

# RacketVision: A Multiple Racket Sports Benchmark for Unified Ball and Racket Analysis

Linfeng Dong<sup>1,2\*</sup>, Yuchen Yang<sup>3,2</sup>, Hao Wu<sup>4,2</sup>, Wei Wang<sup>2</sup>, Yuenan Hou<sup>2</sup>,  
Zhihang Zhong<sup>2†</sup>, Xiao Sun<sup>2†</sup>

<sup>1</sup>Zhejiang University

<sup>2</sup>Shanghai AI Laboratory

<sup>3</sup>Fudan University

<sup>4</sup>University of Science and Technology of China

## Abstract

We introduce RacketVision, a novel dataset and benchmark for advancing computer vision in sports analytics, covering table tennis, tennis, and badminton. The dataset is the first to provide large-scale, fine-grained annotations for racket pose alongside traditional ball positions, enabling research into complex human-object interactions. It is designed to tackle three interconnected tasks: fine-grained ball tracking, articulated racket pose estimation, and predictive ball trajectory forecasting. Our evaluation of established baselines reveals a critical insight for multi-modal fusion: while naively concatenating racket pose features degrades performance, a Cross-Attention mechanism is essential to unlock their value, leading to trajectory prediction results that surpass strong unimodal baselines. RacketVision provides a versatile resource and a strong starting point for future research in dynamic object tracking, conditional motion forecasting, and multi-modal analysis in sports.

## Introduction

Racket sports, typically represented by badminton, tennis, and table tennis, have garnered widespread global participation and attract research for performance analysis (Kulkarni et al. 2023; D’Ambrosio et al. 2024; Gossard, Ziegler, and Zell 2025). These sports encompass structurally defined computer vision tasks, while presenting perception challenges due to the rapid motion of both the ball and players, as well as the complex human-object interactions inherent in racket-based gameplay. However, existing datasets (Huang et al. 2019; Sun et al. 2020; Tarashima et al. 2023) focus narrowly on ball tracking within a single sport at a time, falling short in two critical aspects: 1) They fail to leverage shared ball motion patterns across different sports. 2) Despite the racket being a central component, racket-specific annotations and analysis are lacking. This is crucial not only for sports analysis but also for complex neural avatar modeling (Chen et al. 2024; Xu et al. 2025; Zhan et al. 2025a,b). These shortcomings limit the development of comprehensive racket sports analysis methods.

\*Work done during internship at Shanghai AI Laboratory.

†Corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

To address this gap, we introduce RacketVision, a multiple racket sports benchmark for unified ball and racket analysis. RacketVision first expands the range of sports types for unified model training, aiming to uncover shared priors across racket sports. Specifically, it comprises a collection of 1,672 video clips (435,179 frames, 12,755 seconds) spanning badminton, tennis, and table tennis. In task design, RacketVision progressively proposes three tasks with corresponding annotations, enabling a more comprehensive decomposition of racket sport analysis. Beyond the existing *ball tracking* task, RacketVision defines racket keypoints and supports a novel *racket pose estimation* task. RacketVision further proposes an integrative task, *ball trajectory prediction*, empowering downstream applications, such as tactic analysis (Wang et al. 2024; Yuchen et al. 2025b), robotics (D’Ambrosio et al. 2024; Ma et al. 2025), etc.

In our evaluation, we establish extensive baselines for three tasks and analyze the impact of multi-sport training and multi-modal information under various fusion strategies. Our experiments reveal key insights that, training on all three sports generally enhances model generalization on perception tasks. More importantly, we uncover a nuanced relationship between multi-modal data and performance in trajectory prediction: a naive concatenation of racket pose features was found to be detrimental, performing worse than a ball-only baseline. However, by introducing a sophisticated Cross-Attention fusion mechanism, our LSTM-based model successfully leverages the racket information, ultimately outperforming the strong ball-only baseline across all three sports. This highlights that the value of racket pose data is critically dependent on the fusion architecture’s ability to intelligently integrate contextual cues.

Our contributions are threefold:

- We present RacketVision, a large-scale, multi-sport benchmark with detailed annotations for balls and rackets, supporting cross-sport analysis.
- We define three interconnected tasks, formulating key challenges of computer vision in sports analytics.
- We establish strong baseline solutions and conduct detailed evaluations, revealing key insights into multi-sport learning and the critical role of fusion architecture in multi-modal sports analysis.




Figure 1: Visual examples of annotated data samples in RacketVision across the three sports. Each panel displays annotations for the ball’s position (red dot) and the racket’s bounding box (orange rectangle). The insets of each panel provide a schematic of the five keypoints defined for each specific racket type, which are used for the racket pose estimation task.

## Related Work

### Racket Sport Datasets

Existing datasets in racket sports have primarily focused on ball tracking. As summarized in Table 1, while foundational datasets such as TrackNet (Huang et al. 2019), TrackNetv2 (Huang et al. 2019), and WASB (Tarashima et al. 2023) enabled deep learning approaches for single-sport ball tracking, they presented opportunities for expansion in terms of scene diversity, sport variety, and annotation scope. Building upon these efforts, RacketVision provides a significant leap in scale and diversity, featuring substantially more games and frames across three distinct sports. Critically, it introduces the first large-scale annotations for racket pose (R) in addition to ball positions (B), enabling novel multi-modal analysis beyond simple ball tracking.

Dataset	Res	#S	#G	#F	#A	AT
GolfBall	720p	1	1	2k	2k	B
OpenTTGames	1080p	1	12	55k	55k	B
TrackNet	720p	1	10	19k	19k	B
TrackNetv2	720p	1	19	78k	78k	B
RacketVision	1080p	3	461	435k	88k	B,R

Table 1: Comparison of racket sports datasets. Res stands for resolution. #G, #F and #A stands for number of Games, Frames and Annotations. #S stands for number of sport types. AT stands for annotation types, where B is ball position, R is our first proposed racket pose annotation.

### Racket Sport Analysis Methods

Research on sport analysis methods has evolved from the basic task of 2D ball tracking to more sophisticated tasks centered on humans and games (Xia et al. 2024; Rao et al. 2025; Dong et al. 2024; Yuchen et al. 2024, 2025a). The development of robust 2D trackers has progressed from early CNN-based detectors (Reno et al. 2018), to specialized architectures for small objects (Jedrzejczak et al. 2019), and data-efficient semi-supervised learning (Vandeghen, Cioppa, and Van Droogenbroeck 2022). The importance of this task is further underscored by large-scale benchmarks like SoccerNet (Deliège et al. 2021; Cioppa et al. 2022). Recent

studies have explored 3D trajectory and spin reconstruction from monocular videos (Gossard, Ziegler, and Zell 2025; Kienzle et al. 2025), hit anticipation (Etaat et al. 2025), and stroke recognition using trajectory data alone (Kulkarni et al. 2023). Racket-centric studies rely on specialized hardware like high-speed or stereo cameras (Chen et al. 2013; Gao 2019) or complex proxies like human keypoints to handle occlusions (Zheng et al. 2023). However, these advanced methods have been constrained by the lack of large-scale, unified benchmarks. RacketVision addresses this gap, providing a public benchmark to train and evaluate general-purpose models (Jiang et al. 2023; Xu et al. 2022; Carion et al. 2020; Cao et al. 2017; Jocher and Qiu 2024) on these complex, interconnected analysis tasks, thereby lowering the barrier for future research.

### RacketVision Dataset

We introduce RacketVision, a large-scale video dataset designed to foster research in sports analytics across multiple racket sports: table tennis, tennis, and badminton. The benchmark provides a comprehensive resource for the interconnected tasks of ball tracking, racket pose estimation, and trajectory prediction. A detailed statistical breakdown for each sport is provided in Table 2.

### Data Collection and Annotation

The data collection and annotation pipeline, illustrated in Figure 2, was designed to ensure data quality, diversity, and annotation efficiency. The process begins with sourcing video from 942 top-level professional game broadcasts on YouTube, covering badminton, tennis, and table tennis to capture a wide variety of players and match dynamics.

In the first stage of the pipeline, a team of crowd-sourced annotators segments these raw videos into valid clips. A clip is defined as a continuous segment of 5-10 seconds where the ball is actively in play, which focuses the dataset on the most analytically relevant portions of the game. In the second stage, we employ a sparse annotation strategy to balance cost and quality. Instead of annotating every frame, 20% of the frames within each clip are evenly sampled for manual labeling by a different group of annotators. This approach maintains high temporal diversity while reducing redundancy. As illustrated in Figure 1, annotators labeled the

Sport	#Games	#Clips	#Frames	Length (s)	Ball Anno.	Racket Anno.
Table Tennis	50	780	170,027	3,878	19,495	6,648
Tennis	431	431	150,399	4,285	21,544	7,395
Badminton	461	461	114,753	4,592	23,003	10,578
Total	942	1672	435,179	12,755	64,042	24,621

Table 2: Statistical breakdown of the RacketVision dataset, detailing the number of games, clips, frames, total duration (in seconds), and the count of ball and racket annotations for each sport and in total.

ball’s position as a single point (red dot) with a visibility flag for each sampled frame. Rackets were annotated with both a bounding box (orange rectangle) and five specific keypoints designed to capture the pose of each racket type.

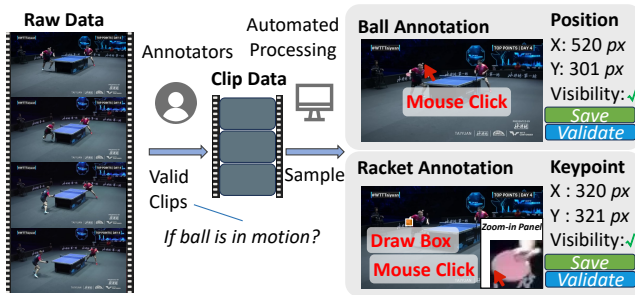


Figure 2: The two-stage annotation pipeline for RacketVision. First, crowd-sourced annotators segment valid clips from raw videos where the ball is in motion. Second, on sparsely sampled frames from these clips, another group of annotators labels the ball’s position as well as the racket’s bounding box and keypoints using a specialized interface.

## Dataset Structure

Each sample in RacketVision is a short video clip accompanied by a metadata file. The metadata includes the sport type and the indices of the annotated frames. To support baseline models that leverage background information, we also pre-process and provide a median frame for each clip. This serves as a stable background reference, which is particularly useful for distinguishing the small, fast-moving ball from the environment. All annotations, including ball positions, racket bounding boxes, and racket keypoints, are provided at the pixel level.

## Tasks

The RacketVision benchmark is structured around three interconnected tasks that form a comprehensive pipeline for sports analysis, progressing from low-level perception to high-level prediction: ball tracking, racket pose estimation, and ball trajectory prediction. Together, they serve a dual purpose: to drive innovation in sports analytics and to provide a framework for studying multi-modal, dynamic human-object interactions.

Figure 3 provides a detailed schematic of the relationship and workflow between these tasks. As illustrated, the process begins with training the two foundational perception models, the Ball Tracker and the Racket Pose Estimator, directly on the sparse, manually-labeled ground-truth frames provided in our dataset. Subsequently, these trained perception models are deployed on full video clips to generate dense, per-frame predictions, or “soft labels,” of ball and racket positions. These continuous sequences are then segmented into historical and future data segments, forming the rich training dataset for the final high-level task: the Ball Trajectory Predictor. This pipeline structure not only defines the dependencies between the tasks but also represents a practical workflow for building a complete sports analysis system. The following sections will provide a formal problem definition for each task.

## Ball Tracking

**Problem Formulation.** Given RGB frame  $I_t \in \mathbb{R}^{H \times W \times 3}$  at time  $t$  (single-frame setting) or a sequence of frames  $\{I_{t-N}, \dots, I_t\}$  with  $N = 5$  (multi-frame setting), predict the ball’s position  $(x_t, y_t) \in \mathbb{R}^2$  and visibility flag  $v_t \in \{0, 1\}$  in the target frame  $I_t$ .

**Settings.** The task has two settings: (1) *single-frame*, using only the target frame  $I_t$  to test static detection capabilities; and (2) *multi-frame*, using the target frame and five preceding frames  $\{I_{t-5}, \dots, I_t\}$  to leverage temporal context for improved robustness against occlusions and motion blur (Rozumnyi et al. 2021; Zhong et al. 2022, 2023).

## Racket Pose Estimation

**Problem Formulation.** Given RGB frame  $I_t \in \mathbb{R}^{H \times W \times 3}$ , predict the bounding box  $(x_{\min}, y_{\min}, x_{\max}, y_{\max}) \in \mathbb{R}^4$  and five keypoints  $\{(x_i, y_i)\}_{i=1}^5 \in \mathbb{R}^{10}$  for each racket in the frame.

**Settings.** The task uses a single-frame setting, predicting Bbox and keypoints from  $I_t$ . This setting focuses on static pose estimation, suitable for the dataset’s high-resolution frames and diverse racket orientations.

## Ball Trajectory Prediction Given History

**Problem Formulation.** Given a history of ball positions over  $N$  frames  $\{(x_{t-N+1}, y_{t-N+1}), \dots, (x_t, y_t)\} \in \mathbb{R}^{N \times 2}$  (ball-only setting) or ball positions plus racket poses  $\{(x_{t-N+1}, y_{t-N+1}, \{(x_i, y_i)\}_{i=1}^5), \dots, (x_t, y_t, \{(x_i, y_i)\}_{i=1}^5)\}$  (ball + racket setting), predict the ball’s trajectory over the next  $M$  frames  $\{(\hat{x}_{t+1}, \hat{y}_{t+1}), \dots, (\hat{x}_{t+M}, \hat{y}_{t+M})\} \in \mathbb{R}^{M \times 2}$ .

**Settings.** The task has two settings for input data modality: (1) *ball-only*, using the  $N$ -frame ball position history, which tests trajectory modeling based on ball dynamics; and (2) *ball + racket*, incorporating racket pose history (5 keypoints), which accounts for player interactions and improves prediction accuracy. We also use two settings for (1) *long trajectory*, that set  $N = 80$  and  $M = 20$ , and (2) *short trajectory*, that set  $N = 20$  and  $M = 5$ .

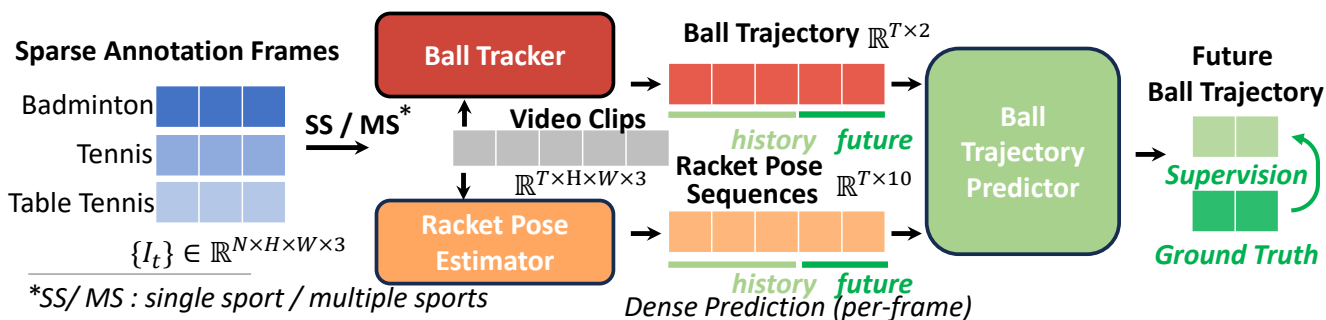


Figure 3: An overview of the task pipeline in RacketVision. Initially, the *Ball Tracker* and *Racket Pose Estimator* are trained using sparse ground-truth annotations. These models then process full video clips to generate dense trajectory data (soft labels), which serves as the training input for the final *Ball Trajectory Predictor*.

## Experiments and Baseline Solutions

We evaluate the RacketVision dataset on tasks defined in Sec.: ball tracking, racket pose estimation, and ball trajectory prediction given history. For each task, we define evaluation metrics, introduce baseline models, present experimental results, and analyze key observations. Figure 3 shows the relationships between the 3 tasks.

### Ball Tracking

**Evaluation Metrics.** We use four standard metrics to evaluate ball tracking performance:

- *Precision (Prec.)*: The ratio of correctly predicted ball positions (within a distance threshold) to the total number of predictions.
- *Recall (Rec.)*: The ratio of correctly predicted ball positions to the total number of ground-truth visible balls.
- *Mean Distance Error (MDE)*: The average Euclidean distance in pixels between predicted and ground-truth positions for visible balls, assuming a 1920x1080 resolution. A lower value is better.
- *mAP@50 (mAP)*: Mean Average Precision at IoU threshold of 0.5, evaluating the overall detection accuracy.

**Baseline Models.** We evaluate three representative baseline models:

- *RTMDet* (Lyu et al. 2022): A state-of-the-art real-time object detection model, adapted for ball detection.
- *YOLO11* (Jocher and Qiu 2024): A state-of-the-art vision model for real-time object detection.
- *WASB* (Tarashima et al. 2023): A strong baseline specifically designed for sports ball tracking, which internally incorporates background modeling.
- *TrackNetV3* (Chen and Wang 2023): A specialized heatmap-based network for tracking small, high-speed objects in sports, which can leverage temporal context from multiple frames.

**Experimental Setup.** Our experiments are designed to investigate three key axes of performance: the choice of model architecture, the benefit of multi-sport training, and the impact of techniques like background modeling (BM) and

multi-frame inputs (#F). Due to the extensive search space, we focused our multi-frame and multi-sport experiments primarily on the *TrackNetV3* architecture.

**Experimental Results.** Table 3 summarizes the performance of all evaluated models and settings. Figure 4 provides a visual example of the tracking performance of our best model.




Figure 4: The visualization of ball tracking result of MS-TrackNetV3 (with BM, #F=4) on table tennis. The red dots are sparse ground-truth ball position annotations, while the green dots are model predictions. The yellow line shows the combined path of ground-truth and predictions, illustrating the complete ball trajectory within the clip.

**Observations and Analysis.** Table 3 summarizes the performance of all evaluated models and settings. Figure 4 provides a visual example of the tracking performance of our best model.

- **Multi-Sport training generally enhances model generalization, especially on detection-oriented metrics.** By comparing the best multi-sport model (MS-TrackNetV3, #F=4, bold results) with the best single-sport model (TrackNetV3, #F=4, underlined results), we observe a clear trend of improved generalization. For example, the MS model boosts mAP by a significant 19.2% in Tennis (81.9 vs. 68.7) and 14.6% in Badminton (83.1 vs. 72.5). However, this broad generalization sometimes comes at the cost of hyper-specialized precision on a single sport.

Model	BM	#F	Table Tennis				Tennis				Badminton			
			Prec.	Rec.	MDE	mAP	Prec.	Rec.	MDE	mAP	Prec.	Rec.	MDE	mAP
RTMDet	×	1	0.844	0.724	35.8	68.1	0.862	0.477	22.5	46.7	0.801	0.721	38.2	64.2
YOLO11	×	1	0.877	0.595	30.6	67.2	0.881	0.416	17.5	58.5	0.858	0.548	29.3	72.1
WASB	✓	1	0.843	<u>0.733</u>	12.9	51.8	0.937	0.803	3.62	66.0	0.907	0.803	4.03	58.7
TrackNetV3	×	1	0.793	0.639	22.5	51.0	0.904	0.725	7.95	58.9	0.876	0.706	7.39	57.8
TrackNetV3	✓	1	0.869	0.724	10.3	67.7	0.941	<u>0.830</u>	3.07	65.6	<u>0.924</u>	0.797	3.34	67.7
TrackNetV3	✓	4	<u>0.898</u>	0.725	<u>6.63</u>	<u>68.3</u>	<u>0.962</u>	0.797	<u>1.66</u>	<u>68.7</u>	0.922	<u>0.865</u>	<u>2.07</u>	<u>72.5</u>
MS-RTMDet	×	1	0.852	0.733	29.6	69.3	0.885	0.503	19.3	48.3	0.823	0.736	31.0	64.5
MS-WASB	✓	1	0.867	0.742	8.53	56.2	0.943	0.802	3.73	65.1	0.924	0.807	2.13	63.3
MS-TrackNetV3	×	1	0.804	0.672	19.6	52.4	0.912	0.769	6.58	50.2	0.830	0.727	6.95	59.4
MS-TrackNetV3	✓	1	0.890	0.731	7.57	59.1	<b>0.962</b>	0.820	<b>1.70</b>	66.6	0.906	0.824	3.68	67.5
MS-TrackNetV3	✓	4	<b>0.924</b>	<b>0.762</b>	<b>3.41</b>	<b>71.1</b>	0.945	<b>0.880</b>	1.96	<b>81.9</b>	<b>0.915</b>	<b>0.865</b>	<b>1.54</b>	<b>83.1</b>

Table 3: Ball Tracking Experimental Results on RacketVision. BM represents whether add background median into input. Models starts with MS- are trained on all three sports, while others are trained on one sport. The bold results are the best results of MS-models, the underline results are the best results of models trained on single sport.

For instance, the single-sport model retains a slight edge in Tennis precision (0.962 vs. 0.945) and MDE (1.66 vs. 1.96). This suggests that training on diverse data forces the model to learn more robust features, enhancing its ability to find the ball under varied conditions (higher Recall and mAP), occasionally at the expense of pinpoint localization accuracy on a specific domain.

- **Background modeling is a highly effective technique for reducing localization error.** Incorporating a median frame for background subtraction (BM=✓) provides a powerful prior for distinguishing the small, fast-moving ball from a static or complex background. This is most evident in the Mean Distance Error (MDE). For example, comparing TrackNetV3 (#F=1) with and without BM, background modeling reduces MDE by a remarkable 54.0% for Table Tennis, 61.4% for Tennis, and 54.8% for Badminton. This consistently large improvement underscores the value of providing the model with an explicit representation of the static scene to mitigate false positives and improve localization.
- **Leveraging temporal context with multi-frame inputs boosts overall detection accuracy but reveals performance trade-offs.** Using multiple frames (#F=4) allows the model to leverage motion cues, which is particularly effective for improving recall and mAP in complex scenarios. The best overall results in our benchmark are achieved by MS-TrackNetV3 with 4 frames. However, the claim that multi-frame input is universally superior requires nuance. For instance, while using 4 frames with MS-TrackNetV3 in Tennis boosts Recall (0.880 vs. 0.820), the single-frame version achieves a slightly better MDE (1.70 vs. 1.96). This indicates that while temporal context is crucial for detecting the ball during challenging rallies (improving mAP), it can occasionally introduce minor jitter or motion blur that slightly affects the precision of the final predicted coordinate compared to a clean single frame.

## Racket Pose Estimation

**Evaluation Metrics.** We use two metrics to evaluate racket pose estimation:

- *Percentage of Correct Keypoints@0.2 (PCK):* Percentage of predicted keypoints that fall within a normalized distance threshold of 0.2 times the bounding box size from their corresponding ground-truth positions.
- *Mean Per-Joint Position Error (MPJPE):* Average Euclidean distance in pixels between predicted and ground-truth keypoint positions across all visible keypoints.
- *mean Object Keypoint Similarity (mOKS):* Mean Object Keypoint Similarity score that measures the similarity between predicted and ground-truth keypoint configurations.
- *Normalized Mean Error (NME):* Normalized mean error calculated by dividing the average keypoint position error by the distance between left and right keypoints, providing scale-invariant evaluation of pose estimation accuracy.
- *mAP@50 (mAP): Mean Average Precision at IoU=0.5, evaluating detection accuracy:* Average precision computed at Intersection over Union (IoU) threshold of 0.5, measuring detection accuracy for moderately overlapping predictions with ground-truth bounding boxes.

**Baseline Models.** We evaluate a top-down baseline in single-sport and multi-sport training settings:

- *RTMPose* (Jiang et al. 2023): A real-time top-down pose estimation model, optimized for keypoint detection in single-frame inputs. We adopt RTMDet (Lyu et al. 2022) as detector to generate bounding box.

**Experimental Results.** Table 4 summarizes the performance of baselines under the single-frame setting. Table 5 compares the performance of the best model on 5 keypoints. Figure 5 provides a visual example of the racket pose estimation result of our best model.

Train	Sport	Pose Estimation				Det
		PCK	MPJPE	mOKS	NME	mAP
SS	Table Tennis	75.6	10.6	0.453	0.279	72.4
	Tennis	83.7	5.87	0.574	0.245	73.4
	Badminton	82.1	5.45	0.601	0.259	69.8
MS	Table Tennis	81.8	9.71	0.498	0.254	78.4
	Tennis	<b>89.6</b>	5.34	0.630	<b>0.223</b>	<b>79.4</b>
	Badminton	88.5	<b>5.00</b>	<b>0.668</b>	0.235	75.5

Table 4: Main Performance Metrics Comparison. MS representing model trained on multiple sports, while SS representing single sport.

Sport	Individual Keypoints (%)				
	Top	Bottom	Handle	Left	Right
Table Tennis	97.6	97.3	<b>97.9</b>	64.8	64.8
Tennis	98.6	98.9	92.6	<b>79.7</b>	<b>80.1</b>
Badminton	<b>99.4</b>	<b>99.7</b>	97.3	74.6	75.5

Table 5: The comparison of PCK@0.2 performance of MS RTMPose model on different racket pose keypoints.




Figure 5: Visualization result of racket pose estimation of MS RTMPose model on tennis clip.

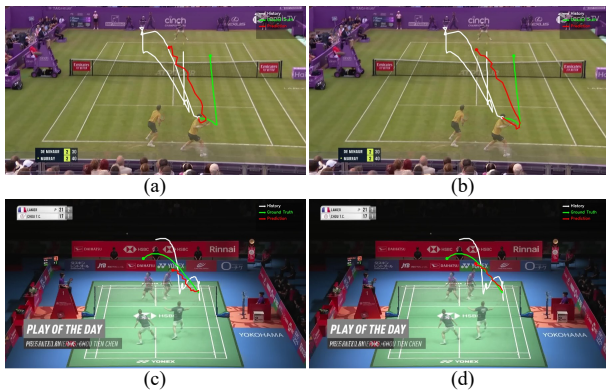


Figure 6: Qualitative comparison of long trajectory prediction. We compare the baseline LSTM Ball-Only model in (a), (c) with our proposed Cross-Attention LSTM Ball+Racket model in (b), (d).

- *Multi-sport training outperforms single-sport.* The MS model achieves superior performance with PCK@0.2 improvements of 6.17%, 6.36%, and 5.97% for table tennis, badminton, and tennis respectively. Tennis reaches the highest overall PCK@0.2 of 89.69% under multi-sport training, while badminton excels in pose quality metrics with the lowest MPJPE (5.00px) and highest mOKS (0.668), demonstrating the effectiveness of cross-sport knowledge transfer.
- *Side keypoint detection poses significant challenges for accurate racket pose estimation.* While structural keypoints (top, bottom, handle) achieve excellent accuracy above 92%, side keypoints (left, right) exhibit substantially lower performance ranging from 64.85% to 80.11%. This disparity stems from the inherent difficulty of detecting side edges, which are often occluded by hand grip, subject to motion blur during rapid movements, and highly sensitive to viewing angles.

## Ball Trajectory Prediction Given History

**Evaluation Metrics.** We use two metrics for trajectory prediction:

- *Average Displacement Error (ADE):* The average Euclidean distance between the predicted and ground-truth ball positions over all  $M$  future frames. The error is measured in pixels, assuming a resolution of 1920x1080.
- *Final Displacement Error (FDE):* The Euclidean distance between the predicted and ground-truth ball positions at the final frame ( $t + M$ ).

**Baseline Models.** We evaluate two backbone architectures, LSTM and Transformer, under three different input and fusion settings. This allows us to analyze not only the performance of the backbones but also the effectiveness of different multi-modal fusion strategies.

### • Backbones:

- *LSTM* (Hochreiter and Schmidhuber 1997): A 2-layer recurrent neural network that models temporal dependencies sequentially through a stateful, recurrent mechanism.
- *Transformer* (Vaswani et al. 2017): A 4-layer encoder-only model that captures global dependencies in the sequence in parallel via its self-attention mechanism.

### • Input and Fusion Methods:

- *Ball-Only:* A strong unimodal baseline where the model only receives the historical ball coordinates as input.
- *Concat Fusion:* A naive multi-modal baseline. The embeddings of the ball coordinates and racket poses are concatenated along the feature dimension before being fed into the backbone. This method treats all features equally at every time step.
- *Cross-Attention Fusion:* The ball trajectory sequence acts as the *Query* and the racket pose sequence acts as the *Key* and *Value*. This allows the model to dynamically weigh and extract the most relevant racket

Setting	Model	Input	Fuse Method	Table Tennis		Tennis		Badminton	
				ADE	FDE	ADE	FDE	ADE	FDE
<i>Short Trajectory Prediction (History=20, Future=5)</i>									
Short	LSTM	Ball	-	41.9	64.0	23.8	37.6	37.5	60.7
Short	LSTM	Ball+Racket	Concat	58.1	86.6	29.3	45.3	45.7	70.7
Short	LSTM	Ball+Racket	CrossAttn	<b>38.3</b>	<b>60.4</b>	<b>22.8</b>	<b>35.7</b>	<b>37.0</b>	<b>59.3</b>
Short	Transformer	Ball	-	46.9	67.8	35.7	47.3	43.9	64.8
Short	Transformer	Ball+Racket	Concat	60.2	89.1	39.5	51.7	50.4	79.3
Short	Transformer	Ball+Racket	CrossAttn	43.6	65.7	27.8	45.8	40.1	63.7
<i>Long Trajectory Prediction (History=80, Future=20)</i>									
Long	LSTM	Ball	-	113.9	184.3	62.5	108.7	118.7	194.7
Long	LSTM	Ball+Racket	Concat	139.9	198.9	76.8	125.0	134.5	203.3
Long	LSTM	Ball+Racket	CrossAttn	<b>101.3</b>	<b>161.3</b>	<b>55.5</b>	<b>94.7</b>	<b>114.6</b>	<b>187.6</b>
Long	Transformer	Ball	-	145.3	207.5	89.9	144.8	142.7	228.3
Long	Transformer	Ball+Racket	Concat	152.8	218.8	107.9	177.8	146.8	239.7
Long	Transformer	Ball+Racket	CrossAttn	127.3	195.6	74.8	118.4	122.5	200.2

Table 6: Trajectory Prediction Experimental Results on RacketVision. The table shows the performance of different models and fusion methods under Short (History=20, Future=5) and Long (History=80, Future=20) prediction settings. The best result for each metric within each setting is highlighted in **bold**.

pose information for each time step of the ball’s trajectory, effectively filtering noise and focusing on critical events like impacts.

**Experimental Results.** Table 6 summarizes the performance of baselines under the short and long trajectory prediction setting. Figure 6 provides a visual comparison of the trajectory prediction results between best models with ball-only input and ball+racket input.

- **Naive Fusion Degrades Performance.** As shown in Table 6, simply concatenating racket pose features consistently leads to worse performance than the ball-only baseline for both LSTM and Transformer backbones. This is likely because a large portion of trajectory samples in our dataset capture the ball in mid-flight, where racket information is absent or irrelevant. The Concat method indiscriminately fuses this noisy or uninformative data, which hinders the model’s ability to learn the primary trajectory dynamics.
- **Cross-Attention Excels at Predicting Critical Events.** The LSTM model equipped with Cross-Attention fusion is the best-performing model overall, demonstrating that racket information is highly valuable when integrated intelligently. The qualitative results in Figure 6 reveal precisely why this method is effective. For the Tennis sample ((a) vs. (b)), the Cross-Attention model leverages the racket’s position to more accurately predict the trajectory’s turning point. Similarly, for the Badminton sample ((c) vs. (d)), the model correctly infers the post-hit direction from the racket’s pose. This shows that the Cross-Attention mechanism successfully learns to identify and heavily weigh racket features during critical “event” frames (i.e., hits), which are decisive for the subsequent trajectory.

- **The Nature of Trajectory Data Explains Overall Gains.** While Cross-Attention provides a clear advantage during player-ball interactions, the overall statistical improvement in ADE/FDE over the strong ball-only baseline is noticeable but not dramatic. In Short Badminton, ADE improves from 37.5 to 37.0. This can be attributed to the dataset’s composition: many samples, especially in the short-trajectory setting, consist entirely of the ball in flight, where no informative racket interaction occurs. In these common cases, the Cross-Attention model correctly learns to ignore the racket modality, effectively behaving like the ball-only model.

## Conclusion

In this work, we introduced RacketVision, a large-scale, multi-sport benchmark designed to advance sports analytics. By providing the first large-scale dataset with detailed annotations for both ball position and racket pose, we formulated three interconnected tasks—ball tracking, racket pose estimation, and ball trajectory prediction—to address key challenges in perception and motion forecasting.

Through extensive evaluation, we not only established strong performance benchmarks but also uncovered a critical insight into multi-modal fusion for trajectory prediction. We demonstrated that naively incorporating racket pose data via simple concatenation was detrimental to performance. However, a Cross-Attention architecture successfully unlocked the value of this contextual information, reversing the performance degradation and ultimately surpassing strong unimodal baselines. This key finding definitively demonstrates the dual importance of both the novel racket pose data and the advanced fusion architecture required to leverage it.

## Acknowledgments

The work is supported by Shanghai Artificial Intelligence Laboratory.

## References

- Cao, Z.; Simon, T.; Wei, S.-E.; and Sheikh, Y. 2017. Real-time multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 7291–7299.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *European conference on computer vision (ECCV)*, 213–229. Springer.
- Chen, G.; Xu, D.; Fang, Z.; Jiang, Z.; and Tan, M. 2013. Visual Measurement of the Racket Trajectory in Spinning Ball Striking for Table Tennis Player. *IEEE Transactions on Instrumentation and Measurement*, 62(11): 2901–2911.
- Chen, Y.; Zhan, Y.; Zhong, Z.; Wang, W.; Sun, X.; Qiao, Y.; and Zheng, Y. 2024. Within the dynamic context: Inertia-aware 3d human modeling with pose sequence. In *European Conference on Computer Vision*, 491–508. Springer.
- Chen, Y.-J.; and Wang, Y.-S. 2023. Tracknetv3: Enhancing shuttlecock tracking with augmentations and trajectory rectification. In *Proceedings of the 5th ACM International Conference on Multimedia in Asia*, 1–7.
- Cioppa, A.; Deliege, A.; Giancola, S.; Ghanem, B.; Van Droogenbroeck, M.; Gade, R.; and Moeslund, T. B. 2022. SoccerNet-Tracking: Multiple Object Tracking Dataset and Benchmark in Soccer Videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 3491–3501.
- D’Ambrosio, D. B.; Abeyruwan, S. W.; Graesser, L.; Iscen, A.; Amor, H. B.; Bewley, A.; Reed, B.; Reymann, K.; Takayama, L.; Tassa, Y.; et al. 2024. Achieving human level competitive robot table tennis. In *7th Robot Learning Workshop: Towards Robots with Human-Level Abilities*.
- Deliege, A.; Cioppa, A.; Giancola, S.; Seikavandi, M. J.; Dueholm, J. V.; Nasrollahi, K.; Ghanem, B.; Moeslund, T. B.; and Droogenbroeck, M. V. 2021. SoccerNet-v2 : A Dataset and Benchmarks for Holistic Understanding of Broadcast Soccer Videos. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- Dong, L.; Wang, W.; Qiao, Y.; and Sun, X. 2024. Lucidaction: A hierarchical and multi-model dataset for comprehensive action quality assessment. *Advances in Neural Information Processing Systems*, 37: 96468–96482.
- Etaat, D.; Kalaria, D.; Rahmanian, N.; and Sastry, S. 2025. LATTE-MV: Learning to Anticipate Table Tennis Hits from Monocular Videos. ArXiv:2503.20936 [cs].
- Gao, Y. 2019. Real-time 6D Racket Pose Estimation and Classification for Table Tennis Robots. *International Journal of Robotic Computing*, 1(1): 23–39.
- Gossard, T.; Ziegler, A.; and Zell, A. 2025. TT3D: Table Tennis 3D Reconstruction.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation*, 9(8): 1735–1780.
- Huang, Y.-C.; Liao, I.-N.; Chen, C.-H.; İk, T.-U.; and Peng, W.-C. 2019. Tracknet: A deep learning network for tracking high-speed and tiny objects in sports applications. In *2019 16th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, 1–8. IEEE.
- Jedrzejczak, K.; Twardowski, M.; Gryka, A.; and Tabor, J. 2019. DeepBall: Deep Neural-Network Ball Detector. *arXiv preprint arXiv:1902.07304*.
- Jiang, T.; Lu, P.; Zhang, L.; Ma, N.; Han, R.; Lyu, C.; Li, Y.; and Chen, K. 2023. RtmPose: Real-time multi-person pose estimation based on mmpose. *arXiv preprint arXiv:2303.07399*.
- Jocher, G.; and Qiu, J. 2024. Ultralytics YOLO11.
- Kienzle, D.; Schön, R.; Lienhart, R.; and Satoh, S. 2025. Towards Ball Spin and Trajectory Analysis in Table Tennis Broadcast Videos via Physically Grounded Synthetic-to-Real Transfer.
- Kulkarni, K. M.; Jamadagni, R. S.; Paul, J. A.; and Shenoy, S. 2023. Table Tennis Stroke Detection and Recognition Using Ball Trajectory Data. ArXiv:2302.09657.
- Lyu, C.; Zhang, W.; Huang, H.; Zhou, Y.; Wang, Y.; Liu, Y.; Zhang, S.; and Chen, K. 2022. RtmDET: An empirical study of designing real-time object detectors. *arXiv preprint arXiv:2212.07784*.
- Ma, Y.; Cramariuc, A.; Farshidian, F.; and Hutter, M. 2025. Learning coordinated badminton skills for legged manipulators. *Science Robotics*, 10(102): eadu3922.
- Rao, J.; Wu, H.; Jiang, H.; Zhang, Y.; Wang, Y.; and Xie, W. 2025. Towards Universal Soccer Video Understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Reno, V.; Mosca, N.; Marani, R.; Nitti, M.; D’Orazio, T.; and Stella, E. 2018. Convolutional Neural Networks Based Ball Detection in Tennis Games. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2338–2344.
- Rozumnyi, D.; Oswald, M. R.; Ferrari, V.; and Pollefeys, M. 2021. Shape from blur: Recovering textured 3d shape and motion of fast moving objects. *Advances in Neural Information Processing Systems*, 34: 29972–29983.
- Sun, N.-E.; Lin, Y.-C.; Chuang, S.-P.; Hsu, T.-H.; Yu, D.-R.; Chung, H.-Y.; and İk, T.-U. 2020. Tracknetv2: Efficient shuttlecock tracking network. In *2020 International Conference on Pervasive Artificial Intelligence (ICPAI)*, 86–91. IEEE.
- Tarashima, S.; Haq, M. A.; Wang, Y.; and Tagawa, N. 2023. Widely Applicable Strong Baseline for Sports Ball Detection and Tracking. In *BMVC*.
- Vandeghen, R.; Cioppa, A.; and Van Droogenbroeck, M. 2022. Semi-Supervised Training To Improve Player and Ball Detection in Soccer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 3481–3490.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.

Wang, Z.; Veličković, P.; Hennes, D.; Tomašev, N.; Prince, L.; Kaisers, M.; Bachrach, Y.; Elie, R.; Wenliang, L. K.; Piccinini, F.; et al. 2024. TacticAI: an AI assistant for football tactics. *Nature communications*, 15(1): 1906.

Xia, H.; Yang, Z.; Zou, J.; Tracy, R.; Wang, Y.; Lu, C.; Lai, C.; He, Y.; Shao, X.; Xie, Z.; et al. 2024. Sportu: A comprehensive sports understanding benchmark for multimodal large language models. *arXiv preprint arXiv:2410.08474*.

Xu, W.; Zhan, Y.; Zhong, Z.; and Sun, X. 2025. Sequential Gaussian Avatars with Hierarchical Motion Context. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13592–13603.

Xu, Y.; Zhang, J.; Zhang, Q.; and Tao, D. 2022. Vitpose: Simple vision transformer baselines for human pose estimation. *Advances in neural information processing systems*, 35: 38571–38584.

Yuchen, Y.; Linfeng, D.; Wei, W.; Zhihang, Z.; and Xiao, S. 2025a. Learnable SMPLify: A Neural Solution for Optimization-Free Human Pose Inverse Kinematics. *arXiv:2508.13562*.

Yuchen, Y.; Wei, W.; Yifei, L.; Linfeng, D.; Hao, W.; Mingxin, Z.; Zhihang, Z.; and Xiao, S. 2025b. SGA-INTERACT: A 3D Skeleton-based Benchmark for Group Activity Understanding in Modern Basketball Tactic. *arXiv:2503.06522*.

Yuchen, Y.; Xuanyi, L.; Xing, G.; Zhihang, Z.; and Xiao, S. 2024. X as Supervision: Contending with Depth Ambiguity in Unsupervised Monocular 3D Pose Estimation. *arXiv:2411.13026*.

Zhan, Y.; Xu, W.; Zhu, Q.; Niu, M.; Ma, M.; Liu, Y.; Zhong, Z.; Sun, X.; and Zheng, Y. 2025a. R3-Avatar: Record and Retrieve Temporal Codebook for Reconstructing Photorealistic Human Avatars. *arXiv preprint arXiv:2503.12751*.

Zhan, Y.; Zhu, Q.; Niu, M.; Ma, M.; Zhao, J.; Zhong, Z.; Sun, X.; Qiao, Y.; and Zheng, Y. 2025b. Towards Explicit Exoskeleton for the Reconstruction of Complicated 3D Human Avatars. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 14259–14269.

Zheng, Y.; Zhou, W.; Zou, T.; and Zhang, H. 2023. A Method for Table Tennis Bat Trajectories Reconstruction with the Fusion of Human Keypoint Information. In *2023 8th IEEE International Conference on Network Intelligence and Digital Content (IC-NIDC)*, 71–75. Beijing, China: IEEE. ISBN 9798350317923.

Zhong, Z.; Cao, M.; Ji, X.; Zheng, Y.; and Sato, I. 2023. Blur interpolation transformer for real-world motion from blur. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5713–5723.

Zhong, Z.; Sun, X.; Wu, Z.; Zheng, Y.; Lin, S.; and Sato, I. 2022. Animation from blur: Multi-modal blur decomposition with motion guidance. In *European Conference on Computer Vision*, 599–615. Springer.