player detection results, i.e., the bounding boxes. We use the upper half of a bounding box since it usually covers the torso of a player. Each bounding box is first uniformly scaled to  $20 \times 20$  and then the center of size  $16 \times 16$  is cropped. This should reduce the influence of the surrounding grass in the considered area. Since the jersey colors differ greatly, it is sufficient to use the average over color channels (HSV color space). It can be assumed that

field players are most frequently detected and that this is roughly balanced between both teams. Furthermore, due to the previous segmentation of the playing field, only a few detections are expected which are not field players. *DBScan* requires two parameters:  $\epsilon$ , which is the maximum normalized (color)  $L2$  distance between two detections to be assigned to the same cluster, and  $n_{cls} \in [0, 0.5]$ , which specifies how many of all detections must belong together to form a cluster (maximum of 0.5 due to two main clusters). Since the optimal value for  $\epsilon$  will be different for each match, a grid search for randomly selected frames of each sequence of the match is performed to determine the parameter. In contrast to previous work that generally utilizes color histograms [30, 55] to reduce the input feature space, we apply the average over pixels without any performance decline. The value for  $\epsilon$  is selected, for which the cost function  $c(\epsilon) = |X_{(O)ther}| + ||X_A| - |X_B||$  is minimal and restricted to form exactly two clusters ( $X_A$  and  $X_B$ ). The cost function should ensure that the clusters A and B, which represent the two teams, are about the same size and that there are as few as possible unassigned detections ( $X_{(O)ther}$ ).

## 4. Experimental Results

All individual components are evaluated individually, while the main task of *player position estimation* is evaluated at the end. The main test sets that are used both to evaluate the sports field registration and *player position estimation* are introduced in Section 4.1. As shot boundary and shot type classification (Section 3.1) are common preprocessing steps in video data, we refer to the supplemental material. Section 4.2 and 4.3 focus on the evaluation of player detection and team assignment, while the evaluation of sports field registration is reported in Section 4.4. Finally, the main task is evaluated by comparing the estimated positional data with the ground-truth data (Section 4.5). Further implementation details, supplemental material, and source code are provided at [https://github.com/MM4SPA/vid2pos\\_soccer](https://github.com/MM4SPA/vid2pos_soccer).

### 4.1. Main Datasets

To evaluate the main task of *player position estimation*, synchronized video and positional data are needed. To indicate the generalizability, we use a total of four datasets that are primarily designed only for testing, i.e., no training nor fine-tuning of individual modules is performed on this or closely related data.

Common broadcast videos of different resolutions (SD and HD) and seasons (2012, 2014) are available as well another type of video – the tactic-cam (TC): this camera recording is without any cuts and usually covers a wider range of the pitch. Since the tactic-cam is located next to the main TV camera and usually covers the majority of players, it is usually used for video analysis. In general, each dataset

contains four halves from four matches from the German Bundesliga in 25 Hz temporal resolution with synchronized positional data. Our datasets are referred to as *TV12* (2012, SD resolution), *TV14* (2014, HD), *TC14* (HD), and *TV14-S* that covers the broadcast videos of the same matches as *TC14*. Due to temporal inconsistencies in the raw video of *TV14-S* to the positional data, these videos are synchronized using the visible game clock. The position data are considered as ground truth since they are generated by a calibrated (multi-)camera system that covers the entire field. However, this system can be inaccurate in some cases [39]. An error of one meter is to be assumed in the data provided to us.

The quality of the field registration is essential for the accurate prediction of the player positions, but as there is only one limited dataset for sports field registration in soccer [16], we manually estimate ground-truth homography matrices for a subset of our datasets. In particular, 25 representative and challenging images per match are chosen to cover a wide range of camera settings resulting in 100 annotated images per test set. The remaining modules, i.e., shot boundary and shot type classification, player detection, and team assignment are trained and evaluated on other publicly available datasets and introduced in their respective sections.

## 4.2. Evaluating Player Detection

Player detection and the usage of the homography matrix enable the extraction of two-dimensional coordinates for players. While a general object detector like *Faster R-CNN* [42] localizes the bounding box for each object, this information is not necessarily needed, rather the exact position is of interest. To assess the performance of *CenterTrack* [61] on soccer data, we compare it to another specialized network [21] for this domain and to a general object detection framework that is fine-tuned [42] for the soccer domain. We note that alternative solutions such as [28] exist and a comparison is generally possible. However, it is out of scope of our paper to re-implement and test several variants especially if a satisfactory quality is achieved with the selected solution.

**Datasets & Setup:** Due to the lack of publicly available datasets for training and evaluation, Komorowski et al. [21] train their network on two small-scale datasets [9, 28] where the training and test data is separated by frame-wise shuffling and subsequential random selection (80% training, 20% test). *CenterTrack* can exploit temporal information to track players. However, to the best of our knowledge, there exists only one dataset in the domain of soccer with tracking information (*ISSIA-CNR* [9]), but it contains a very limited number of scene perspectives from multiple static cameras and is thus inappropriate for our system. For a fair comparison with the alternative approach, we follow the train-test split of Komorowski et al. [21] where individual

	ISSIA-CNR [9]	Soccer Player [28]
Faster R-CNN [42]	87.4	<b>92.8</b>
FootAndBall [21]	<b>92.1</b>	88.5
<b>CenterTrack</b> [61]	90.1	90.2

Table 1. Performance evaluation for player detection: The average precision in percent is measured on two subsets from the *ISSIA-CNR* and *Soccer Player* dataset.

frames are used for training. The publicly available *ISSIA-CNR* [9] dataset contains annotated sequences from several matches captured by six static cameras (in 30 Hz and FHD resolution) comprising 3000 frames per camera. *Soccer Player* [28] is a dataset created from two professional matches where each match is recorded by three HD broadcast cameras with 30 Hz and bounding boxes are annotated for approximately 2000 frames. For evaluation, we report the average precision (AP) according to [38]. In the final step of *CenterTrack* bounding boxes are estimated, which makes AP a suitable metric to compare the performance of object detectors, even though the size of the bounding box is not relevant to extract positional data. We refer to the supplemental material for details about the training process.

**Results:** The results on the test set for our fine-tuned *Faster R-CNN* [42], Komorowski et al. [21]’s model and the fine-tuned *CenterTrack* [61] are reported in Table 1. Since *Faster R-CNN* and *FootAndBall* perform well on only one test set and perform significantly worse on the other, this suggests a lack of generalizability whereas *CenterTrack* achieves good results on both data sets. As *CenterTrack* benefits from training with tracking data [61], we are confident that results can further be improved, but choose this model for our pipeline as it already provides good results.

## 4.3. Evaluating Team Assignment

In this experiment, we evaluate the team assignment that relies on detected bounding boxes.

**Dataset & Setup:** In contrast to very small datasets [9, 28], Yu et al. [59]’s dataset provides a good diversity regarding the environmental setting (camera movements, lighting conditions, different matches, jersey colors, etc.). Therefore, a subset from their dataset is manually annotated with respect to team assignment. To bypass errors in the player detection, bounding boxes and player assignment are manually annotated for a set of frames containing multiple shot perspectives and matches. Team assignment is annotated for three categories, *team A*, *team B*, and *other* including referees and goalkeepers (due to its sparsity). We randomly select one frame for a total of ten shots captured by the main camera for each match. We took 20 matches that were already used to evaluate the temporal segmentation resulting in 200 frames for evaluation. As mentioned before, the aim

is to find two main clusters, and we found empirically that  $n_{cls} = 0.2$  provides good results for this task.

**Metrics:** Istasse et al. [18] proposed micro accuracy for this task, but this metric only considers labels from both teams and is insufficient in our case, since it can be misleading when the algorithm assigns uncertain associations to the class *other*. To prevent this, referees and goalkeepers must be excluded from the object detection or an alternative metric needs to be defined. For this reason, we additionally consider the macro accuracy for all three classes.

**Results:** Our simple method performs well, both in terms of macro accuracy (0.91) for the three classes and micro accuracy (0.93) for the two team classes. We found that most errors are players that are assigned to *other* (goalkeepers, referees). This leads to the conclusion that field players are assigned correctly with a high probability in most cases. In comparison to an end-to-end approach for team assignment of Istasse et al. [18], where the overall performance is evaluated on basketball data, a similar micro accuracy (0.91) is reported. However, the domain basketball differs much from soccer making a direct comparison difficult.

#### 4.4. Importance of Sports Field Registration

As already introduced, many approaches rely on manually annotated ground-truth data for training. There exist only one public benchmark dataset (*World-Cup2014 (WC14)* [16]). While the test set follows the same data distribution as the training data, in particular, the camera hyper-parameters (location, focal-length, etc.), generalization capabilities are not investigated by existing solutions [4, 19, 37, 46]. Primarily, to indicate the adaptability of Chen and Little [4]’s approach (Section 3.2) to different environmental settings, we explore several hyper-parameters on our target test sets (see Section 4.1). Additionally we compare them with recent work.

**Metrics:** Since the visible part of the pitch is of interest for application, we report the intersection over union ( $IoU_{part}$ ) score to measure the calibration accuracy. It is computed between the two edge images using the predicted homography and the ground-truth homography on the visible part of the image.

**Camera Hyper-parameters:** In general, we assume that the recommended parameters [4] (derived from *WC14* [16]) for generating synthetic training data fit for many soccer stadiums. However, we also evaluate slight modifications of the base camera parameters which are available in *WC14*: camera location distribution  $\mathcal{N}(\mu = [52, -45, 17]^T, \sigma = [2, 9, 3]^T)$  in meters, i.e., the average location from all stadiums (origin is the lower left corner flag of the pitch); focal length ( $\mathcal{N}(3018, 716)mm$ ) and pan ( $\mathcal{U}(-35^\circ, 35^\circ)$ ), tilt ( $\mathcal{U}(-15^\circ, -5^\circ)$ ) ranges. We extend the pan and tilt range to  $(-40^\circ, 40^\circ)$  and  $(-20^\circ, -5^\circ)$ , respectively, in all mod-

Approach	WC14 [16]		TV12	TV14	TC14	TV-S
	Mean	Med.	Mean (std)	Mean	Mean	Mean
[4] repr. w.o. refinement	88.3	90.2	63.6 (34.8)	80.2	82.7	84.5
[4] reproduced	93.6	96.5	66.5 (37.4)	85.2	88.7	90.2
$\mathcal{U}_{focal\ len.}$	92.2	96.6	64.3 (34.8)	85.5	<b>92.5</b>	89.6
$\mathcal{U}_{focal\ len.} + 2x\ num.\ cam.$	94.6	96.2	59.5 (39.1)	82.4	89.9	<b>91.0</b>
$\mathcal{U}_{focal\ length} + \mathcal{U}_{xyz}$	92.2	95.6	61.1 (38.3)	<b>87.4</b>	87.1	89.8
Jiang et al. [19] (repr.*)	95.1	96.7	72.1	72.5	65.0	76.6
CCBV [46]	94.2	95.4	-	-	-	-
Student CCBV [5]†	88.5	92.3	-	-	-	-
Teacher CCBV [46] [5]	96.6	98.7	-	-	-	-
Nie et al. [37] keypoints	95.8	97.2	-	-	-	-
Nie et al. [37] alignment	95.9	97.1	-	-	-	-

Table 2. Evaluation of multiple candidates for the sports field registration on several test sets using  $IoU_{part}$ . \*official released model; †no fine-tuning on WC14 but learned from private teacher model; Gray colored: private training data.

els. As the tactic-cam obviously covers a wider range (especially focal length as seen in Figure 2 A,D,E), we also test versions, where we uniformly sample from the focal length parameters ( $\mathcal{U}_{focal\ length}(1000, 6000)$ ) and from the locations ( $\mathcal{U}_{xyz}([45, -66, 10]^T, [60, -17, 23]^T)$ ), and double the number of training images to 100 000. Training process for line segmentation and homography estimation remain unchanged and we refer to Chen and Little [4] for implementation details or the publicly available source code.

**Results:** As reported in Table 2 the reproduced results (base parameters) from Chen and Little [4] at *WC14* are of similar quality compared to other methods [19, 37, 46]. We observe a noticeable drop in  $IoU_{part}$  on our test sets where the camera parameters (especially the camera position  $(x, y, z)$ ) are unknown. For the *TV12* test set all configurations fail on challenging images. This further indicates that the original parameters are optimized for the camera dataset distribution in *WC14*. However, on the remaining three test sets, the approach of Chen and Little [4] is able to generalize, whereas an alternative solution [20] fails. Due to the non-availability of (private) training data a comparison with [37, 46] is not fair (colored gray). Yet, these approaches seem to yield comparable results. A (student) CCBV [46] model from [5] is trained on the output of a teacher model. As it was originally trained on a large-scale and private dataset, noticeable lower transfer performance is observed on *WC14*. In summary, with slight changes in the hyper-parameters, the approach from Chen and Little [4] is suitable for the applicability to new data without fine-tuning by human annotations.

#### 4.5. Player Position Estimation

This section investigates the performance for player position estimation. Besides, errors of individual modules, i.e., sports field registration, player detection, and team assignment as well as compounding errors of the system are

discussed. We choose the full datasets as introduced in Section 4.1. Despite the shot boundary and shot type classification provide good results (see Appendix), we eliminate their influence by considering manually annotated shots as the results for position estimation depend on this pre-processing step. False-negative errors lead to a lower number of relevant frames for the system’s output and for evaluation, while false-positive errors (e.g., close-ups) primarily produce erroneous output for homography estimation.

**Metrics:** We measure the distance (in meters) between the estimated positions and the actual positions by taking the mean and median over all individual frames ( $d_{\text{mean}}$ ,  $d_{\text{med}}$ ) and additionally report how many frames have an error of less or equal than  $l \in \{2.0, 3.0\}$  meters ( $acc_l$ ). As previously mentioned, the sensor devices that capture position data (used as ground-truth also in other works [33]) can be slightly inaccurate. A domain expert confirmed, that errors in our system of less or equal than  $l \leq 2m$  can be considered as correct results and that errors of less than  $3m$  can still be meaningful for some sports analysis applications.

*Matching estimated positions to ground-truth:* Most of the time only a subset of players is visible in the broadcast videos and there is no information about which player is visible at a certain frame – making evaluation complex. As there is no direct mapping between predicted and ground-truth positions and the number of detections may vary, the resulting linear sum assignment problem first minimized using the *Hungarian Method* [23]. Its solution provides a set of distances for each field player visible in the frame  $t$ , formally  $\mathbb{D}_t = \{d_1, \dots, d_n\}$  where  $n$  is the number of players and  $d_i = \|\hat{\mathbf{p}} - \mathbf{p}\|_2$  is the distance between the estimated position  $\hat{\mathbf{p}} \in \mathbb{R}^2$  for the  $i$ -th player to its actual (ground-truth) position  $\mathbf{p}$ . To aggregate the player distances of one frame, the use of the average distance as an error metric can be misleading as an outlier, e.g., a false-positive player detection (like a substitute player or goalkeeper) can be matched to a ground truth position with high distance (Fig. 2 D,E,H). These outliers can drastically affect the average distance and lead to wrong impressions. To efficiently reject outliers without using an error threshold as an additional system parameter, we propose to report the average distances of the best 80-percent position estimates. Detailed results for this aggregation are included in the supplemental material.

*Player mismatch (pm) due to homography estimation & player detection errors:* Despite the self-verification (sv) step (Section 3.3) that discards erroneous homography estimations, we cannot directly evaluate whether the remaining homography matrices are correct, since some errors are not considered (e.g., wrong focal length as in Fig. 2 D). To analyze the impact of very inaccurate homography estimations and major errors in player detection, we utilize ground-truth data to isolate these types of failures. We re-

Dataset	sv	pm	Ratio	Team Assignment Constraint							
				no				yes			
				$d_{\text{mean}}$	$d_{\text{med}}$	$acc_2$	$acc_3$	$d_{\text{mean}}$	$d_{\text{med}}$	$acc_2$	$acc_3$
TV12			1.00	4.15	1.84	0.56	0.71	5.33	3.47	0.27	0.43
	✓		0.90	2.52	1.66	0.61	0.77	4.06	3.31	0.30	0.46
	✓	✓	0.79	2.10	1.55	0.64	0.80	3.39	2.99	0.34	0.51
TV14			1.00	3.33	1.78	0.56	0.70	4.95	2.20	0.31	0.46
	✓		0.84	2.82	1.71	0.55	0.71	3.57	1.81	0.34	0.51
	✓	✓	0.72	2.29	1.64	0.60	0.77	3.17	1.71	0.35	0.53
TC14			1.00	3.71	1.20	0.74	0.83	3.32	1.39	0.65	0.78
	✓		0.89	1.81	1.14	0.79	0.88	2.16	1.34	0.68	0.81
	✓	✓	0.78	1.66	1.13	0.79	0.88	1.92	1.29	0.71	0.81
TV14-S			1.00	2.47	1.36	0.69	0.81	3.19	2.44	0.43	0.58
	✓		0.92	1.89	1.29	0.73	0.85	2.89	2.34	0.44	0.59
	✓	✓	0.75	1.73	1.27	0.75	0.87	2.78	2.32	0.45	0.60

Table 3. Results regarding mean ( $d_{\text{mean}}$ ) and median error ( $d_{\text{med}}$ ) in meters and fraction of frames with an error of less or equal than  $l$  meters ( $acc_l$ ) of the total system on several datasets. *Ratio* indicates how many frames are kept for evaluation after applying different criteria (system output: only with sv).

project all ground-truth positions to the image space according to the estimated homography matrix. If the number of detected players differs significantly from the actual players then the homography is probably erroneous (called player mismatch: pm). We also define a tolerance range of 5% of the image borders to include players that are at the boundary to avoid penalizing smallest discrepancies in the estimation of the homography matrix. Finally, we discard all frames for evaluation that do not satisfy the following condition:

$$\alpha_t := 1 - \zeta < \frac{|\mathbb{D}_t^{\text{real}}|}{|\mathbb{D}_t^{\text{gt}}|} < 1 + \zeta \quad (1)$$

$\alpha_t$  is the indicator function whether a frame  $t$  is discarded based on the ratio of detected players  $|\mathbb{D}_t^{\text{real}}|$  and expected players  $|\mathbb{D}_t^{\text{gt}}|$ . For example, assuming that ten players are claimed to be visible, but only six are detected by the system, then we want to discard such discrepancies and set  $\zeta = 0.3$ . Furthermore, we incorporate a constraint to measure the results after team assignment by differentiating between the teams before the linear sum assignment and report the mean distance over both teams per frame.

**Results:** Table 3 shows the results for each dataset while taking the best performing models from Table 2. The results are summarized for all matches by taking the mean of the per match results. The results after self-verification sv are the output of our system and set its tolerance area to  $\rho = 3m$ . The results for other thresholds are reported in the supplemental material. With the pm criteria the impact of erroneous sports field registration is analyzed. As evaluated, the applied model for sports field registration provides good results, however, for a couple of frames the  $IoU_{\text{part}}$  is below 90%. Our sv-process is able to discard some of



Figure 2. Qualitative results of the proposed system for the extraction of positional player data ordered from low (left) to high error (right): The top row presents the output without considering teams. The green triangles correspond to the predicted positions ( $\nabla$ ) of players and the black points to the ground-truth positions ( $\bullet$ ); team assignments are colored red and blue. For the input image the ground-truth positions are re-projected according to the estimated homography matrix; in the sports field some grid points are highlighted.

these frames as the error drops significantly. For the remaining frames, the pipeline provides promising results on all datasets. The  $pm$  criterion demonstrates the high impact of the sports field registration. Even for marginal errors in the homography estimation, i.e., ca. 95% IoU, the absolute error in meter (mean) is about 1 m when back-projecting known keypoints [37]. Hence, the applied reduction of bounding boxes to one point does not substantially affect the error in meters. The qualitative examples in Figure 2 (with applied  $sv$ ) primarily show the output of the pipeline and support the choice of our metrics. In Figure 2 (I, J), the output is obviously erroneous, but not discarded in the  $sv$  process demonstrating the importance of an accurate sports field registration. Furthermore, the influence of false-positive field players (A, D, E), and incorrect team identification is visible (G, H, I). Since the quantitative results are weaker with team assignment, this suggests a lack of generalizability to the test data for the player detection and team assignment module. Indeed, player tracking is not covered which would lead to more stable predictions across multiple frames, and temporal consistency of the sports field registration is not evaluated quantitatively. However, the sports field registration appears to provide stable results even without explicitly treating the temporal component but could be post-processed in an additional step [24, 47].

In summary, we claim that our system outputs promising results in many cases providing a first baseline to conduct

various automatic analyses, for instance, regarding formation detection [2, 36] or space control [12, 41].

## 5. Conclusions & Future Work

In this paper, we have presented a fully-automated system for the extraction of positional data from broadcast soccer videos with interchangeable modules for shot boundary detection, shot type classification, player detection, field registration, and team assignment. All components as well as their impact on the overall performance were evaluated. We investigated which parts of the pipeline influence each other and how they could be improved, e.g., by fine-tuning a specific module with more appropriate data. A relatively small error in meters should allow sports analysts to study team behavior. Indeed, the adaptation to other sports would definitely be interesting. In the future, we also plan to integrate a tracking module. However, additional steps for player re-identification (within and across shots) are necessary to allow player-based analysis across a match.

## Acknowledgement

This project has received funding from the German Federal Ministry of Education and Research (BMBF – Bundesministerium für Bildung und Forschung) under 01IS20021B and 01IS20021A.

## References

- [1] Simon Baker and Iain Matthews. Lucas-kanade 20 years on: A unifying framework. *International journal of computer vision*, 56(3):221–255, 2004.
- [2] Alina Bialkowski, Patrick Lucey, Peter Carr, Yisong Yue, Sridha Sridharan, and Iain Matthews. Large-scale analysis of soccer matches using spatiotemporal tracking data. In *2014 IEEE International Conference on Data Mining*, pages 725–730. IEEE, 2014.
- [3] P. S. Bradley and J. D. Ade. Are current physical match performance metrics in elite soccer fit for purpose or is the adoption of an integrated approach needed? *International Journal of Sports Physiology and Performance*, 13(5):656–664, 2018.
- [4] Jianhui Chen and James J. Little. Sports camera calibration via synthetic data. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 2497–2504. Computer Vision Foundation / IEEE, 2019.
- [5] Anthony Cioppa, Adrien Delière, Floriane Magera, Silvio Giancola, Olivier Barnich, Bernard Ghanem, and Marc Van Droogenbroeck. Camera calibration and player localization in soccernet-v2 and investigation of their representations for action spotting. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2021, virtual, June 19-25, 2021*, pages 4537–4546. Computer Vision Foundation / IEEE, 2021.
- [6] Leonardo Citraro, Pablo Márquez-Neila, Stefano Savare, Vivek Jayaram, Charles Dubout, Félix Renaut, Andres Hsufura, Horesh Ben Shitrit, and Pascal Fua. Real-time camera pose estimation for sports fields. *Mach. Vis. Appl.*, 31(3):16, 2020.
- [7] Carlos Cuevas, Daniel Quilon, and Narciso García. Automatic soccer field of play registration. *Pattern Recognition*, 103:107278, 2020.
- [8] V Di Salvo, Ramon Baron, H Tschan, FJ Calderon Montero, N Bachl, and F Pigozzi. Performance characteristics according to playing position in elite soccer. *International Journal of Sports Medicine*, 28(03):222–227, 2007.
- [9] Tiziana D’Orazio, Marco Leo, Paolo Spagnolo, Pier Luigi Mazzeo, Nicola Mosca, Massimiliano Nitti, and Arcangelo Distante. 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):1804–1818, 2009.
- [10] B. Drust, G. Atkinson, and T. Reilly. Future perspectives in the evaluation of the physiological demands of soccer. *Sports Medicine*, 37(9):783–805, 2007.
- [11] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *kdd*, volume 96, pages 226–231, 1996.
- [12] Javier Fernandez and Luke Bornn. Wide open spaces: A statistical technique for measuring space creation in professional soccer. In *Sloan sports analytics conference*, volume 2018, 2018.
- [13] J. Garganta. Trends of tactical performance analysis in team sports: bridging the gap between research, training and competition. *Portuguese Journal of Sport Sciences*, 9(1):81–89, 2009.
- [14] Michael Gygli. Ridiculously fast shot boundary detection with fully convolutional neural networks. In *2018 International Conference on Content-Based Multimedia Indexing (CBMI)*, pages 1–4. IEEE, 2018.
- [15] Raia Hadsell, Sumit Chopra, and Yann LeCun. Dimensionality reduction by learning an invariant mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2, pages 1735–1742. IEEE, 2006.
- [16] Namdar Homayounfar, Sanja Fidler, and Raquel Urtasun. Sports field localization via deep structured models. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 4012–4020. IEEE Computer Society, 2017.
- [17] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.
- [18] Maxime Istasse, Julien Moreau, and Christophe De Vleeschouwer. Associative embedding for team discrimination. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 2477–2486. Computer Vision Foundation / IEEE, 2019.
- [19] Wei Jiang, Juan Camilo Gamboa Higuera, Baptiste Angles, Weiwei Sun, Mehrsan Javan, and Kwang Moo Yi. Optimizing through learned errors for accurate sports field registration. In *IEEE Winter Conference on Applications of Computer Vision, WACV 2020, Snowmass Village, CO, USA, March 1-5, 2020*, pages 201–210. IEEE, 2020.
- [20] Wei Jiang, Juan Camilo Gamboa Higuera, Baptiste Angles, Weiwei Sun, Mehrsan Javan, and Kwang Moo Yi. Optimizing through learned errors for accurate sports field registration. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 201–210, 2020.
- [21] Jacek Komorowski, Grzegorz Kurzejamski, and Grzegorz Sarwas. Footandball: Integrated player and ball detector. In Giovanni Maria Farinella, Petia Radeva, and José Braz, editors, *Proceedings of the 15th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications, VISIGRAPP 2020, Volume 5: VIS-APP, Valletta, Malta, February 27-29, 2020*, pages 47–56. SCITEPRESS, 2020.
- [22] Maria Koshkina, Hemanth Pidaparthy, and James H Elder. Contrastive learning for sports video: Unsupervised player classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4528–4536, 2021.
- [23] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.
- [24] Antje Linnemann, Sebastian Gerke, Sven Kriener, and Patrick Ndjiki-Nya. Temporally consistent soccer field registration. In *2013 IEEE International Conference on Image Processing*, pages 1316–1320. IEEE, 2013.
- [25] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott E. Reed, Cheng-Yang Fu, and Alexander C.

- Berg. SSD: single shot multibox detector. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part I*, volume 9905 of *Lecture Notes in Computer Science*, pages 21–37. Springer, 2016.
- [26] Jakub Lokoc, Gregor Kovalčík, Tomáš Soucek, Jaroslav Moravec, and Premysl Cech. A framework for effective known-item search in video. In Laurent Amsaleg, Benoit Huet, Martha A. Larson, Guillaume Gravier, Hayley Hung, Chong-Wah Ngo, and Wei Tsang Ooi, editors, *Proceedings of the 27th ACM International Conference on Multimedia, MM 2019, Nice, France, October 21-25, 2019*, pages 1777–1785. ACM, 2019.
- [27] Benedict Low, Diogo Coutinho, Bruno Gonçalves, Robert Rein, Daniel Memmert, and Jaime Sampaio. A systematic review of collective tactical behaviours in football using positional data. *Sports Medicine*, pages 1–43, 2020.
- [28] Keyu Lu, Jianhui Chen, James J. Little, and Hangen He. 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, 2017.
- [29] Keyu Lu, Jianhui Chen, James J Little, and Hangen He. Lightweight convolutional neural networks for player detection and classification. *Computer Vision and Image Understanding*, 172:77–87, 2018.
- [30] Wei-Lwun Lu, Jo-Anne Ting, James J Little, and Kevin P Murphy. Learning to track and identify players from broadcast sports videos. *IEEE transactions on pattern analysis and machine intelligence*, 35(7):1704–1716, 2013.
- [31] R. Mackenzie and C. Cushion. Performance analysis in football: a critical review and implications for future research. *Journal of Sports Sciences*, 31(6):639–76, 2013.
- [32] Mehrtash Manafifard, Hamid Ebadi, and H Abrishami Moghaddam. A survey on player tracking in soccer videos. *Computer Vision and Image Understanding*, 159:19–46, 2017.
- [33] Daniel Memmert, Koen APM Lemmink, and Jaime Sampaio. Current approaches to tactical performance analyses in soccer using position data. *Sports Medicine*, 47(1):1–10, 2017.
- [34] Daniel Memmert and Dominik Raabe. Data analytics in football: Positional data collection, modelling and analysis. *Journal of Sport Management*, 2018.
- [35] Rabia A Minhas, Ali Javed, Aun Irtaza, Muhammad Tariq Mahmood, and Young Bok Joo. Shot classification of field sports videos using alexnet convolutional neural network. *Applied Sciences*, 9(3):483, 2019.
- [36] Eric Müller-Budack, Jonas Theiner, Robert Rein, and Ralph Ewerth. “does 4-4-2 exist?” –: An analytics approach to understand and classify football team formations in single match situations. In *Proceedings of the 2nd International Workshop on Multimedia Content Analysis in Sports, MM-Sports '19*, page 25–33, New York, NY, USA, 2019. Association for Computing Machinery.
- [37] Xiaohan Nie, Shixing Chen, and Raffay Hamid. A robust and efficient framework for sports-field registration. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1936–1944, 2021.
- [38] Rafael Padilla, Sergio L Netto, and Eduardo AB da Silva. A survey on performance metrics for object-detection algorithms. In *2020 International Conference on Systems, Signals and Image Processing (IWSSIP)*, pages 237–242. IEEE, 2020.
- [39] Svein A Pettersen, Håvard D Johansen, Ivan AM Baptista, Pål Halvorsen, and Dag Johansen. Quantified soccer using positional data: A case study. *Frontiers in physiology*, 9:866, 2018.
- [40] Robert Rein and Daniel Memmert. Big data and tactical analysis in elite soccer: future challenges and opportunities for sports science, dec 2016.
- [41] Robert Rein, Dominik Raabe, and Daniel Memmert. “which pass is better?” novel approaches to assess passing effectiveness in elite soccer. *Human movement science*, 55:172–181, 2017.
- [42] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. In Corinna Cortes, Neil D. Lawrence, Daniel D. Lee, Masashi Sugiyama, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 91–99, 2015.
- [43] Melike Şah and Cem Direkoğlu. Evaluation of image representations for player detection in field sports using convolutional neural networks. In *International Conference on Theory and Applications of Fuzzy Systems and Soft Computing*, pages 107–115. Springer, 2018.
- [44] Hugo Sarmento, Filipe Manuel Clemente, Duarte Araújo, Keith Davids, Allistair McRobert, and António Figueiredo. What performance analysts need to know about research trends in association football (2012–2016): A systematic review. *Sports medicine*, 48(4):799–836, 2018.
- [45] Mattia Savardi, Alberto Signoroni, Pierangelo Migliorati, and Sergio Benini. Shot scale analysis in movies by convolutional neural networks. In *2018 25th IEEE International Conference on Image Processing (ICIP)*, pages 2620–2624. IEEE, 2018.
- [46] Long Sha, Jennifer Hobbs, Panna Felsen, Xinyu Wei, Patrick Lucey, and Sujoy Ganguly. End-to-end camera calibration for broadcast videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13627–13636, 2020.
- [47] Rahul Anand Sharma, Bharath Bhat, Vineet Gandhi, and C. V. Jawahar. Automated top view registration of broadcast football videos. In *2018 IEEE Winter Conference on Applications of Computer Vision, WACV 2018, Lake Tahoe, NV, USA, March 12-15, 2018*, pages 305–313. IEEE Computer Society, 2018.
- [48] M. Sigari, H. Soltanian-Zadeh, V. Kiani, and A. Pourreza. Counterattack detection in broadcast soccer videos using camera motion estimation. In *The International Symposium on Artificial Intelligence and Signal Processing (AISP)*, pages 101–106, 2015.
- [49] SkillCorner. <https://www.skillcorner.com/>. Accessed: 01/06/2021.
- [50] Tomáš Soucek and Jakub Lokoc. Transnet V2: an effective

- deep network architecture for fast shot transition detection. *CoRR*, abs/2008.04838, 2020.
- [51] Metrica Sports. <https://metrica-sports.com/>. Accessed: 01/06/2021.
- [52] StatsPerform. <https://www.statsperform.com/team-performance/football-performance/optical-tracking>. Accessed: 01/06/2021.
- [53] Shitao Tang, Litong Feng, Zhanghui Kuang, Yimin Chen, and Wei Zhang. Fast video shot transition localization with deep structured models. In *Asian Conference on Computer Vision*, pages 577–592. Springer, 2018.
- [54] Wenjing Tong, Li Song, Xiaokang Yang, Hui Qu, and Rong Xie. Cnn-based shot boundary detection and video annotation. In *2015 IEEE international symposium on broadband multimedia systems and broadcasting*, pages 1–5. IEEE, 2015.
- [55] Xiaofeng Tong, Jia Liu, Tao Wang, and Yimin Zhang. Automatic player labeling, tracking and field registration and trajectory mapping in broadcast soccer video. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(2):1–32, 2011.
- [56] TRACAB. <https://tracab.com>. Accessed: 01/06/2021.
- [57] track160.com. <https://www.track160.com/>. Accessed: 01/06/2021.
- [58] Lifang Wu, Shuai Zhang, Meng Jian, Zhe Lu, and Dong Wang. Two stage shot boundary detection via feature fusion and spatial-temporal convolutional neural networks. *IEEE Access*, 7:77268–77276, 2019.
- [59] Junqing Yu, Aiping Lei, Zikai Song, Tingting Wang, Hengyou Cai, and Na Feng. Comprehensive dataset of broadcast soccer videos. In *IEEE 1st Conference on Multimedia Information Processing and Retrieval, MIPR 2018, Miami, FL, USA, April 10-12, 2018*, pages 418–423. IEEE, 2018.
- [60] Lijing Zhang, Yao Lu, Ge Song, and Hanfeng Zheng. RC-CNN: reverse connected convolutional neural network for accurate player detection. In Xin Geng and Byeong-Ho Kang, editors, *PRICAI 2018: Trends in Artificial Intelligence - 15th Pacific Rim International Conference on Artificial Intelligence, Nanjing, China, August 28-31, 2018, Proceedings, Part II*, volume 11013 of *Lecture Notes in Computer Science*, pages 438–446. Springer, 2018.
- [61] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Tracking objects as points. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part IV*, volume 12349 of *Lecture Notes in Computer Science*, pages 474–490. Springer, 2020.