

3D Ball Localization From A Single Calibrated Image

Gabriel Van Zandycke^{1,2} and Christophe De Vleeschouwer¹

¹ UCLouvain ICTEAM/ELEN Belgium <firstname>.<lastname>@uclouvain.be

² Sportradar <f>.<lastname>@sportradar.com

Abstract

Ball 3D localization in team sports has various applications including automatic offside detection in soccer, or shot release localization in basketball. Today, this task is either resolved by using expensive multi-views setups, or by restricting the analysis to ballistic trajectories. In this work, we propose to address the task on a single image from a calibrated monocular camera by estimating ball diameter in pixels and use the knowledge of real ball diameter in meters. This approach is suitable for any game situation where the ball is (even partly) visible. To achieve this, we use a small neural network trained on image patches around candidates generated by a conventional ball detector. Besides predicting ball diameter, our network outputs the confidence of having a ball in the image patch. Validations on 3 basketball datasets reveals that our model gives remarkable predictions on ball 3D localization. In addition, through its confidence output, our model improves the detection rate by filtering the candidates produced by the detector. The contributions of this work are (i) the first model to address 3D ball localization on a single image, (ii) an effective method for ball 3D annotation from single calibrated images, (iii) a high quality 3D ball evaluation dataset annotated from a single viewpoint. In addition, the code to reproduce this research will be made freely available at <https://github.com/gabriel-vandycke/deepsport>

1. Introduction

Ball 3D localization has largely been studied [13], leading to industrial products like the automated line calling system in tennis or the goal line technology in soccer [2]. Most existing solutions rely on multiple view points [6, 16, 18, 20–24] to triangulate the 2D ball positions detected in individual frames, thereby offering high reliability in case of occlusions, which are frequent and severe in some team sports like basketball or American football. Obviously, due to the exploitation of multiple synchronized cameras, those systems are expensive and complex to in-

stall, compared to single viewpoint setups. Yet, there are many practical cases where ball 3D localization is valuable even though only a single viewpoint is available. This typically happens when the end-user is interested in the 3D reconstruction and analysis of the play e.g. to collect statistics related to the 3D movement of the ball and players.

Several works have addressed the question of 3D ball localization using single calibrated viewpoints. They generally rely on 3D ballistic trajectory fitting from the 2D detections [4, 5, 14, 17, 22, 25, 26, 28]. This approach, very well described in [28], is however limited to detections on video sequences (multiple consecutive images), when ball is in free fall (ballistic trajectories).

Nonetheless, in many game situations, the ball is only partially visible or not in free fall. Therefore, in our work, we drop assumptions of ballistic trajectory, time consistency, and clean visibility, and we aim at localizing the ball in 3D from a *single* calibrated image. In particular, we validate our method on basketball games where ballistic trajectories – typically long passes and shots – only represent less than 10% of game time¹. In short, our method estimates ball diameter in pixels and uses knowledge of the real ball size and camera calibration to locate the ball in the 3D space. We use a two-stage approach where a small convolutional neural network (CNN) is trained to predict ball diameter from candidates generated by a detector. In addition, the CNN is also trained to confirm/refute the presence of a ball, thereby improving the performance of the global image-wise ball detector. To the best of our knowledge, our approach is the first one to make ball 3D localization possible from a single viewpoint without resorting to temporal analysis and ballistic trajectory assumption.

Since an appropriate dataset is crucial for the training of our model, we also present a method to produce high quality ball 3D annotations in single view setups, and used it to create a 3D localization dataset made freely available here: <https://www.kaggle.com/gabrielvandycke/ballistic-raw-sequences>.

¹Evaluated on our dataset presented in Section 3.1

2. Method

We propose to formulate the ball 3D localization task in a single calibrated image as the combination of two image analysis problems, namely the ball detection and the ball diameter estimation problems. Calibration information and knowledge of the actual ball diameter, in meters, are then used to compute the ball localization in the 3D space from its diameter and position in the image pixels space.

2.1. Ball detection in image space

To detect balls in the image space, we adopt the State-of-the-Art solution *BallSeg* [30] that uses a segmentation approach. Specifically, the model, based on an ICNet architecture [32], is trained to predict a mask of the ball. At test-time, the top- k highest peaks are identified on the output heatmap Ψ to produce k ball candidates.

2.2. Ball diameter estimation in pixels

A baseline using a Hough circle transform. A natural and well-known approach to estimate the size of a circular and highly contrasted object in an image consists in applying the Hough-circle transformation (HCT) [10] to the image gradient magnitude. As a baseline, to estimate the ball diameter in the image space, we propose to use the Hough transformation on Ψ , over a spacial neighborhood of the top ($k=1$) ball candidate.

In practice, several pre-processing steps are applied prior to the Hough transformation. A morphological opening operation with a circular filter of diameter ρ is first applied to Ψ . Edges are then detected using adaptive thresholding with hysteresis [3], with high and low thresholds τ_h and τ_l . Finally, the Hough-circle transformation is computed from candidate diameters bounded in a range defined by the physical limits of the scenes (see Section 4).

Our proposed CNN-based method. In this work, we propose to adopt a Convolutional Neural Network (CNN) to regress the ball diameter, in pixels, from an image patch I centered on the candidates proposed by any conventional ball detector. For a detector that produces k bounding boxes on the initial input image, when $k > 1$, it is likely that most candidates are false positives. Hence, our CNN is trained to predict whether the input patch contains a ball or not, in addition to regressing the ball diameter. As demonstrated in Section 4, this offers the possibility to increase the ball detection rate at constant false positive rate, compared to a detector configured to output the most confident detection.

The CNN model \mathcal{M}_θ adopted in this work follows a VGG [27] architecture that takes image patches I as input, and extracts a feature vector which is connected to three fully connected layers inspired by Le Cun *et al.* in [19].

The last layer has 2 neurons that produce the regression $\hat{d} := \mathcal{M}_\theta^d(I)$ and classification $\hat{c} := \mathcal{M}_\theta^c(I)$ outputs.

We trained \mathcal{M}_θ using batch stochastic gradient descent over θ . The classification output \hat{c} , is supervised with binary cross entropy loss

$$\mathcal{L}^c(c, \hat{c}) = -c \cdot \log(\hat{c}) + (1 - c) \cdot \log(1 - \hat{c}) \quad , \quad (1)$$

with $c \in \{0, 1\}$ being the target class where 1 means I is centered on a ball. The regression output is supervised with Huber Loss [12] defined by

$$\mathcal{L}^d(d, \hat{d}) = \begin{cases} \frac{1}{2}(d - \hat{d})^2 & \text{for } |d - \hat{d}| \leq \delta \\ \delta(|d - \hat{d}| - \frac{1}{2}\delta) & \text{otherwise} \end{cases} \quad , \quad (2)$$

where $d \in \mathbb{R}_0^+$ is the target diameter in pixels, and the meta-parameter δ represents the threshold at which the loss turns from linear to quadratic. The two losses were weighted with a parameter α

$$\mathcal{L} = \alpha \cdot \mathcal{L}^d + (1 - \alpha) \cdot \mathcal{L}^c \quad , \quad (3)$$

but \mathcal{L}^d was ignored for non ball inputs.

At test time, given k candidates predicted by our detector, we extract k crops I_i , with $i \in [1, k]$, from the initial RGB input image. The final ball detection candidate is the one that has maximum c_i .

2.3. Turning ball diameter into 3D localization

For a camera defined by a focal length f , pixel density μ_x and μ_y in x and y direction respectively, a skew coefficient γ , and a principal vector (u_x, u_y) , we have the camera matrix

$$K := \begin{bmatrix} f \cdot \mu_x & \gamma & u_x \\ 0 & f \cdot \mu_y & u_y \\ 0 & 0 & 1 \end{bmatrix} \quad . \quad (4)$$

Given the diameter d and ball position (b_x, b_y) in pixels in the image space, the 3D rays of the ball center \mathbf{b}^c and two diametrically opposed ball edges \mathbf{e}_-^c and \mathbf{e}_+^c , expressed in the camera coordinate system, are

$$\mathbf{b}^c = K^{-1} \cdot \mathcal{R} \left(\begin{bmatrix} b_x \\ b_y \\ 1 \end{bmatrix} \right) \quad \text{and} \quad (5)$$

$$\mathbf{e}_\pm^c = K^{-1} \cdot \mathcal{R} \left(\begin{bmatrix} b_x \\ b_y \pm \frac{d}{2} \\ 1 \end{bmatrix} \right) \quad , \quad (6)$$

where $\mathcal{R} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ is the function that rectifies image coordinates to handle lens distortions.

For a camera placed at $\mathbf{c}^o \in \mathbb{R}^3$ in the world coordinates system, with an orientation defined by the rotation matrix R , and the true ball diameter ϕ , the 3D localization in the ball is then given by \mathbf{b}^o :

$$\mathbf{b}^o = R^T \cdot \frac{\phi \cdot \mathbf{b}^c}{\mathbf{e}_+^c - \mathbf{e}_-^c} + \mathbf{c}^o \quad (7)$$

3. Datasets

Our deep learning method requires a dataset to train and evaluate on. However, not many public datasets with 3D ball annotations exist. Most of the works mentioned in the introduction rely on private datasets, usually captured in one or few different environments [6, 18, 20, 23, 24, 26, 28]. To our knowledge, there are only three public datasets with annotated 3D ball position:

Basket-APIDIS [31] is a 15 minutes women basketball sequence captured by 7 un-synchronized cameras at about 22 frames per second in a single venue equipped with by a rather weak lighting setup, resulting in poorly contrasted images. Ball center was manually labelled by P. Parisot *et al.* on the 7 streams for 3 minutes of the game. A pseudo ground-truth of 3D localization was built by triangulation and time interpolation on a re-sampled stream of 800×600 pixels at 25 frames per second [22]. The difference with the acquisition frame rate implies alignment issues between the images and the corresponding ball 3D annotations due to duplicated frames.

Soccer-ISSIA [9] is a 2 minutes soccer sequence captured from 6 1920×1080 synchronized cameras, at 25 frames per second. Ball position was annotated with a semi automated process on the 6 streams and 3D localization can be retrieved using known calibration. However, as it is designed for player tracking, the ball is often out of the field of view when flying.

DeepSport [1, 29] is a basketball dataset composed of 314 high-resolution panoramic images captured in 15 different basketball arenas during professional matches. Image resolution varies in each arena but generally reaches around 4500×1500 pixels. Ball 3D annotation is given by the annotation of ball center and its vertical projection on the court in the image space (see Figure 1). An extended version of the dataset containing 1200 additional annotated images of similar quality is available under non disclosure agreement (NDA) at [1].

The soccer context does not offer the complexity encountered in basketball and the 2 minutes sequence from ISSIA dataset lacks in diversity. Hence, in this work, we focus on basketball, where balls are frequently occluded, often poorly contrasted, and suffer from noise and motion blur due to indoor lighting conditions. This makes it a challenging sport to address the ball 3D localization task. Despite a scene captured by multiple viewpoints, the APIDIS dataset does not offer accurate ball 3D localization. Indeed, due to the lack of synchronization between the multiple image streams, we observed dozen of pixels differences between the 3D position projected in image space and the visual appearance of the ball. Additionally the dataset does not provide any information about ball visibility, in particular when it is hidden by the players as often encountered

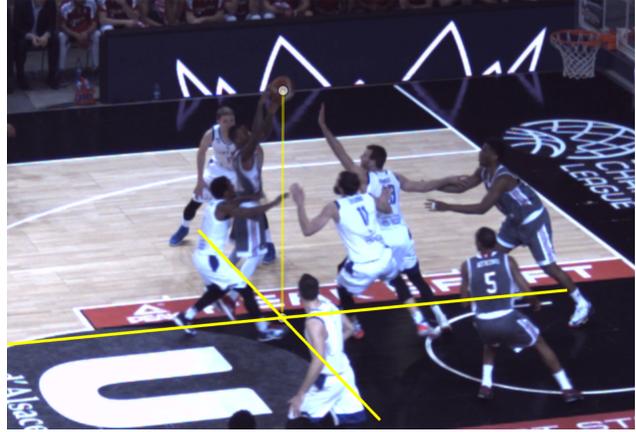


Figure 1. Annotation interface for ball 3D localization in calibrated images. A first point is clicked on the ball center, the second point is clicked at its vertical projection on the ground. The calibration is used to constrain the projection in a vertical line, and a cross aligned with the field geometry is drawn on the ground to help the annotator localizing the ball in the scene.

in basketball scenes. Therefore, we decided to build on the DeepSport dataset for training, as it offers the largest environment diversity with different angles of view, ball color, court color, and lighting conditions.

3.1. Annotation quality assessment

The lack of multiple viewpoints in the DeepSport dataset makes the 3D annotation quality uncertain. Yet, as our experiments presented in section 4.2 reveal, a few pixels diameter error can translate in meters in the 3D space. Hence, we conducted a quality assessment of the annotations resulting from the approach adopted in [1], and compared it against a more natural approach that consists in annotating the ball diameter and using the relation defined by Eq. 7 to retrieve the 3D position.

To build a ground-truth of 3D ball localization using a single calibrated camera, we propose to annotate multiple consecutive images when the ball follows a ballistic trajectory, and fit a motion model on the annotations to regularize and thereby denoise the initial 3D annotations.

We captured 2 sequences during professional basketball games with a static calibrated monocular camera in 2 different arenas, and identified 35 ballistic trajectories composed of a total of 233 images. After an initial annotation using the center+projection approach, trajectories were individually inspected on a visualization made from the first and last image of the trajectory, in which each annotation was displayed (see Figure 2). Annotations were then manually adjusted such that the trajectory derived from the motion model fitting would match the scene.

In practice, we used the second degree model \mathcal{F} , ignor-

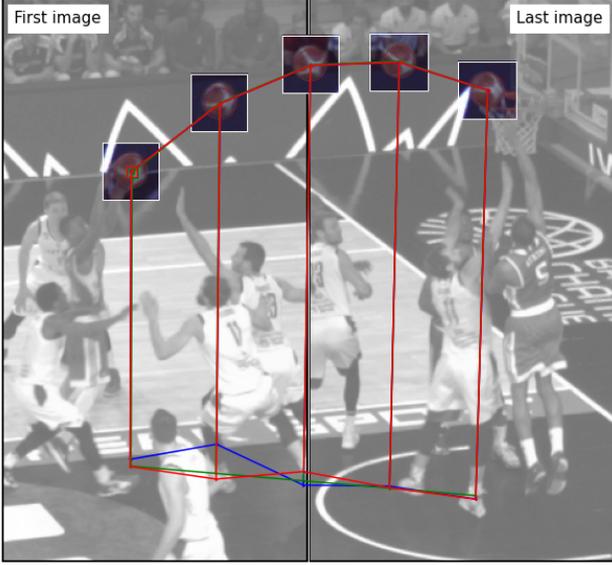


Figure 2. Process for achieving high quality ground-truth for ball 3D localization on calibrated images. (i) Ballistic trajectories are identified. (ii) Images are annotated individually using the center+projection approach (blue annotations). (iii) Trajectories are visualized and inaccurate annotations are manually annotated again (red annotations). (iv) A motion model is fitted to those annotations to obtain near perfection annotations (green annotations).

ing friction, ball spin and other effects:

$$\mathcal{F}_{\mathbf{p}_0, \dot{\mathbf{p}}_0}(t) = \mathbf{p}_0 + \dot{\mathbf{p}}_0 t - \begin{bmatrix} 0 \\ 0 \\ g \end{bmatrix} \frac{t^2}{2}, \quad (8)$$

where \mathbf{p}_0 and $\dot{\mathbf{p}}_0$ are the initial position and initial velocity vectors, and $g = 9.81m/s^2$ is the gravity constant.

Given a trajectory composed of N annotated positions $\hat{\mathbf{p}}_i$ at times t_i , with $i \in [1, N]$, the motion model parameters are obtained by the following relationship

$$\mathbf{p}_0, \dot{\mathbf{p}}_0 = \operatorname{argmin}_{\mathbf{p}_0, \dot{\mathbf{p}}_0} \sum_{i=1}^N \left(\mathcal{F}_{\mathbf{p}_0, \dot{\mathbf{p}}_0}(t_i) - \hat{\mathbf{p}}_i \right)^2. \quad (9)$$

Finally, the high quality positions \mathbf{p}_i are eventually computed by evaluating $\mathcal{F}_{\mathbf{p}_0, \dot{\mathbf{p}}_0}$ at $t = t_i$.

These high quality annotations can be considered as ground-truth to compare the distribution of errors resulting from the two annotation approaches. With a reasoning similar to Equation 7, we can compute the diameter in pixels d that corresponds to the position \mathbf{p} , using the camera calibration and knowledge of the true diameter ϕ .

As shown in Figure 3, the ball center+projection approach is significantly better for ball 3D annotation in single calibrated viewpoints.

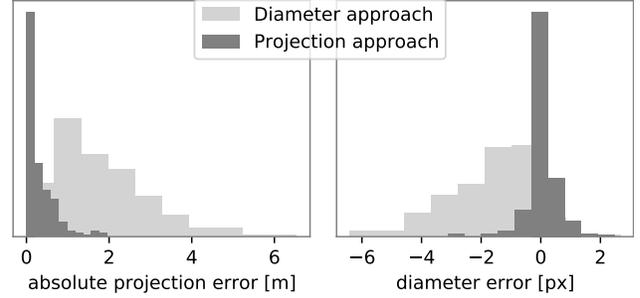


Figure 3. Comparison of two approaches to annotate ball 3D localization on single calibrated views. The diameter approach consists in annotating ball center and diameter, and the projection approach consists in annotating ball center and vertical projection on the ground. The diameter approach shows higher error and more specifically a tendency to under estimate the true diameter. The projection approach is more accurate and translates to a diameter error of less than 1 pixel.

3.2. Used datasets

As mentioned earlier, we used the DeepSport dataset for training, as it offers the largest environment diversity. In contrast to the DeepSport dataset where many balls are partly occluded in cluttered scenes as often encountered in basketball games, our high quality evaluation set presented in Section 3.1 is limited to samples where the ball moves freely, on ballistic trajectories and does not include occluded balls or balls held in player hands. For this reason, in the rest of the paper, the evaluation on the DeepSport dataset was generally preferred. Nevertheless, the APIDIS dataset and our accurate evaluation set have been used to illustrate the generalization capabilities of our model and assess the impact of annotation errors. A summary of the different datasets used in this research is provided in Table 1, and samples of each dataset are given in Figure 9.

	Images	Scenes	ball size range [px]	Resolution
DeepSport [29]	314	15	14 – 37	2 – 5 Mpx
DeepSport ext. [1]	1514	49	14 – 45	2 – 5 Mpx
Our evaluation set	233	2	19 – 40	2336 × 1752
APIDIS [22]	4019	1	5 – 27	800 × 600

Table 1. Summary of the different datasets used in this research. The DeepSport dataset was used both for training and for evaluation, as it offers the largest environment diversity with high quality annotations. The other datasets were used for evaluation only. The extended version of DeepSport was used to assess the impact of the training set size on the performances. The numbers reported for the APIDIS dataset correspond to the images in which ball could be projected. Our evaluation set, presented in Section 3.1, is clean from any occlusions.

4. Experiments

Implementation details and metrics used to conduct this research are presented hereafter.

Dataset As mentioned in Section 3, we use the DeepSport dataset for training, and have adopted the split defined in [29], where each fold contains images from basketball arenas exclusive to that fold. The fold A remained unseen and was used for testing. In addition, we evaluated the models on the APIDIS dataset and our high quality evaluation set presented in Section 3.

Ball detector *BallSeg* was trained as recommended in [30], with a random scaling strategy where ball size was kept between 14 and 37 pixels, corresponding to the range of ball sizes in the DeepSport dataset (see Table 1). To analyze the impact of working on smaller balls, we also considered halving the size range in specific trainings (see Section 4.5). In our context, we are only interested in the ball being played. Therefore, only the top ($k = 1$) candidate was kept, unless stated otherwise.

Baseline method A 64×64 pixels neighbourhood of Ψ around the detection candidate was used to apply the different image processing steps; with $\rho = 37$, $\tau_l = 10$ and $\tau_h = 20$, selected by maximizing the performances on a validation set. The range of candidate diameters for the Hough transformation was constrained, with a 2 meters margin, by the physical boundary of the basketball court.

Proposed method In practice, although our CNN method can be combined with any ball detector we used the detections produced by *BallSeg* as input to our model, allowing a fair comparison with the baseline. As described in Section 2.2, \mathcal{M}_θ is composed of a VGG feature extractor and 3 additional fully connected layers. We took VGG16 pre-trained on ImageNet [7] from TensorFlow library [8], and used Glorot initializer [11] to initialize the last layers. Our CNN was trained using stochastic gradient descent over θ , on batches of 4 different images, with $k = 4$ candidates output by *BallSeg*, making an effective batch size of 16 image patches I . Training took 100 epochs using Adam optimizer [15], with a learning rate schedule starting at 10^{-4} and reducing by a factor 2 for 2 epochs every 10 epochs starting at epoch 50. The loss parameters $\delta = 1.0$ and $\alpha = 0.5$ were selected empirically using the validation set. Image patches of 64×64 pixels were used (see Section 4.6).

Metrics To assess the diameter estimation accuracy, we used the mean absolute error ($\text{MAE}_{[\text{px}]}$) in pixels between the labels d and the predictions \hat{d} . The corresponding

projection error ($\text{MAE}_{[\text{m}]}$) is the distance in meters between the real position and the estimated position, computed with Equation 7, projected on an horizontal plane. It depends on the camera pixel density, focal length and distance to the scene. We also computed the relative distance error ($\text{MAE}_{[\%]}$) that represent the ratio between the projection error and the distance to the cameras. To assess the ball detection performance, we used the Area under the Curve (AuC) associated to the Receiver Operating Characteristic (ROC) curve, plotting the True Positive (TP) rate against the False Positive (FP) rate for a varying detection threshold, as defined in [30]. Here, the TP rate measures the fraction of images in which the ball is correctly detected while the FP rate measures the mean number of false candidates that are detected per image.

Experiments were repeated with 8 different initializations of the networks (for both *BallSeg* and the randomly initialized layers of our CNN). Unless stated otherwise, the mean metric over the 8 repetitions was reported, along with the associated standard deviation.

4.1. Ball diameter estimation

The comparison between the baseline using a Hough-circle transformation and our CNN approach is presented in Figure 4. In this setup, we evaluated on the true positive candidates produced by the detector *BallSeg*. We observe that our CNN-based approach reduces the mean absolute diameter estimation error $\text{MAE}_{[\text{px}]}$ by 3 pixels on each dataset, reaching less than 2 pixels on the DeepSport testset and our evaluation set.

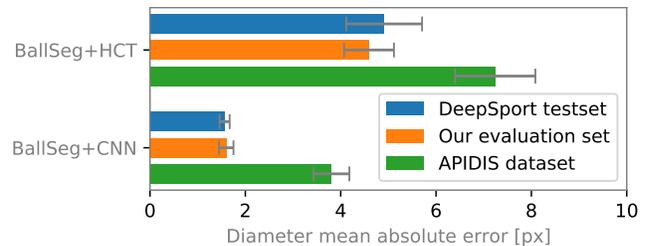


Figure 4. Our CNN method is compared against the Hough-circle transformation baseline (HCT) for diameter estimation on different Basketball evaluation sets. Our method reduces the mean absolute error by a factor two or more.

The fact that the prediction accuracy does not improve between the DeepSport testset and our accurate evaluation set suggests that errors are not caused by annotation errors. Moreover, when referring to Figure 3, we observe that a 2 pixels error is close to the estimated annotation error, meaning that our model achieves an accuracy comparable to human annotator.

	TP _[%]	MAE _[px]	MAE _[m]	MAE _[%]
<i>BallSeg</i> + HCT	47 ± 7 83 ± 2	4.9 ± .8 4.6 ± .5	6.3 ± 1.0 5.1 ± .5	28 ± 5 24 ± 4
<i>BallSeg</i> + CNN	47 ± 7 83 ± 2	1.6 ± .1 1.6 ± .2	2.3 ± .2 1.8 ± .2	10 ± .9 10 ± .7
Oracle + CNN	100 ± 0 100 ± 0	1.9 ± .1 1.5 ± .1	2.8 ± .2 1.7 ± .1	12 ± .6 10 ± .5

Table 2. Error analysis for three different approaches on two different evaluations sets: *DeepSport testset* first, and *Our Evaluation set* second. *BallSeg* + HCT is the baseline that consists in applying Hough-circle transformation on the heatmap Ψ output by *BallSeg*. *BallSeg* + CNN is our proposed method that estimate diameter or image patches I around detection candidates generated by *BallSeg*. Oracle + CNN is our proposed method applied on oracle detections instead of using a detector. The proportion of TP from the detector (TP_[%]) is given for each dataset. The MAE_[px] on oracle detections reaches 1.5 pixels on our evaluation set.

4.2. Ball 3D localization error

The estimated diameter in pixels, along with camera calibration and knowledge of the ball real diameter, allow computing the ball 3D position using Equation 7. The error analysis given in Table 2 shows that a diameter error (MAE_[px]) of 1.6 pixels translates to a projection error (MAE_[m]) of less than two meters on our high-precision evaluation set. However this relationship highly depends on the camera pixel density, focal length and distance to the scene. Relative to the distance to the camera (MAE_[%]), it corresponds to a 10% error. Recall that our image-based approach provides raw estimations, before any possible time filtering.

4.3. Impact study of the detector

As described in Section 2, we adopt a two-stage approach where the balls diameter estimated at the second stage depends on the detections proposed by the first stage. To assess the performance of our CNN independently of the detection stage, the quality of diameter estimation was also evaluated on oracle detections by using input image patches I around the true ball positions. The error analysis given in Table 2 shows that the performances using *BallSeg* detections (1.6 pixels on both datasets) are close to the ones obtained using oracle detections (1.9 and 1.5 pixels).

This study reveals that our CNN gets the ability to correctly estimate the ball size on images where the ball is challenging to detect. Moreover, the similarity between the *DeepSport* testset (subject to occlusions, see samples in [30]) and on our evaluation set (free from occlusions), indirectly reveals that our method is robust to occlusions.

4.4. Ball detection quality

With its classification output \mathcal{M}_θ^c , our CNN model can be used to revise the detection confidence produced by the detector it is combined with. Furthermore, by considering k ball candidates from the detector and keeping the candidate from image patch I_j such that

$$j = \operatorname{argmax}_{i \in [1, k]} \mathcal{M}_\theta^c(I_i) \quad , \quad (10)$$

our CNN becomes able to recover some of the detections missed when the detector adopts a top-1 detection strategy. We conducted experiments with *BallSeg* for different values of k and, as revealed by Figure 5, our model improves the detection quality by a large margin compared to *BallSeg*-top1 alone. More specifically, the ROC curves in Figure 5 reveal two positive effects. First, our CNN is more effective in discriminating between true positives and false positives, as attested by the larger TP-rate at small FP-rate compared to *BallSeg* for $k = 1$. Second, the number of true positives is largely increased with $k > 1$.

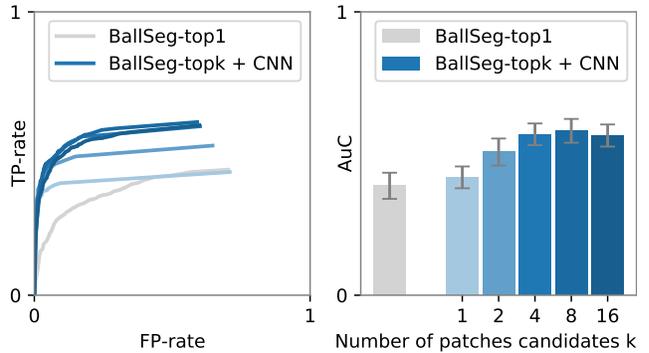


Figure 5. Impact of k on the detection quality of *BallSeg*-top k combined with our CNN approach. (Left) ROC curves obtained by varying the detection threshold, on the accumulated detection candidates over the 8 repetitions (see Section 4). (Right) Area under the ROC curves for the different configurations. Our CNN is able to recover balls missed by *BallSeg*-top1 but who ranked in the top k with $k \geq 2$.

4.5. Impact of ball resolution.

In the *DeepSport* training dataset, ball sizes range from 14 to 37 pixels. However, they are many other applications for 3D ball localization where ball appears smaller, due to lower resolution images like in the APIDIS dataset, or simply because the physical size of the ball is smaller, as encountered for other sports like volleyball or tennis. To evaluate the error distribution one could expect from smaller balls, we trained our CNN on half resolution images and evaluated it on sub-sampled images from the *DeepSport* testset where images were randomly scaled such that balls

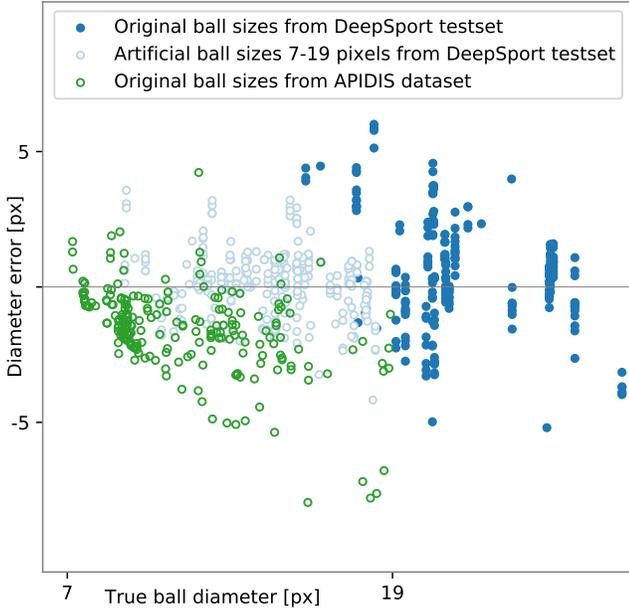


Figure 6. Diameter error distribution with respect to ball true diameter. Only 4 out of 8 repetitions (see Section 4) are displayed to declutter the figure. Our CNN was trained on half resolution images (empty circles) and evaluated both on the DeepSport testset where images were randomly scaled such that balls appear smaller, and the APIDIS dataset where balls larger than 19 pixels or smaller than 7 pixels were ignored. This experiments shows that our CNN approach performs very well at small diameters, validating the approach for many more applications. The distribution shift noticeable on the APIDIS dataset may be caused by the limited quality of that dataset (see Section 3).

appear between 7 and 19 pixels. Results presented in Figure 6 reveal that the error distribution is centered around 0 and shows a standard deviation similar compared to the model working at normal resolution. The model was also evaluated on the APIDIS dataset where balls are smaller (see Table 1). Referring to Figure 6, it appears that the model underestimates the diameter from the balls of the APIDIS dataset.

Results presented in Figure 6 reveal that our model handles smaller resolution balls from the DeepSport dataset very well, validating the approach for many more applications. We conjecture that the distribution shift observed on the APIDIS dataset is caused by the poorly contrasted images of that dataset (see Figure 9).

4.6. Impact of image patch size

Multiple patch sizes were considered. As the Figure 7 reveals, 64×64 image patches resulted in the most accurate estimation of the ball diameter. This study was conducted on oracle detections from our high-accuracy evaluation set, clean from any occlusions.

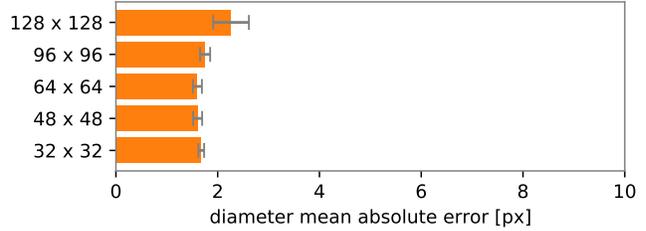


Figure 7. Diameter estimation evaluated on different input image patch I sizes. This study was conducted on oracle detections from our high-accuracy evaluation set, clean from any occlusions.

4.7. Impact of training set size

We also evaluated our method by training both the detector *BallSeg* and our CNN on the extended DeepSport dataset shared under NDA by [1] (see Table 1). Figure 8 reveals that our CNN benefits from being trained on a larger dataset, but also that the baseline experiences a significant gain with a drop of more than 50% of the mean absolute diameter error on all datasets. Visual inspection of the output heatmap Ψ shows that ball is much better segmented, resulting in a better diameter estimation using the Hough circle transform approach. Evaluation on the APIDIS dataset consistently yields weaker performances, corroborating the observation that this dataset suffers from low contrasted images and poor annotations.

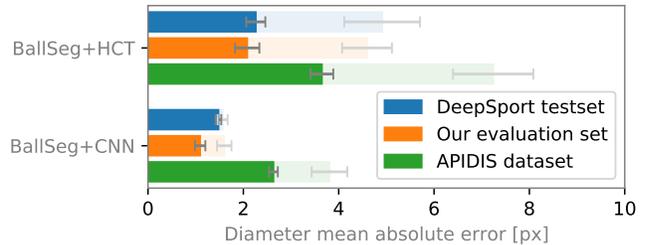


Figure 8. Evaluation of diameter estimation when training the models (both *BallSeg* and our CNN) on the extended DeepSport dataset (see Table 1). Results from Figure 4 are shown dimmed for an easier comparison. *BallSeg+HCT* sees a huge performance gain from being trained on a larger dataset, but our CNN model performs better with a pixel error down to one pixel on our evaluation set.

5. Discussion

The method presented in this work achieves a mean absolute diameter error of 1.6 pixels in the best conditions (ball free from occlusion on high resolution images), which translates to a projection error of 1.8 meters. This might appear as being a lousy approximation. However, in many applications (like 3D analysis of the play to collect statistics), having an automated estimation of ball 3D localization

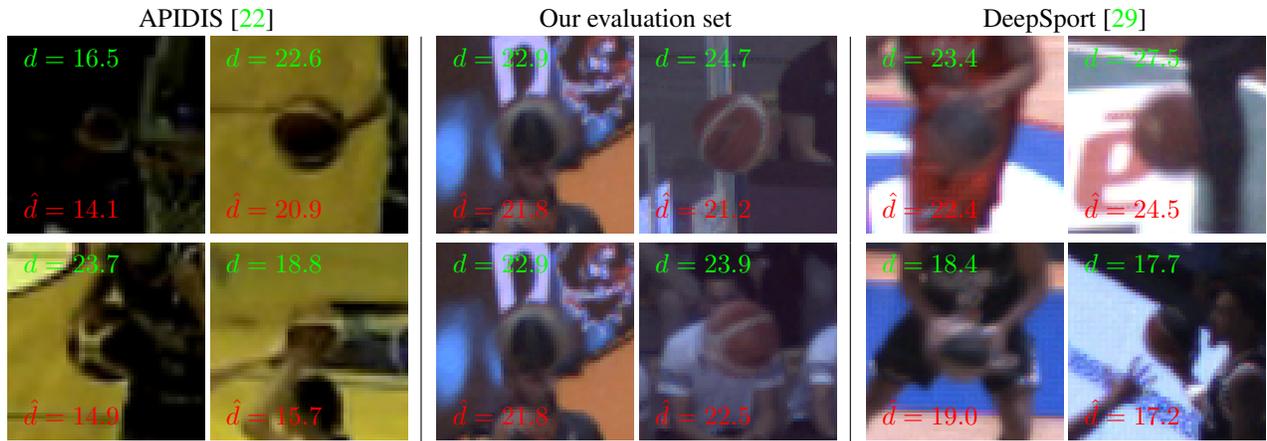


Figure 9. Image samples extracted from the three datasets. d is the true ball diameter in pixels and \hat{d} is the prediction from one of our models. Samples from the DeepSport dataset were extracted from the fold A (see Section 4). Samples from APIDIS dataset were scaled to match ball size range set during training.

from a single camera can be extremely valuable, even with an imprecision of two meters. Besides, our image-based approach provides raw estimations that could be improved based on temporal filtering when needed. However, investigating such temporal regularization goes beyond the scope of our paper. Additionally, for industrial applications, collecting a dedicated dataset is generally required. And, as our experiments reveal, there is margin for improvements by training on a larger dataset.

6. Conclusion

This work has presented the first method to localize a ball in the 3D space from a single monocular image. Our approach is not restricted to balls following a ballistic trajectory, and works on any image where the ball is (event partly) visible, provided that we have camera calibration (which is mandatory for 3D localization) and know the real ball diameter (which is generally the case for many applications). Our method combines a small CNN with an arbitrary ball detector, and compares favorably against a baseline using classical image processing techniques. Additionally, our CNN improves the detection capabilities of the detector it is combined with.

To conduct our study, we have also compared two methods to annotate the ball 3D position on a single image. The results of this comparison demonstrated the benefit of annotating the ball center and its projection on the ground compared to a natural approach of annotating ball diameter.

Acknowledgement

Part of this work has been funded by the Walloon Region project DeepSport and by Synergy Sports Technology, a division of Sportradar. C. De Vleeschouwer is a Research

Director of the Fonds de la Recherche Scientifique - FNRS.

References

- [1] Image and Signal Processing Group (UCL) — Softwares: <https://sites.uclouvain.be/ispgroup/Softwares/DeepSport>. 3, 4, 7
- [2] Baljinder Bal and Gaurav Dureja. Hawk Eye: A Logical Innovative Technology Use in Sports for Effective Decision Making. *Sport Science Review*, 21(1-2):107–119, may 2012. 1
- [3] John Canny. A Computational Approach to Edge Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-8(6):679–698, 1986. 2
- [4] Hua Tsung Chen, Ming Chun Tien, Yi Wen Chen, Wen Jiin Tsai, and Suh Yin Lee. Physics-based ball tracking and 3D trajectory reconstruction with applications to shooting location estimation in basketball video. *Journal of Visual Communication and Image Representation*, 20(3):204–216, 2009. 1
- [5] Hua Tsung Chen, Wen Jiin Tsai, Suh Yin Lee, and Jen Yu Yu. Ball tracking and 3D trajectory approximation with applications to tactics analysis from single-camera volleyball sequences. *Multimedia Tools and Applications*, 60(3):641–667, 2012. 1
- [6] Xina Cheng, Norikazu Ikoma, Masaaki Honda, and Takeshi Ikenaga. Simultaneous physical and conceptual ball state estimation in volleyball game analysis. *2017 IEEE Visual Communications and Image Processing, VCIP 2017*, 2018-Janua:1–4, 2018. 1, 3
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. pages 248–255. Institute of Electrical and Electronics Engineers (IEEE), mar 2010. 5
- [8] TensorFlow Developers. TensorFlow. feb 2022. 5
- [9] T. D’Orazio, M. Leo, N. Mosca, P. Spagnolo, and P. L. Mazzeo. A semi-automatic system for ground truth gener-

- ation of soccer video sequences. In *6th IEEE International Conference on Advanced Video and Signal Based Surveillance, AVSS 2009*, pages 559–564, 2009. 3
- [10] Richard O. Duda and Peter E. Hart. Use of the Hough Transformation to Detect Lines and Curves in Pictures. *Communications of the ACM*, 15(1):11–15, jan 1972. 2
- [11] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Journal of Machine Learning Research*, volume 9, pages 249–256, 2010. 5
- [12] Peter J. Huber. Robust Estimation of a Location Parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, mar 1964. 2
- [13] Paresh R. Kamble, Avinash G. Keskar, and Kishor M. Bhurchandi. Ball tracking in sports: a survey. *Artificial Intelligence Review*, 52(3):1655–1705, 2019. 1
- [14] Taeone Kim, Yongduek Seo, and Ki Sang Hong. Physics-based 3D position analysis of a soccer ball from monocular image sequences. *Proceedings of the IEEE International Conference on Computer Vision*, pages 721–726, 1998. 1
- [15] Diederik P. Kingma and Jimmy Lei Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*. International Conference on Learning Representations, ICLR, dec 2015. 5
- [16] K. C.Amit Kumar, Pascaline Parisot, and Christophe De Vleeschouwer. Demo: Spatio-temporal template matching for ball detection. In *2011 5th ACM/IEEE International Conference on Distributed Smart Cameras, ICDS 2011*, 2011. 1
- [17] Mikel Labayen, Igor G. Olaizola, Naiara Aginako, and Julian Florez. Accurate ball trajectory tracking and 3D visualization for computer-assisted sports broadcast. *Multimedia Tools and Applications*, 73(3):1819–1842, dec 2014. 1
- [18] Christoph H. Lampert and Jan Peters. Real-time detection of colored objects in multiple camera streams with off-the-shelf hardware components. *Journal of Real-Time Image Processing*, 7(1):31–41, 2012. 1, 3
- [19] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2323, 1998. 2
- [20] Andrii Maksai, Xinchao Wang, and Pascal Fua. What players do with the ball: A physically constrained interaction modeling. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2016-Decem, pages 972–981, 2016. 1, 3
- [21] Pascaline Parisot and Christophe De Vleeschouwer. Graph-based filtering of ballistic trajectory. In *Proceedings - IEEE International Conference on Multimedia and Expo*, 2011. 1
- [22] Pascaline Parisot and Christophe De Vleeschouwer. Consensus-based trajectory estimation for ball detection in calibrated cameras systems. *Journal of Real-Time Image Processing*, 16(5):1335–1350, 2019. 1, 3, 4, 8
- [23] Gopal Pingali, Agata Opalach, and Yves Jean. Ball tracking and virtual replays for innovative tennis broadcasts. *Proceedings-International Conference on Pattern Recognition*, 15(4):152–156, 2000. 1, 3
- [24] Jinchang Ren, James Orwell, Graeme A. Jones, and Ming Xu. Real-time modeling of 3-D soccer ball trajectories from multiple fixed cameras. *IEEE Transactions on Circuits and Systems for Video Technology*, 18(3):350–362, 2008. 1, 3
- [25] Hubert Shum and Taku Komura. A spatiotemporal approach to extract the 3D trajectory of the baseball from a single view video sequence. Technical report, 2004. 1
- [26] Hugo Silva, André Dias, José Almeida, Alfredo Martins, and Eduardo Silva. Real-Time 3D Ball Trajectory Estimation for RoboCup Middle Size League Using a Single Camera. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 7416 LNCS:586–597, 2011. 1, 3
- [27] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*. International Conference on Learning Representations, ICLR, sep 2015. 2
- [28] Jonas Sköld. Estimating 3D-trajectories from Monocular Video Sequences Estimating 3D-trajectories from Monocular Video Sequences. 2015. 1, 3
- [29] Gabriel Van Zandycke. DeepSport dataset: <https://www.kaggle.com/gabrielvanzandycke/deepsport-dataset>, 2021. 3, 4, 5, 8
- [30] Gabriel Van Zandycke and Christophe De Vleeschouwer. Real-time CNN-based segmentation architecture for ball detection in a single view setup. *MMSports 2019 - Proceedings of the 2nd International Workshop on Multimedia Content Analysis in Sports, co-located with MM 2019*, pages 51–58, 2019. 2, 5, 6
- [31] C De Vleeschouwer, Fan Chen, and Damien Delannay. Distributed video acquisition and annotation for sport-event summarization. *NEM summit*, 2008. 3
- [32] Hengshuang Zhao, Xiaojuan Qi, Xiaoyong Shen, Jianping Shi, and Jiaya Jia. ICNet for Real-Time Semantic Segmentation on High-Resolution Images. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 11207 LNCS, pages 418–434. Springer Verlag, 2018. 2