

## **Uplifting table tennis: a robust, real-world application for 3D trajectory and spin estimation**

**Daniel Kienzle, Katja Ludwig, Julian Lorenz, Shin'ichi Satoh, Rainer Lienhart**

### **Angaben zur Veröffentlichung / Publication details:**

Kienzle, Daniel, Katja Ludwig, Julian Lorenz, Shin'ichi Satoh, and Rainer Lienhart. 2026. "Uplifting table tennis: a robust, real-world application for 3D trajectory and spin estimation." In IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) 2026, 6-10 March 2026, Tucson, AZ, USA. Piscataway, NJ: IEEE. <https://doi.org/10.48550/arXiv.2511.20250>.

### **Nutzungsbedingungen / Terms of use:**

**licgercopyright**

Dieses Dokument wird unter folgenden Bedingungen zur Verfügung gestellt: / This document is made available under these conditions:

**Deutsches Urheberrecht**

Weitere Informationen finden Sie unter: / For more information see:

<https://www.uni-augsburg.de/de/organisation/bibliothek/publizieren-zitieren-archivieren/publiz/>



# Uplifting Table Tennis: A Robust, Real-World Application for 3D Trajectory and Spin Estimation

Daniel Kienzle  
University of Augsburg  
daniel.kienzle@uni-a.de

Katja Ludwig  
University of Augsburg  
katja.ludwig@uni-a.de

Julian Lorenz  
University of Augsburg  
julian.lorenz@uni-a.de

Shin'Ichi Satoh  
National Institute of Informatics & University of Tokyo  
satoh@nii.ac.jp

Rainer Lienhart  
University of Augsburg  
rainer.lienhart@uni-a.de

## Abstract

Obtaining the precise 3D motion of a table tennis ball from standard monocular videos is a challenging problem, as existing methods trained on synthetic data struggle to generalize to the noisy, imperfect ball and table detections of the real world. This is primarily due to the inherent lack of 3D ground truth trajectories and spin annotations for real-world video. To overcome this, we propose a novel two-stage pipeline that divides the problem into a front-end perception task and a back-end 2D-to-3D uplifting task. This separation allows us to train the front-end components with abundant 2D supervision from our newly created TTHQ dataset, while the back-end uplifting network is trained exclusively on physically-correct synthetic data. We specifically re-engineer the uplifting model to be robust to common real-world artifacts, such as missing detections and varying frame rates. By integrating a ball detector and a table keypoint detector, our approach transforms a proof-of-concept uplifting method into a practical, robust, and high-performing end-to-end application for 3D table tennis trajectory and spin analysis.

## 1. Introduction

Table tennis is a dynamic sport demanding exceptional precision, speed, and strategic thinking. For athletes, coaches, and sports scientists, understanding the intricate 3D trajectory and spin of the ball is paramount for in-depth performance analysis and technique refinement. Such detailed insights, however, are notoriously difficult to obtain from conventional broadcast or monocular video footage, which only provides 2D observations. The rapid motion of the ball, coupled with occlusions, varying lighting conditions, and diverse camera angles, poses significant challenges for

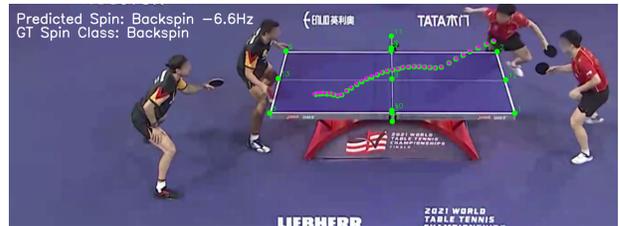


Figure 1. Qualitative example prediction of the full pipeline for a serve trajectory. The green dots represent the front-end detections for 2D ball positions and table keypoints. The magenta dots represent the predicted 3D ball trajectory from the back-end.

accurate 3D reconstruction. Prior research has demonstrated the feasibility of training neural networks on synthetically generated data to reconstruct 3D trajectories and spin [8, 22, 31]. However, a significant gap remains in transitioning these models from clean, synthetic inputs to the noisy, sparse, and imperfect detections found in real-world videos. This discrepancy severely limits their practical deployment. This paper addresses these limitations by presenting a comprehensive pipeline designed to robustly infer the 3D trajectory and initial spin of a table tennis ball directly from monocular video footage. To the best of our knowledge, this work constitutes the first learning-based application of a complete pipeline for 3D table tennis analysis. Our solution is built on a novel two-stage framework that solves the fundamental problem of missing 3D ground truth in real-world broadcast videos. We achieve this by dividing the problem into a front-end perception stage and a back-end uplifting stage, which allows us to train each component with different, readily available supervision. We also introduce a new dataset, essential for training our models.

Our main contributions are:

- We develop a state-of-the-art **2D ball detector** leveraging the efficient Segformer++ architecture [20], specifically optimized for processing high-resolution images.

- We introduce a novel **2D table keypoint detector**, enabling precise localization of the table boundaries within diverse video frames. These keypoints provide essential contextual information for the 2D-to-3D uplifting model.
- We present a **2D-to-3D uplifting model** that takes the detected 2D ball trajectory and table keypoints as input, outputting the 3D trajectory and initial spin of the ball. Even though the model is solely trained on physically-correct synthetic data, it achieves zero-shot generalization to real-world scenarios. We specifically adapt this model to robustly handle real-world detection noise, missing detections, and varying frame rates, making it compatible with our presented detectors.
- We introduce the **TTHQ dataset**, a novel high-quality, high-resolution dataset featuring meticulously annotated 2D ball trajectories, annotated table keypoints, spin information, and comprehensive meta information, all sourced from publicly available YouTube videos. This dataset is instrumental for training and evaluating our 2D detectors and for benchmarking the integrated pipeline.

## 2. Related Work

### Ball and Table Keypoint Detection

General object detectors [4, 32, 34] can be adapted for ball detection, but heatmap-based methods, common in 2D pose estimation [25, 29, 42, 44], have become the de facto standard for ball detection. The TrackNet model family [6, 19, 38] and the state-of-the-art WASB [39], demonstrate strong performance. In addition to the ball position, we also want to detect specific table keypoints in the image, which is also sometimes performed for camera calibration in sports analytics [16, 26, 35]. However, existing methods in table tennis [10, 15] often lack the precision and comprehensive set of points required for direct integration into our pipeline. We address these limitations by leveraging the **Segformer++** architecture [20]. This modern, transformer-based approach is uniquely suited for our task due to its efficiency in processing **high-resolution images**, which is crucial for capturing tiny ball details and thin table edges. It is trained using a heatmap-based approach, which proves exceptionally effective for precise ball localization.

### 3D Trajectory and Spin Estimation

Reconstructing 3D trajectories from monocular video is challenging. While controlled multi-camera setups offer high accuracy via triangulation [27, 30, 33, 43], they are impractical for broadcast footage. Monocular methods often rely on physics-based model fitting to observed 2D trajectories [5, 10, 15, 18, 24] or single-frame estimates using cues like observed ball size or height [2, 3, 21, 41], but these approaches are susceptible to errors from inaccurate identification of key events, explicit camera calibration or insufficient video quality. The use of deep learning has

shown great promise in overcoming these issues by directly predicting 3D trajectories [8, 31]. Especially the work of *Kienzle et al.* [22] is of great importance for this paper, demonstrating that a network trained only on synthetic data can achieve zero-shot generalization to real-world footage. However, this is only a proof-of-concept, not considering the front-end detection task and, thus, lacking the necessary robustness for a practical application. Their approach is not designed to handle **real-world imperfections** such as imperfect 2D detector outputs, missing detections due to occlusions, and varying video frame rates. Our work directly addresses this gap. We **re-engineer and adapt** this foundational uplifting model to function with imperfect real-world inputs and integrate it into a complete application pipeline. We also expand the synthetic training dataset to include a wider range of match situations, such as serves and faults, which is essential for building a truly useful application.

Estimating ball spin is equally vital for player analysis but remains a significant challenge due to its indirect observability. Existing solutions often rely on specialized hardware such as event cameras [13, 28] or high-speed cameras [12]. To our knowledge, [22] introduced the first straightforward learning-based method to estimate ball spin from standard monocular video. Building upon this, we further adapt and improve this spin estimation capability, integrating it robustly within our pipeline.

### Datasets

Robust applications are highly dependent on high-quality datasets. Existing table tennis datasets like TTST [22] and Blurball [14] have been instrumental in previous research, but they lack the necessary scale, resolution, or annotation diversity to train and validate a complete, robust pipeline. To solve this, we introduce the **TTHQ dataset**. This is the first comprehensive, high-resolution dataset sourced from real-world broadcast footage that provides all the necessary annotations - 2D ball positions, table keypoints, and spin labels - in sufficient quantity to train and evaluate an end-to-end pipeline. This new dataset is a crucial contribution that will enable future research in this domain.

## 3. Method

### 3.1. General Overview

The lack of 3D supervision in real-world broadcast videos presents a major challenge for learning the ball’s 3D trajectory and spin. We overcome this core problem by proposing a complete, end-to-end pipeline that divides the task into two distinct stages, each solvable with available supervision. An overview of our method is given in Figure 2.

The process begins with the **front-end stage**, which handles the perception and data extraction. Our robust 2D ball detection module is applied to each frame  $n$  at time  $t_n$ , identifying the ball’s 2D position  $\vec{r}_{2D}(t_n)$ . This process yields

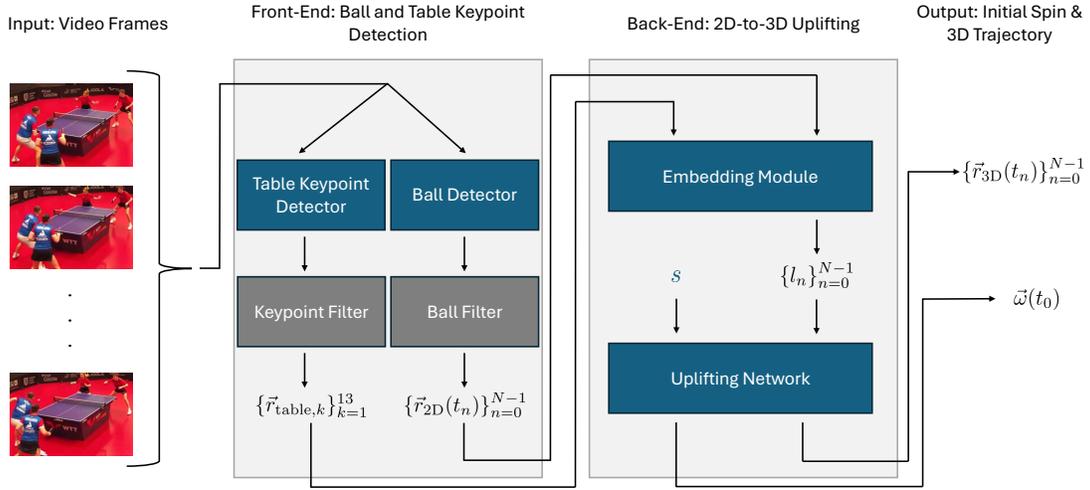


Figure 2. Overview of the proposed pipeline. In the **front-end stage**, we detect the 2D ball position and localize the 13 table keypoints in each frame  $n$  at time  $t_n$ . After robust filtering, we obtain a clean 2D ball trajectory  $\{\vec{r}_{2D}(t_n)\}_{n=0}^{N-1}$  with  $N$  being the number of frames in the trajectory. As we assume a static camera, we obtain a single, time-independent set of table keypoints  $\{\vec{r}_{\text{table},k}\}_{k=1}^{13}$  after filtering. In the **back-end stage**, the coordinates are embedded into a location token  $l_n$  for each timestep  $t_n$ , a learnable spin token  $s$  is prepended, and the sequence is then processed by the uplifting network to predict the 3D trajectory  $\{\vec{r}_{3D}(t_n)\}_{n=0}^{N-1}$  and the initial spin  $\vec{\omega}(t_0)$ . The blue color represents modules with learnable parameters.

a sequence of 2D ball detections given by  $\{\vec{r}_{2D}(t_n)\}_{n=0}^{N-1}$  with  $N$  being the number of frames in the trajectory. Concurrently, the 2D table keypoint detection module localizes the 13 characteristic table points in each frame. To enhance accuracy and stability against detection noise, we apply effective filtering techniques to remove erroneous detections. Moreover, as the camera in table tennis is usually static, we merge the table keypoint detections over time to obtain a single, high-quality set of table keypoints  $\{\vec{r}_{\text{table},k}\}_{k=1}^{13}$  for the entire trajectory.

In the subsequent **back-end stage**, this processed 2D data is fed into our 2D-to-3D uplifting network. The network takes the 2D ball trajectory  $\{\vec{r}_{2D}(t_n)\}_{n=0}^{N-1}$  and the table keypoints as input to predict the initial spin  $\vec{\omega}(t_0) \in \mathcal{R}^3$  and the full 3D trajectory  $\{\vec{r}_{3D}(t_n)\}_{n=0}^{N-1}$  in metric world coordinates. This modular approach is central to our solution: we train the front-end components on our newly created TTHQ dataset, which provides extensive manually-annotated 2D supervision, while the back-end uplifting network is trained solely on synthetic data, similar to [22]. This deliberate separation into two stages solves the problem of missing 3D ground truth for real-world videos, making our complete system a practical application.

### 3.2. Ball Detection

Precise 2D ball detection is a crucial prerequisite for accurate 3D trajectory reconstruction. This task is particularly challenging due to the ball’s tiny size, high speed, and susceptibility to motion blur and occlusions. To address these challenges, we require a model that is both computationally

efficient for high-resolution imagery and robust to these visual artifacts. While existing methods like WASB [39] use the HRNet architecture [37], which is effective but computationally expensive, we instead adopt the more efficient Segformer++ architecture [20]. This transformer-based model leverages token merging [1] to reduce computational cost, enabling us to process high-resolution inputs, which are essential to prevent the ball from vanishing due to downsampling.

To provide the network with vital temporal context for challenging cases like motion blur and occlusions, we take three consecutive images of height  $H$  and width  $W$  as input. These images are concatenated along the channel dimension to form a 9-channel tensor  $I \in \mathbb{R}^{H \times W \times 9}$ . The network then outputs a heatmap where each pixel value indicates the confidence of the ball’s presence in the central frame. We obtain the 2D ball position by first locating the pixel with the highest confidence and then refining its position by fitting a 2D Gaussian to the local neighborhood. Sequentially applying this network to each triple of input frames yields a complete sequence of 2D ball detections given by  $\{\vec{r}_{2D}(t_n)\}_{n=0}^{N-1}$ .

To ensure a clean trajectory, which is crucial for the subsequent uplifting, we present a robust filtering technique to mitigate the impact of false detections and occlusions. Even highly performant models sometimes detect sporadic false positives, such as a player’s feet or paddle, or yield unreliable predictions during occlusions. Common temporal filtering techniques are insufficient for this task, as false positives often occur for multiple consecutive frames, e.g.

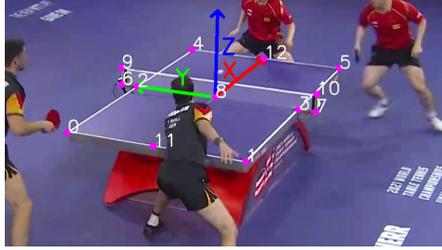


Figure 3. Definition of the 13 Table Keypoints and illustration of the world coordinate system axes.

tracking the paddle movement and leading to wrong trajectories. To address this, we utilize an auxiliary model based on the HRNet architecture. Since the two architectures tend to make different types of errors, we leverage their disagreement to filter out erroneous predictions. A detection is only considered valid if both models’ predictions are in close agreement. If the models agree, we keep the Segformer++ prediction due to its high accuracy. In cases where the models disagree, we simply discard the detection. This simple yet effective filtering step removes nearly all false positives, providing a highly reliable, albeit potentially incomplete, 2D trajectory that is essential for the stable performance of our back-end uplifting network. We provide visualizations of the filtering in Figure 7 of the supplementary material.

### 3.3. Table Keypoint Detection

The table keypoints are a critical input for our uplifting model as they provide a stable spatial reference, enabling the network to understand the camera’s perspective and the ball’s position relative to the table. This approach bypasses the need for explicit camera calibration, which is often a source of error and complexity in other methods. We detect 13 characteristic keypoints, including the four corners, midpoints of each side, and key points along the net. These keypoints are illustrated in Figure 3.

Similar to our ball detection module, we employ the Segformer++ architecture [20] for this task. However, since the table is static, we use a single frame instead of a sequence of 3 frames as input, which is sufficient to capture the required spatial information. The network outputs a heatmap for each of the 13 keypoints, from which we extract their 2D positions using the same Gaussian-fitting approach as for ball detection. Since the different table corners and edges can appear visually similar, the global context of the entire table is essential for accurate keypoint identification, thus, the global receptive field of the Segformer++ architecture is particularly beneficial for this task.

Despite the table’s characteristic appearance, detecting its keypoints is challenging due to frequent occlusions by players, paddles, or the ball itself. We assume a static camera setup during each trajectory, which is standard in table tennis broadcasts, allowing us to aggregate keypoint detec-

tions over time to obtain a highly reliable set of keypoints. We implement a two-step filtering process. First, we use a cross-architectural filtering technique similar to the one described in Section 3.2. We train an auxiliary model based on the HRNet architecture, and compare its predictions with the Segformer++ model. Since these models have different failure modes, detections are only considered valid if the two models are in close agreement. Second, we implement a temporal clustering approach to aggregate and refine detections over the entire trajectory. As some keypoints may be occluded for extended periods, frequently for more than 50% of the trajectory duration, we cannot rely on simple temporal smoothing techniques. Therefore, we cluster the valid detections over time for each keypoint using DBSCAN [9]. The center of the largest cluster is taken as the final, robust keypoint position. If no cluster is found, the keypoint is considered undetectable for the given video. Consequently, after this robust filtering, we obtain a single time-independent, high-quality set of table keypoints  $\{\vec{r}_{\text{table},k}\}_{k=1}^{13}$  for the trajectory. An illustration of this filtering is provided in Figure 8 of the supplementary material.

### 3.4. 2D-to-3D Uplifting

The back-end stage of our pipeline, which performs the 2D-to-3D uplifting, is the central component responsible for reconstructing the 3D data from our front-end 2D detections. It is designed to be trained solely on synthetic data, which provides perfect 3D ground truth for supervision, thus solving the problem of missing real-world 3D labels. We build upon the method from *Kienzle et al.* [22], which demonstrates that a network trained on synthetic data can generalize to real-world scenarios. Given a 2D ball trajectory and the table keypoints, the model predicts the 3D trajectory and initial spin of the ball. Since they do not consider the front-end detection task, their approach handles manually-annotated 2D inputs, which are perfect and complete. This is a significant limitation for practical applications, as real-world detections are often imperfect and incomplete due to occlusions and detection errors. Therefore, we make three key architectural modifications to extend this foundation into a robust, practical application.

#### Robustness to Missing Detections

Our front-end filtering can result in missing detections due to occlusions or false positives. To handle this naturally, we make use of the network’s transformer-based architecture. The attention mechanism of the transformer allows it to process a sequence of inputs of varying lengths, so we can simply remove the filtered-out ball detections and the network can still process the remaining tokens. However, the time information between consecutive detections is crucial for accurate trajectory and spin estimation. Therefore, we employ a custom Rotary Positional Embedding (RoPE) [36] which encodes the exact time stamps of each detection, as

described in more detail below. Consequently, the network can naturally understand the time difference between consecutive detections, even when some are missing. To bridge the gap between training on synthetic data and applying it to real-world scenarios, we simulate missing ball detections during training, increasing the network’s robustness against these common failure cases.

Similarly, we encode the table keypoints using a custom transformer-based embedding module, which allows us to handle varying numbers of visible keypoints, e.g. due to occlusions. We also randomly drop keypoints during training to improve robustness to missing keypoint detections.

### Handling of Varying Frame Rates

Unlike prior work, our system must handle videos with a wide range of frame rates. To enable the network to understand the time difference between consecutive detections, we use a custom Rotary Positional Embedding (RoPE) [36]. Instead of applying a fixed rotation for each position in the input sequence, we rotate each token proportionally to its exact time stamp  $t_n$ . This naturally encodes the time information into the attention mechanism. Further details are provided in Section B of the supplementary material. We randomly sample frame rates during training to improve robustness at test time.

### Network Architecture and Data Flow

Our back-end pipeline begins with an embedding module consisting of 4 transformer blocks as illustrated in Figure 4. We designed this module to process the embedded 2D ball position  $\vec{r}_{2D}(t_n)$  at time  $t_n$  and all visible table keypoints as input. The transformer blocks distill the context from the keypoints into a single **location token**  $l_n \in \mathcal{R}^d$  for each ball detection. By calculating an embedding for each detection and prepending a learnable **spin token**  $s \in \mathcal{R}^d$  to the sequence of location tokens  $\{l_0, \dots, l_{N-1}\}$ , we obtain the final input sequence for the uplifting network.

The uplifting network itself is a two-stage transformer as shown in Figure 5. First, a 2D-to-3D uplifting is performed by  $L - 4$  transformer blocks, followed by a simple MLP head to compute the 3D trajectory  $\{\vec{r}_{3D}(t_n)\}_{n=0}^{N-1}$ . Then, the motion information is distilled into the spin token by 4 additional transformer blocks, followed by another MLP head to predict the initial spin vector  $\vec{\omega}(t_0)$ . This modular design, together with our targeted architectural adaptations, allows our back-end to be exclusively trained on synthetic data and still generalize to real-world footage, effectively solving the challenge of unavailable 3D supervisions.

## 4. Data and Metrics

### TTHQ Dataset

To train and evaluate our method, we introduce the TTHQ (Table Tennis High Quality) dataset, which is created to ad-

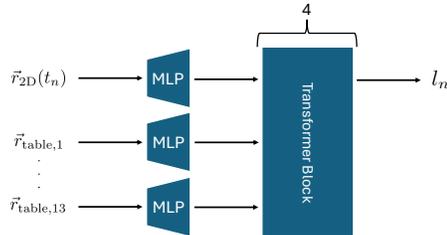


Figure 4. **Embedding Module.** The detected ball position  $\vec{r}_{2D}(t_n)$  at time  $t_n$  and all visible table keypoints are projected into a higher dimensional space by a 2-layer MLP and then processed by a 4-block transformer. Finally, only the token corresponding to the ball position is kept as location token  $l_n$ .

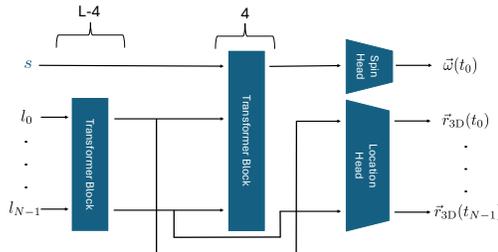


Figure 5. **Uplifting Network.** The input consists of the learnable spin token  $s$  and the sequence of location tokens  $\{l_0, \dots, l_{N-1}\}$  obtained from the embedding module. The first stage consisting of  $L - 4$  transformer blocks computes the 3D trajectory. The second stage consisting of 4 transformer blocks computes the initial spin. To obtain the final three-dimensional output vectors, we apply a small 3-layer MLP as head for both the trajectory and the spin.

dress the shortcomings of existing datasets for our specific task. It consists of high-resolution ( $1920 \times 1080$ ) broadcast videos of professional, semi-professional, and amateur table tennis matches sourced from YouTube, covering a wide range of conditions, including different camera angles, lighting situations, and player skill levels. The TTHQ dataset includes 14 videos of individual matches and 5 highlight videos, where a highlight video is a compilation of selected scenes from multiple matches, increasing the diversity of scenarios significantly. We provide precise 2D annotations for both ball and table keypoint detection modules. The 2D ball position is manually annotated in 9 092 frames, ensuring a diverse set of scenarios with varying ball speeds and motion blur levels. We also annotate the 13 2D table keypoints in 257 frames, placing special emphasis on highlight videos and scenes with varying levels of occlusion to ensure the network can generalize despite the relatively small number of annotations.

For comprehensive evaluation, we designate 3 of the videos for validation and testing. For these videos, we also provide binary spin class labels (topspin or backspin) for 57 trajectories, which we annotated similarly to [22]. This allows us to evaluate the full pipeline’s spin prediction accuracy in a realistic setting.

## Synthetic Data

To train our uplifting network, we generate a large synthetic dataset using the MuJoCo physics engine [40]. We extend the original training data from [22] by including a much wider variety of trajectories. This includes trajectories from both sides, along with serves and various fault types. We also vary the ball’s initial position, velocity, and spin more extensively, and we simulate a wider range of camera parameters to cover more diverse real-world broadcast conditions. Overall, we generate 140 000 synthetic trajectories, with over 4.7 million individual 2D ball positions.

## Existing Datasets

The Blurball dataset [14] consists of 25 videos and provides 50 000 annotated frames with 2D ball positions. While the videos are of lower resolution ( $1280 \times 720$ ) and quality, we use this dataset for pre-training our ball detection module to bootstrap its learning.

The TTST dataset [22] provides 50 trajectories. The 2D ball position is annotated in 1197 frames and the set of 13 table keypoints is annotated once per trajectory. The trajectories are also annotated with a binary spin label (topspin or backspin). We use 16 trajectories for validating the uplifting network. The remaining 34 trajectories are used to test the uplifting network’s performance in a controlled setting.

## Metrics

For our front-end modules, we use the  $ACC@Xpx$  metric, which measures the percentage of correct detections within a radius of  $X$  pixels around the ground truth position.

Given that no 3D ground truth data exists for realistic broadcast videos, a direct evaluation of the back-end’s 3D predictions is not possible. We overcome this challenge by employing a comprehensive set of metrics that indirectly, but robustly, assess our model’s performance. We first evaluate the quality of the predicted 3D trajectory by projecting it back into the image plane. The 2D Reprojection Error (2DRE) is then calculated as the average distance between the reprojected 3D positions and the ground truth 2D positions for each trajectory:

$$2DRE = \frac{1}{N} \sum_{n=0}^{N-1} \left\| \mathcal{P}(\vec{r}_{3D}(t_n)) - \hat{r}_{2D}(t_n) \right\|_2 \quad (1)$$

where  $\mathcal{P} \in \mathcal{R}^{2 \times 3}$  is the camera projection matrix,  $\vec{r}_{3D}(t_n)$  is the predicted 3D position,  $\hat{r}_{2D}(t_n)$  is the ground truth 2D position, and  $N$  is the number of frames in the trajectory. The final reported score is the **mean 2DRE (m2DRE)** computed over all trajectories.

Finally, we evaluate the predicted initial spin using the binary spin annotations in the dataset. Since it is not possible for humans to visually estimate the full 3D spin vector or reliably estimate the strength of the spin, human annotations are inherently limited to the binary classification of topspin vs. backspin. While our network is trained on synthetic data

to predict the full continuous 3D spin vector  $\vec{\omega}_{3D}(t_0) \in \mathcal{R}^3$ , we can only evaluate its performance on real-world data by classifying the spin as topspin if the  $y$ -component of the predicted vector is positive, and as backspin otherwise. We perform this classification using a hard threshold, providing a direct mapping from our continuous prediction to the available ground truth. The performance on this binary task provides compelling evidence that the network has successfully learned the underlying full 3D spin vector. We calculate the **Accuracy (ACC)** and the **macro F1-Score ( $F_1$ )** to measure the classification performance.

## 5. Experiments

In this section, we present a systematic evaluation of our proposed components and the complete pipeline across various datasets. The structure of this section is designed to validate each of our key architectural choices and demonstrate the real-world efficacy of our full system. Implementation and training details can be found in Section A of the supplementary material.

### 5.1. Ball Detection Evaluation

Model	#Params	Input Res.	FPS $\uparrow$	ACC@2px $\uparrow$	ACC@5px $\uparrow$	ACC@10px $\uparrow$
Segformer++ (B0)	$3.7 \cdot 10^6$	$1920 \times 1088$	26	<b>75.0 %</b>	85.9 %	90.3 %
Segformer++ (B2)	$24.7 \cdot 10^6$	$1600 \times 896$	18	<b>75.0 %</b>	87.1 %	<b>91.8 %</b>
WASB (HRNet Small)	$1.5 \cdot 10^6$	$1280 \times 704$	16	72.4 %	<b>87.4 %</b>	91.3 %
VitPose (ViT Small)	$25.9 \cdot 10^6$	$1152 \times 640$	19	38.0 %	50.3 %	52.1 %

Table 1. Comparison of different architectures for **ball detection** and results on the TTHQ test set. The input resolution for each model was chosen individually to ensure comparable GPU RAM usage during training and similar inference speeds during testing. This process allowed us to determine the optimal and maximal input resolution for each architecture. The inference FPS are calculated on a single V100 GPU with batch size 8.

Our core hypothesis for the front-end is that **high-resolution input is paramount for precise ball detection**. Because a table tennis ball is a small, fast-moving object, any loss of visual information can be detrimental. We suggest the **Segformer++** architecture due to its strong performance and efficiency, which allows us to use very high-resolution inputs. To validate this choice, we conduct a comparative evaluation against two distinct architectures: the current state-of-the-art ball detection model **WASB** based on the HRNet architecture and the popular **VitPose** architecture for 2D human pose estimation. More details about the chosen model parameters are given in Section B of the supplementary material.

The results in Table 1 confirm our hypothesis, as the Segformer++ models outperform the other architectures, especially on the strict  $ACC@2px$  metric. This high precision is vital, as it is a prerequisite for accurate 3D trajectory uplifting. While the Segformer++ (B2) model achieves superior overall performance, the smaller B0 variant provides a strong balance of accuracy and speed. The CNN-

based WASB and transformer-based VitPose models perform worse, which we attribute to their lower effective resolution. This confirms our core belief: the ability to process high-resolution images is paramount for accurate table tennis ball detection. Based on these results, we select the **Segformer++ (B2)** model for our final pipeline due to its superior overall performance.

## 5.2. Table Keypoint Detection Evaluation

Model	#Params	Input Res.	FPS $\uparrow$	ACC@2px $\uparrow$	ACC@5px $\uparrow$	ACC@10px $\uparrow$
Segformer++ (B0)	$3.7 \cdot 10^6$	$1920 \times 1088$	26	43.2 %	<b>86.8 %</b>	<b>94.4 %</b>
Segformer++ (B2)	$24.7 \cdot 10^6$	$1600 \times 896$	19	<b>54.3 %</b>	85.3 %	93.0 %
WASB (HRNet Small)	$1.5 \cdot 10^6$	$1280 \times 704$	17	41.1 %	83.8 %	89.3 %
VitPose (ViT Small)	$25.3 \cdot 10^6$	$1152 \times 640$	19	30.0 %	68.5 %	79.7 %

Table 2. Comparison of different architectures for **table keypoint detection** and results on the TTHQ test set. The input resolution for each model was chosen individually to ensure comparable GPU RAM usage during training and similar inference speeds during testing. This process allowed us to determine the optimal and maximal input resolution for each architecture. The inference FPS are calculated on a single V100 GPU with batch size 8.

Our second front-end hypothesis is that additionally to the high-resolution capabilities, a **global receptive field is crucial for distinguishing table keypoints**. The 13 table keypoints are often visually similar, requiring a model to understand the context within the entire table to accurately identify them. This makes a transformer-based architecture with a global receptive field particularly beneficial. We again compare our **Segformer++** architecture against the CNN-based **WASB** and the transformer-based **VitPose** models.

The results, summarized in Table 2, show that the Segformer++ models significantly outperform the other architectures. This performance gap is more pronounced than in the ball detection task. This provides strong evidence for our hypothesis, suggesting that while high resolution is vital, the global context provided by the transformer backbone is essential for accurately distinguishing between the keypoints. Further evidence comes from the improved performance of the VitPose model compared to the ball detection task, even though it still lags behind. Based on these results, we again select the **Segformer++ (B2)** model for our final pipeline due to its superior performance on the critical ACC@2px metric.

## 5.3. Back-End Evaluation

We now turn to the back-end, where our core hypothesis is that our proposed architectural modifications enable the network to **handle varying frame rates and missing detections** while still maintaining high performance. As proof, we perform an ablation study with three distinct model variants on the TTST dataset, with results presented in Table 3. The models we compare are:

- The original architecture from *Kienzle et al.* [22], which uses a different embedding module not designed for miss-

Method	Transforms		Metrics		
	Half FPS	Miss. Det.	ACC $\uparrow$	$F_1$ $\uparrow$	m2DRE $\downarrow$
<i>Kienzle et al.</i> [22]			97.1 %	0.970	2.98 px
Mixed	×	×	<b>100.0 %</b>	<b>1.000</b>	<b>2.49 px</b>
Ours			97.1 %	0.970	3.43 px
<i>Kienzle et al.</i> [22]			76.5 %	0.731	<b>2.71 px</b>
Mixed	✓	×	79.4 %	0.770	3.13 px
Ours			<b>100.0 %</b>	<b>1.000</b>	3.54 px
<i>Kienzle et al.</i> [22]			88.2 %	0.876	24.15 px
Mixed	×	✓	<b>97.1 %</b>	<b>0.970</b>	<b>5.45 px</b>
Ours			<b>97.1 %</b>	<b>0.970</b>	5.56 px
<i>Kienzle et al.</i> [22]			67.7 %	0.598	23.54 px
Mixed	✓	✓	70.6 %	0.646	5.99 px
Ours			<b>97.1 %</b>	<b>0.970</b>	<b>5.75 px</b>

Table 3. Ablation study of the back-end on the TTST dataset [22]. We apply two different transformations to the test set to simulate real-world challenges: The Half FPS transformation drops every second frame to simulate a lower frame rate. The Missing Detections transformation randomly removes 10% of the frames to simulate ball detection failures or frame drops. Additionally, each table keypoint has a 10% chance of being removed.

ing data and a sequence-based positional embedding that does not encode the exact timing of each frame.

- The **mixed architecture**, which uses our proposed transformer-based embedding module but retains the sequence-based positional encoding.
- **Our full architecture**, which combines the robust transformer-based embedding module with the time-proportional RoPE.

To test our hypothesis, we augment the TTST test set with two transformations: The **Half FPS** transformation, which drops every second frame to simulate a lower frame rate. And the **Missing Detections** transformation, which randomly removes 10% of the frames to simulate ball detection failures or frame drops. Additionally, each table keypoint has a 10% chance of being removed to simulate occlusions or table keypoint detection failures.

The results provide clear evidence for our design choices. On the original, unaugmented TTST data, all architectures perform well. However, when faced with the Half FPS transformation, the spin predictions of *Kienzle et al.* degrade significantly. The mixed architecture also suffers a notable performance drop in spin accuracy. This proves that the exact timing of each frame is crucial for accurate spin prediction, and that our time-proportional RoPE is effective in encoding this information. When applying the Missing Detections transformation, the predicted trajectories of *Kienzle et al.* degrade significantly, while the other two architectures maintain high performance. This clearly shows that our transformer-based embedding module, which can naturally process sequences of varying lengths, is effective at handling missing data. Finally, when both transformations are applied, our architecture outperforms the other two in both spin and trajectory metrics, demonstrating its robustness to real-world challenges.

Overall, our proposed architecture performs robustly in all scenarios, yielding high-quality spin and trajectory predictions. This validates our design choices and confirms that our method is a practical solution that can handle the complexities of real-world videos.

#### 5.4. Full Pipeline: Combining Front- and Back-End

Dataset	Table: m2DRE ↓	Ball: m2DRE ↓	Spin: ACC ↑	Spin: $F_1$ ↑
TTHQ	$2.72 \pm 5.71$ px	$12.28 \pm 10.84$ px	89.5 %	0.900
TTST	$5.75 \pm 10.26$ px	$9.41 \pm 16.90$ px	97.1 %	0.974

Table 4. Results of the full pipeline on the TTHQ and TTST datasets. The  $\pm$  indicates the standard deviation computed over all trajectories.

The final and most crucial evaluation is to assess our complete, end-to-end pipeline in a real-world setting. While the individual components have proven their effectiveness, the true test is whether the noise and imperfections of the front-end detections degrade the back-end’s performance significantly. To evaluate this, we apply our full system to 57 TTHQ trajectories with annotated spin labels and additionally calculate results on the 34 TTST test trajectories.

Since the 57 TTHQ test trajectories lack ball and table keypoint annotations, we assess the performance indirectly using reprojection errors. We first estimate the camera matrices from the detected table keypoints using a robust RANSAC-based [11] calibration algorithm [22] to ensure that the calibration is not affected by potential outliers in the keypoint detections. Due to the fixed table geometry, we utilize the known 3D positions of the table keypoints and reproject them back into the image. These reprojected keypoints can then be compared to the detected keypoints to compute a mean 2D reprojection error (m2DRE) in pixels. Similarly, we reproject the predicted 3D ball trajectory back into the image and compare it to the detected 2D ball positions to compute a ball m2DRE. Finally, spin classification is measured directly using the annotated labels.

As shown in Table 4, our pipeline achieves strong performance on both datasets. Both the table and ball reprojection errors are low, indicating that each component of the pipeline works effectively. Especially the spin classification achieves high accuracy and  $F_1$  scores, which confirms that the data from our front-end is of sufficient quality for the back-end to make accurate predictions. When we compare the spin classification results on the TTST dataset with the perfect-input results from Section 5.3, we see a similar high performance. This is a crucial finding, as it proves that using the front-end detections does not significantly degrade the back-end’s spin prediction performance, demonstrating the effectiveness of our complete pipeline as a practical, robust solution for real-world table tennis analysis.



Figure 6. Qualitative results of the **full pipeline** on the TTHQ dataset. The **green** dots represent the valid detections of the front-end for 2D ball positions and table keypoints. The **red** dots represent the reprojected table keypoints inferred from the table geometry and camera calibration. The **blue** dots represent the predicted 3D ball trajectory from the back-end.

## 6. Conclusion

In this work, we have presented a comprehensive pipeline for 3D trajectory and spin estimation from monocular video. Our approach addresses a novel application domain by providing a solution that can accurately analyze table tennis ball motion from unconstrained, noisy broadcast footage. We solve a fundamental challenge in this domain: the lack of 3D ground truth data.

The core of our innovation is a two-stage framework, which provides a clean separation between the 2D perception and 3D uplifting tasks. This modularity is a key contribution, as it allows each stage to be trained with different, readily available supervision: 2D human annotations and 3D synthetic data. Consequently, it eliminates the need for a unified 3D ground truth dataset. Our comparative assessment, including extensive ablation studies, validates the superiority of our architectural choices. We have explicitly shown that our custom time-proportional positional encoding and robust embedding module are essential for handling real-world artifacts like occlusions and varying frame rates.

Ultimately, our work provides a practical, ready-to-use tool for sports analytics. The principles of our two-stage design and our comparative methodology offer a template for similar 3D reconstruction problems in other domains where precise ground truth is a significant barrier.

## References

- [1] Daniel Bolya, Cheng-Yang Fu, Xiaoliang Dai, Peizhao Zhang, Christoph Feichtenhofer, and Judy Hoffman. Token merging: Your ViT but faster. In *International Conference on Learning Representations*, 2023. 3
- [2] Marcello Davide Caio, Gabriel Van Zandycke, and Christophe De Vleeschouwer. Context-aware 3d object localization from single calibrated images: A study of basketballs. In *Proceedings of the 6th International Workshop on Multimedia Content Analysis in Sports*, page 49–54, 2023. 2
- [3] Jordan Calandre, Renaud Péteri, Laurent Mascarilla, and Benoit Tremblais. Extraction and analysis of 3d kinematic parameters of table tennis ball from a single camera. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 9468–9475, 2021. 2
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European Conference on Computer Vision (ECCV)*, page 213–229, 2020. 2
- [5] 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. *Soft Computing*, 20(3): 204–216, 2009. 2
- [6] Yu-Jou Chen and Yu-Shuen Wang. Tracknetv3: Enhancing shuttlecock tracking with augmentations and trajectory rectification. In *Proceedings of the 5th ACM International Conference on Multimedia in Asia*, 2024. 2
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. 11
- [8] Morten Holck Ertner, Sofus Schou Konglevoll, Magnus Ibh, and Stella Graßhof. Synthnet: Leveraging synthetic data for 3d trajectory estimation from monocular video. In *Proceedings of the 7th ACM International Workshop on Multimedia Content Analysis in Sports*, page 51–58, 2024. 1, 2
- [9] Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, pages 226–231, 1996. 4, 12
- [10] Daniel Etaat, Dvij Kalaria, Nima Rahmanian, and Shankar Sastry. Latte-mv: Learning to anticipate table tennis hits from monocular videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2025. 2
- [11] Martin A. Fischler and Robert C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. 8
- [12] Thomas Gossard, Jonas Tebbe, Andreas Ziegler, and Andreas Zell. Spindoe: A ball spin estimation method for table tennis robot. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5744–5750, 2023. 2
- [13] Thomas Gossard, Julian Krismer, Andreas Ziegler, Jonas Tebbe, and Andreas Zell. Table tennis ball spin estimation with an event camera. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 3347–3356, 2024. 2
- [14] Thomas Gossard, Filip Radovic, Andreas Ziegler, and Andrea Zell. Blurball: Joint ball and motion blur estimation for table tennis ball tracking, 2025. arXiv:2509.18387. 2, 6, 11, 13
- [15] Thomas Gossard, Andreas Ziegler, and Andreas Zell. TT3D: Table tennis 3d reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2025. 2
- [16] Marc Gutiérrez-Pérez and Antonio Agudo. No bells just whistles: Sports field registration by leveraging geometric properties. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 3325–3334, 2024. 2
- [17] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. Masked autoencoders are scalable vision learners. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15979–15988, 2022. 11
- [18] Chen Huang, 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 (MTA)*, 60, 2012. 2
- [19] Yu-Chuan Huang, I-No Liao, Ching-Hsuan Chen, Tsi-Uí Ik, and Wen-Chih Peng. Tracknet: A deep learning network for tracking high-speed and tiny objects in sports applications\*. *2019 16th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pages 1–8, 2019. 2
- [20] Daniel Kienzle, Marco Kantonis, Robin Schön, and Rainer Lienhart. Segformer++: Efficient token-merging strategies for high-resolution semantic segmentation. *IEEE International Conference on Multimedia Information Processing and Retrieval (MIPR)*, 2024. 1, 2, 3, 4
- [21] Daniel Kienzle, Julian Lorenz, Katja Ludwig, and Rainer Lienhart. Towards learning monocular 3d object localization from 2d labels using the physical laws of motion. *Proceedings of the International Conference on 3D Vision 2024 (3DV)*, 2024. 2
- [22] Daniel Kienzle, Robin Schön, Rainer Lienhart, and Shin’Ichi Satoh. Towards ball spin and trajectory analysis in table tennis broadcast videos via physically grounded synthetic-to-real transfer. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2025. 1, 2, 3, 4, 5, 6, 7, 8, 11, 12, 13
- [23] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations (ICLR)*, 2015. 11
- [24] Paul Liu and Jui-Hsien Wang. Monotrack: Shuttle trajectory reconstruction from monocular badminton video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 3513–3522, 2022. 2

- [25] Katja Ludwig, Daniel Kienzle, and Rainer Lienhart. Recognition of freely selected keypoints on human limbs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 3530–3538, 2022. 2
- [26] Adrien Maglo, Astrid Orcesi, and Quoc-Cuong Pham. Kalicalib: A framework for basketball court registration. In *Proceedings of the 5th International ACM Workshop on Multimedia Content Analysis in Sports*, pages 111–116, 2022. 2
- [27] Andrii Maksai, Xinchao Wang, and Pascal Fua. What players do with the ball: A physically constrained interaction modeling. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 972–981, 2016. 2
- [28] Takuya Nakabayashi, Kyota Higa, Masahiro Yamaguchi, Ryo Fujiwara, and Hideo Saito. Event-based ball spin estimation in sports. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 3367–3375, 2024. 2
- [29] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European Conference on Computer Vision (ECCV)*, pages 483–499, 2016. 2
- [30] Pascaline Parisot and Christophe Vleeschouwer. Consensus-based trajectory estimation for ball detection in calibrated cameras systems. *Journal of Real-Time Image Processing*, 16, 2019. 2
- [31] Puntawat Ponglertnapakorn and Supasorn Suwajanakorn. Where is the ball: 3D ball trajectory estimation from 2d monocular tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2025. 1, 2
- [32] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 779–788, 2016. 2
- [33] 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. 2
- [34] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems (NIPS)*, 2015. 2
- [35] Jia Cheng Shang, Yuhao Chen, Mohammad Javad Shafiee, and David A. Clausi. Rink-agnostic hockey rink registration. In *Proceedings of the 6th International Workshop on Multimedia Content Analysis in Sports*, pages 73–81, 2023. 2
- [36] Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. 4, 5, 12
- [37] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 3
- [38] Nien-En Sun, Yu-Ching Lin, Shao-Ping Chuang, Tzu-Han Hsu, Dung-Ru Yu, Ho-Yi Chung, and Tsi-Uí Īk. Tracknetv2: Efficient shuttlecock tracking network. In *2020 International Conference on Pervasive Artificial Intelligence (ICPAI)*, pages 86–91, 2020. 2
- [39] Shuhei Tarashima, Muhammad Abdul Haq, Yushan Wang, and Norio Tagawa. Widely applicable strong baseline for sports ball detection and tracking. In *British Machine Vision Conference (BMVC)*, 2023. 2, 3
- [40] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033, 2012. 6
- [41] Gabriel Van Zandycke and Christophe De Vleeschouwer. 3d ball localization from a single calibrated image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 3472–3480, 2022. 2
- [42] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Convolutional Pose Machines. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4724–4732, 2016. 2
- [43] Qingyu Xiao, Zulfiqar Zaidi, and Matthew Gombolay. Multi-camera asynchronous ball localization and trajectory prediction with factor graphs and human poses. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 13695–13702, 2024. 2
- [44] Yufei Xu, Jing Zhang, Qiming Zhang, and Dacheng Tao. ViTPose: Simple vision transformer baselines for human pose estimation. In *Advances in Neural Information Processing Systems*, 2022. 2

# Uplifting Table Tennis: A Robust, Real-World Application for 3D Trajectory and Spin Estimation

## Supplementary Material

### A. Training Details

The front-end modules are trained with a learning rate of  $10^{-3}$  using the Adam optimizer [23] and a batch size of 4. For ball detection, we apply random flipping, rotation, translation, cropping, and color jittering as augmentations. For table keypoint detection, we do not use flipping, as it would change the semantics of the keypoints. Instead, we apply an additional random perspective transformation to simulate different camera angles. Especially for table keypoint detection, we found that augmentations are crucial to achieve generalization ability since the TTHQ dataset is diverse but the number of annotated frames is still limited. We save the model with the best  $ACC@5px$  metric on the validation set during training and use it for testing. Similar to 2D human pose estimation, we simply train the models to predict gaussian heatmaps centered at the ground truth ball and keypoint positions. The ground truth heatmaps are created at a resolution of  $1920 \times 1080$ , and we use  $\sigma = 6px$  for the Gaussian. Comparing the predicted and ground truth heatmaps, we use a simple  $L_2$  loss. All evaluations are performed at a resolution of  $1920 \times 1080$  and the  $ACC@Xpx$  metric is also computed at this resolution.

The ball detection backbones are first initialized with ImageNet [7] (Segformer++ and WASB) or MAE [17] (Vit-Pose) weights, then pre-trained on the Blurball dataset [14], and finally fine-tuned on the TTHQ training set. The table keypoint detection backbones are also initialized similarly, but only fine-tuned on the TTHQ training set, as the Blurball dataset does not contain sufficiently diverse table keypoint annotations.

For the training of the back-end module, we stick closely to the training procedure of Kienzle *et al.* [22], but we incorporate additional augmentations of the synthetic training data. We simulate missing detections by randomly dropping 5% of the ball and table keypoint detections during training. Moreover, we randomly sample the frame rate between 20 FPS and 60 FPS during training. A further difference to the Kienzle *et al.* training procedure is that we use our extended synthetic dataset, which contains 140000 trajectories instead of 50000 and includes more diverse match scenes like serves. In contrast to Kienzle *et al.*, we do not validate on the synthetic data, but instead validate on the validation set of the TTST dataset [22]. We choose the final model as compromise between the best  $F_1$  score for spin prediction and the best m2DRE for trajectory prediction. If multiple epochs achieve the same maximum  $F_1$  score, which happens regularly, we choose the one with the best m2DRE



Figure 7. Ball detection filtering. The green dots are the valid detections of the Segformer++ (B2) model after filtering. The crimson dots are the false positives of the Segformer++ (B2) model that were removed. The orange dots are the predictions of the WASB model for the filtered frames.

among them. To allow for a meaningful interpretation of the m2DRE metric, we evaluate all results at a resolution of  $1920 \times 1080$ , which ensures compatibility with our front-end evaluations. Our models predict the spin in the global world coordinate system, but for evaluation and extraction of the spin class, we transform the spin into the local ball coordinate system as described in [22].

### B. Architectures

**Front-End Architectures:** To ensure a fair comparison of front-end architectures, we tuned 2 significant parameters: The input resolution and the model size. We chose these

parameters such that each model can be trained on a single Nvidia V100 GPU with 32 GB of memory. Moreover, the models are tuned such that their inference speed is as similar as possible. We provide the exact parameters in Table 1 and Table 2.

**Back-End Architecture:** We stick closely to [22] when designing the back-end architecture. We use an embedding size of  $d = 128$  and  $L = 16$  transformer layers with 4 attention heads each. The Spin Head and Location Head both consist of 3 fully connected layers with ReLU activations in between.

**Positional Encoding:** Unlike standard transformer-based architectures that rely on a fixed positional encoding based on the token’s index, our system must handle inputs with varying time stamps and missing detections. To address this, we use a custom Rotary Positional Embedding (RoPE) that encodes the exact time information of each detection into the attention mechanism. The core of our implementation is a modification to the standard RoPE approach, where we replace the discrete position index  $n$  with a highly granular, discretized timestamp  $p_n$ :

$$p_n = \text{round}(t_n / \Delta t) \quad (2)$$

with  $\Delta t = 2$  ms. This discretization is so fine-grained that it effectively represents the exact time for all realistic video frame rates. By replacing  $n$  with  $p_n$  in the original RoPE formulation [36], we obtain the following rotation matrix

$$R = \begin{pmatrix} \cos(p_n \cdot \theta_m) & -\sin(p_n \cdot \theta_m) \\ \sin(p_n \cdot \theta_m) & \cos(p_n \cdot \theta_m) \end{pmatrix} \quad (3)$$

where  $\theta_m = \frac{1}{10000^{2m/d}}$  is the angle,  $d$  is the feature dimension and  $m$  indexes the feature dimensions. For a pair of elements from the feature vector  $\mathbf{x} = (x_1, \dots, x_d)$ , the rotated components are computed as:

$$\begin{pmatrix} x'_{2m} \\ x'_{2m+1} \end{pmatrix} = \begin{pmatrix} \cos(p_n \cdot \theta_m) & -\sin(p_n \cdot \theta_m) \\ \sin(p_n \cdot \theta_m) & \cos(p_n \cdot \theta_m) \end{pmatrix} \begin{pmatrix} x_{2m} \\ x_{2m+1} \end{pmatrix}. \quad (4)$$

This time-based encoding allows the attention mechanism to inherently learn the temporal relationships between detections, regardless of their spacing in the input sequence.

### C. Ball & Table Keypoint Filtering

The back-end stage of our pipeline is designed to be robust against missing ball and table keypoint detections. However, it is relatively sensitive to false positive detections, which is why we apply a filtering step to the raw detections of the front-end. For ball detections, we compare the predictions of the Segformer++ (B2) model with the predictions of the WASB model and only keep detections with a maximum distance of 20 px to a WASB detection. Because the model architectures are fundamentally different,



Figure 8. Table keypoint detection filtering. The green dots are the valid detections after filtering. The crimson dots are the false positives of the Segformer++ (B2) model that were removed. The orange dots are the predictions of the WASB model for the filtered frames.

they usually do not make the same false positive detections, allowing us to filter them out with very high precision. We illustrate this filtering step in Figure 7 for three different example trajectories in the TTHQ test set.

We use this two-model based filtering approach instead of more established filtering methods like Kalman filtering or polynomial fitting, because they fail in a significant failure case: Sometimes, the feet or paddle movements of the players are falsely detected as balls for multiple consecutive frames. These false positive detections cannot be easily differentiated from real ball trajectories, which is why the established filtering methods fail. However, our two-model based filtering approach is very robust against this failure case, resulting in a very high precision of the final ball detections.

For table keypoint detections, we apply a similar filtering step. We only keep the detections of the Segformer++ (B2) model that are within a distance of 10 px to a WASB detection. Additionally, we utilize the static nature of the table keypoints and apply the DBSCAN clustering algorithm [9] to the remaining detections of each keypoint over the entire trajectory. The final filtered detections are then the cluster centers of the largest cluster for each keypoint. This results in an extremely robust filtering of false positive detections. We illustrate this filtering step in Figure 8 for two different example trajectories in the TTHQ test set.

### D. Comparison of Datasets

We provide a comparison of existing table tennis datasets in Table 5.

Dataset	Resolution	# Balls	# Table Keypoints	# Spin Class
TTHQ (ours)	1920×1080	9092	257	57
TTST [22]	1920×1080	1197	50	50
Blurball [14]	1280×720	50000	50000*	0

Table 5. Comparison of table tennis datasets. We report the video resolution, number of frames with ball annotations, number of frames with table keypoint annotations, and number of trajectories with spin class annotations. The \* indicates, that table keypoints are not annotated, but can be generated from the provided homography. However, while many frames with table annotations can be obtained this way, the variety of the frames is extremely low, because the dataset is sourced from a limited number of videos.

The **Blurball dataset** [14] has the largest number of ball annotations, and using the provided homographies, pseudo table keypoints can be generated. As the dataset is sourced from 24 different videos, the table keypoints lack diversity and are, thus, not suitable for training a robust table keypoint detection model. However, since ball detection training data needs to capture a wide variety of conditions due to the dynamic nature of the ball, the Blurball dataset is still a valuable resource for training ball detection models. Due to the low resolution of the videos, we do not use this dataset for fine-tuning, but utilize it for pre-training our ball detection models.

The **TTST dataset** [22] has a high resolution and provides annotations for balls, table keypoints, and spin. Because the scenes are not very diverse and the number of annotated frames is relatively low, the dataset is not suitable for training robust front-end models. However, it is still a valuable resource for validating the performance of our pipeline. Especially valuable are the dense annotations for each trajectory, which allow us to evaluate the performance of the back-end model independently of the front-end.

Our **TTHQ dataset** provides high resolution videos with a high variety of scenes and numerous annotations for balls, table keypoints, and spin. This makes it the first dataset that is suitable for training robust front-end models. Moreover, we can also utilize the dataset for the validation of the spin predictions. While the annotations are not dense like in TTST, the high variety of scenes and situations makes it an interesting benchmark for evaluating the overall performance of the full pipeline.

## E. Reproducibility and Open Resources

To foster reproducibility and facilitate further research in the field, we provide our code at <https://kiedani.github.io/WACV2026/>.