

EVENT-BASED BATTING IMPACT ESTIMATION

Ryotaro Ishida[†], Wataru Ikeda[†], Ryosei Hara[†]
Akemi Kobayashi[‡], Toshitaka Kimura[‡], Mariko Isogawa[†]

[†]Keio University, [‡]NTT Communication Science Laboratories

ABSTRACT

Estimating the precise timing of batting impact is crucial for understanding the rapid sensorimotor control. However, this task is challenging for RGB cameras due to insufficient temporal resolution and motion blur. Similarly, Inertial Measurement Units (IMUs) are impractical for actual matches due to sensor intrusiveness and their limited temporal precision. To overcome these limitations, we propose a novel framework leveraging event-based cameras, which offer microsecond resolution and high dynamic range, to estimate impact timing based on the weighted centroid distance between the detected ball and bat. To address the domain gap between event frames and RGB images that degrades segmentation accuracy, we generate high-density event frames. We then introduce a mask refinement network that leverages these frames and bidirectional mask information, optimized using a novel loss function. Experiments on real-world datasets demonstrate that our method achieves superior accuracy under challenging conditions, including low-light environments and severe occlusions, outperforming baselines by reducing the Mean Absolute Error by approximately 63%.

Index Terms— event-based camera, sports analytics, impact estimation, mask optimization

1. INTRODUCTION

In recent years, computer vision has rapidly advanced sports analytics, providing quantitative insights that were previously inaccessible [1]. In sports motion analysis, precise estimation of impact timing is critical not only for decision making [2] but also for evaluating player performance in fast-paced sports [3, 4, 5]. Particularly in baseball, accurate temporal localization enables a detailed breakdown of the kinematic chain during the swing phase. However, the high-speed nature of the game makes this challenging: a ball traveling at 100 km/h covers ~ 2.8 cm per frame even at 1,000 fps. This spatial deviation is too large to pinpoint the exact moment of contact. Consequently, analysis at even higher temporal resolutions is required to capture sub-millisecond contact durations accurately.

Despite its importance, existing sensing modalities have limitations. Standard RGB cameras (typically on the order of 10^1 – 10^2 fps) suffer from insufficient temporal resolution [3, 4].




Fig. 1. Comparison of sensing modalities.

Fast-moving impacts often occur between frames, and the resulting motion blur obscures the precise boundaries of the ball and bat, making precise timing estimation infeasible from intensity image frames alone. While high-speed cameras can mitigate this, they incur prohibitive hardware costs and generate massive data volumes. Furthermore, unconstrained environments, such as variable lighting in night games, introduce severe challenges for these intensity-based approaches, complicating robust tracking.

Alternative modalities also face fundamental limitations. Acoustic-based methods [6, 5] are susceptible to environmental noise and synchronization latency. Wearable sensors, such as IMUs, have been employed to detect impact timing via acceleration spikes [7]. However, these sensors measure the vibrational response triggered by the collision, rather than directly observing the precise moment of contact initiation. Furthermore, the operational burden of attaching devices to the body or equipment of the athlete makes them intrusive. Consequently, there is a strong demand for a robust, non-intrusive vision system capable of sub-millisecond impact estimation under challenging illumination and occlusion.

To overcome these limitations, we leverage event-based cameras [8]. As bio-inspired sensors that asynchronously output pixel-level brightness changes, they offer microsecond-level temporal resolution, high dynamic range, and inherent resistance to motion blur. As shown in Figure 1, this enables capturing sharp bat and ball motion even in high-speed and low-light scenarios. However, utilizing event data is challenging due to its spatial sparsity and unique noise. Although recent foundation models trained on massive RGB datasets excel in video domains, they are inherently designed for dense

RGB textures. This mismatch results in performance degradation when applying them directly to event frames. Existing event-based detection and segmentation methods [9, 10] face critical constraints. They typically rely on events accumulated over long time windows (e.g., 10–60 ms). Consequently, these methods fail to produce reliable masks during the sparse, sub-millisecond impact interval. Furthermore, approaches like EventSAM [10] lack explicit temporal modeling, leading to unstable mask predictions that deviate significantly from previous frames.

To address this, we propose the first event-based impact estimation framework capable of operating at 10,000 fps. Our approach follows a coarse-to-fine strategy that effectively adapts robust RGB segmentation models to the high-speed event domain. First, to bridge the domain gap, we specially construct dense event frames even at this high temporal resolution. Since independent frame-wise predictions lack temporal consistency, we employ a bidirectional tracking mechanism to propagate stable masks across time. Subsequently, we introduce a mask refinement network. This network uses the coarse mask as a prior and leverages high-frequency event edges to sharpen object boundaries at the pixel level, maintaining precision even under occlusions. Finally, to estimate the impact timing, we identify the frame where the distance between the weighted centroids of the ball and bat is minimized. We validate our method on a real-world dataset including challenging low-light and occluded scenarios.

The main contributions of this paper are summarized as follows: **1)** We are the first to propose the task of estimating the impact timing between a fast-moving ball and a returning bat at high temporal resolution using event-based sensing. **2)** To this end, we propose a framework that estimates impact timing based on the weighted centroids of the ball and bat. We also introduce a coarse segmentation module that temporally estimates the mask regions of the ball and bat in a bidirectional manner, as well as a mask refinement network that further refines these mask regions for more accurate estimation. **3)** As this is a novel task, we also construct the first event-based dataset dedicated to high-speed baseball analysis and demonstrate the robustness of our method in challenging real-world environments, including low-light conditions and severe occlusions.

2. RELATED WORK

2.1. Impact Timing Estimation for Sports

Precise impact timing estimation is crucial for tasks ranging from video-to-audio generation [6] to sports performance analysis [3, 4, 5]. While RGB-based “Action Spotting” methods [3, 4] have advanced, they remain fundamentally constrained by standard frame rates. Specifically, the exact moment of impact is either entirely missed between frames or obscured by severe motion blur. Consequently, existing approaches must often rely on auxiliary modalities like IMUs to estimate impact timing [7].

In contrast, event-based cameras offer microsecond-level temporal resolution, showing great success in high-speed motion analysis such as video interpolation [11, 12] and robust object detection [9]. Most recently, Kase *et al.* [13] pioneered the use of event cameras to locate impact positions in tennis. However, their timing estimation relies on simple heuristics based on event counts and polarity changes. Such raw signal metrics are highly susceptible to environmental noise and occlusion. To overcome these limitations, this paper leverages a segmentation-based approach coupled with bidirectional tracking. By enforcing temporal consistency from both directions, our method ensures robust detection even under such challenging conditions.

2.2. Object Boundary Extraction for Contact Detection

Accurate impact timing estimation requires precise boundary localization, as slight mask deviations cause significant errors during trajectory intersection. Thus, robust boundary extraction is a prerequisite for this task. Foundation models for segmentation, such as SAM 3 [14], have demonstrated remarkable zero-shot generalization capabilities for unseen objects and complex scenes. Various refinement approaches [15, 16] further aim to enhance mask boundary precision. However, these methods predominantly rely on standard camera intensity frames. In high-speed impacts, motion blur obscures the object edges, making contact detection difficult with RGB data alone.

Event-driven adaptations like EventSAM [10] incorporate high temporal resolution into the segmentation pipeline. However, since these methods rely on accumulating events to form structural features, they face a critical trade-off in high-speed impact scenarios. Specifically, extending the integration window induces motion blur, whereas narrowing it results in spatially sparse inputs that cause segmentation breakdown. Furthermore, these methods typically lack mechanisms for temporal consistency, making them unstable for tracking fast-moving objects. To resolve these limitations, we propose a framework that leverages the semantic robustness of large-scale RGB segmentation models. To bridge the domain gap, we convert raw event streams into dense event frames that maintain sufficient structural density even at high temporal resolutions (10,000 fps). This approach allows us to utilize the RGB model’s generalization capability for robust tracking, effectively decoupling feature density from the integration time constraints.

2.3. Vision Datasets for Sports

While event-based vision shows promise in sports, such as the table tennis dataset by Alberico *et al.* [17], existing benchmarks focus on trajectory tracking rather than microsecond-level impact mechanics. Conversely, traditional baseball datasets like MLB-YouTube [18] rely on RGB video, which suffers from temporal aliasing unsuitable for precise contact analysis. Therefore, we introduce the first public dataset combining event-based vision with high-speed baseball impacts,




Fig. 2. The overview of proposed method

capturing critical dynamics across diverse environments.

3. PROPOSED METHOD

We propose a coarse-to-fine impact timing estimation framework that integrates the high temporal resolution and high dynamic range of an event-based camera with SAM 3. The proposed method consists of four main stages: Event Representation, Coarse Segmentation, Mask Refinement, and Impact Estimation. The Event Representation stage converts the raw event stream into a dense event frame sequence $\mathcal{E}$. The Coarse Segmentation stage then employs SAM 3 to obtain initial coarse masks $\mathbf{M}_{\text{fwd}}$ and $\mathbf{M}_{\text{bwd}}$. The Mask Refinement stage utilizes a U-Net-based network, taking the event frames E_k and coarse masks $\mathbf{M}_{\text{fwd}}$, $\mathbf{M}_{\text{bwd}}$ as inputs to generate refined probability maps $\mathbf{P}_{\text{ball}}$, $\mathbf{P}_{\text{bat}}$, and $\mathbf{P}_{\text{background}}$. The Impact Estimation stage determines the contact moment between the ball and the bat based on the spatiotemporal analysis of these refined probability maps.

3.1. Event Representation

Let $\mathcal{S} = \{e_i\}_{i=1}^N$ be the asynchronous event stream (often referred to as an event cloud), where each event $e_i = (x_i, y_i, t_i, p_i)$ encodes the pixel location, timestamp, and polarity [8]. To capture the impact timing with high precision, we convert $\mathcal{S}$ into a dense frame sequence $\mathcal{E} = \{E_k\}_{k=1}^K$, where each frame $E_k \in \mathbb{R}^{1 \times H \times W}$. We set the frame interval to $\Delta t = 0.1$ ms, corresponding to a temporal resolution of 10,000 fps. To ensure sufficient pixel density, we apply a sliding window accumulation. Specifically, the k -th frame E_k aggregates positive polarity events ($p_i = 1$) within the time window $[t_k - T_{\text{win}}, t_k]$ to prevent trailing streaks from degrading contour precision, where $t_k = k\Delta t$ and $T_{\text{win}} = 1.0$ ms is the accumulation duration. This approach yields dense event frames with enhanced ball and bat contours, providing sufficient structural information for segmentation while

maintaining fine-grained temporal resolution.

3.2. Coarse Segmentation

To obtain initial segmentation masks, we prompt SAM 3 with ball/bat point coordinates. Since the model expects standard video rates, we input the 10,000 Hz event frames $\mathcal{E}$ as a time-stretched 30 fps sequence, aligning apparent motion with the pre-trained priors of SAM 3. To mitigate occlusions and mask merging at impact, we employ a bi-directional inference strategy. Processing $\mathcal{E}$ in both temporal directions yields forward and backward probability maps $\mathbf{M}_{\text{fwd}}, \mathbf{M}_{\text{bwd}} \in \mathbb{R}^{2 \times H \times W}$ for the ball and bat.

3.3. Mask Refinement Network

Network Architecture. While the initial masks generated by SAM 3 are robust, they remain general-purpose and are not specifically optimized for the challenges arising from object collisions, such as the merging of masks or identity switching during the contact between the ball and the bat. To address this, we introduce a Refinement Network designed to sharpen object boundaries by leveraging the initial masks as spatial priors. We employ a U-Net [19] architecture with a ResNet-18 [20] backbone as the encoder.

The input to the network is defined as a tensor $\mathbf{X} \in \mathbb{R}^{5 \times H \times W}$, constructed by concatenating the event frame E_k with the forward and backward probability maps $\mathbf{M}_{\text{fwd}}$ and $\mathbf{M}_{\text{bwd}}$. This configuration enables the network to simultaneously exploit high-frequency edge cues from the event stream and bidirectional spatiotemporal priors provided by the initial tracking. Finally, the network outputs refined probability maps $\mathbf{P} \in \mathbb{R}^{3 \times H \times W}$, representing the pixel-wise likelihoods for the ball, bat, and background.

Loss Function. To ensure pixel-level classification accuracy while guaranteeing the geometric consistency of the object

masks, we employ a composite loss function $\mathcal{L}$, defined as:

$$\mathcal{L} = \lambda_{ce}\mathcal{L}_{ce} + \lambda_{dice}\mathcal{L}_{dice} + \lambda_{smooth}\mathcal{L}_{smooth} + \lambda_{circ}\mathcal{L}_{circ}, \quad (1)$$

where λ terms denote the balancing weights. $\mathcal{L}_{ce}$ represents the Weighted Cross-Entropy Loss and $\mathcal{L}_{dice}$ is the Dice Loss [21]. These standard terms maximize fundamental segmentation performance. To address the significant class imbalance between the background and the small ball object, we assign higher weights to the ball class in $\mathcal{L}_{ce}$.

To mitigate mask fragmentation caused by event noise, we introduce a Smoothness Loss ($\mathcal{L}_{smooth}$). This term acts as a regularizer to suppress isolated noise pixels by minimizing the anisotropic total variation of predicted probability maps $\mathbf{P}$. It is computed as the mean absolute difference between adjacent pixels:

$$\mathcal{L}_{smooth} = \frac{1}{N} \sum_{x,y} (|\mathbf{P}_{x+1,y} - \mathbf{P}_{x,y}| + |\mathbf{P}_{x,y+1} - \mathbf{P}_{x,y}|), \quad (2)$$

where N is the total number of pixels.

Furthermore, we impose a shape constraint specifically on the ball class using a Circularity Loss ($\mathcal{L}_{circ}$). Given that the projection of the ball is circular, we utilize the isoperimetric ratio to enforce compactness. We approximate the perimeter C using the gradients of the probability map $\mathbf{P}_{ball}$ and the area A as the sum of probabilities. The loss is defined as:

$$\mathcal{L}_{circ} = \frac{1}{N} \frac{C^2}{4\pi A + \epsilon}, \text{ where } C = \sum |\nabla \mathbf{P}_{ball}|, A = \sum \mathbf{P}_{ball}. \quad (3)$$

Here, ϵ is a small constant introduced for numerical stability. Minimizing this ratio encourages the predicted mask to approach a perfect circle, for which the ratio is 1, thereby preventing physically implausible deformations.

3.4. Impact Estimation

Centroid Localization and Distance Measurement. Based on the refined probability maps, we estimate object positions using weighted centroids. This approach avoids quantization errors inherent in binary mask centers. The centroid $\mathbf{c} = (c_x, c_y)$ for each object is computed as:

$$c_x = \frac{\sum_{i,j} x \cdot \mathbf{P}(i,j)}{\sum_{i,j} \mathbf{P}(i,j)}, \quad c_y = \frac{\sum_{i,j} y \cdot \mathbf{P}(i,j)}{\sum_{i,j} \mathbf{P}(i,j)} \quad (4)$$

The relative distance $d(t)$ is defined as the Euclidean distance between centroids $\mathbf{c}_{ball}(t)$ and $\mathbf{c}_{bat}(t)$. This strategy combines temporal consistency with sub-pixel precision, yielding a coherent trajectory sequence.

Timing Determination. The impact corresponds to the instant of maximum proximity between the two objects. Since the rotational mechanics of a bat swing induce non-linear distance changes, linear fitting approaches like RANSAC are often insufficient. Therefore, leveraging the 10,000 Hz resolution, we

determine the impact timing t_{impact} by directly minimizing the relative distance:

$$t_{impact} = \arg \min_t d(t) \quad (5)$$

This model-free approach effectively utilizes the high-frequency sampling to capture the precise collision frame without relying on geometric assumptions that may not hold during complex swing dynamics.

4. EXPERIMENT

4.1. Experimental Setup and Dataset

Data Collection. Since no public event-based dataset exists for high-speed baseball analysis, we constructed a novel dataset utilizing a DAVIS346 [22] event-based camera (346×260). A pitching machine was positioned at a standard distance of about 18.44 m from the home plate, pitching to three batters: a right-handed former professional player and two left-handed participants. The impact sequence of the batter was recorded from the side view. The camera position was slightly adjusted per session to optimize the field of view.

Dataset. Table 1 details the seven scenarios collected to evaluate robustness. The dataset covers varying conditions: illumination (low: ≈ 18 lux, mod: ≈ 435 lux, high: ≈ 1491 lux), occlusion (safety net), ball material (soft/hard), pitch type (fast/slider), and speed (100–145 km/h). In total, the dataset comprises 40 clips and 23,779 event frames.

Ground Truth Annotation. Initially, we attached an IMU to the bat knob, sampling at 1000 Hz, synchronized via an LED trigger. However, preliminary analysis revealed a variable latency of 0.4–7.1 ms in the IMU signal, rendering it insufficient for validating our sub-millisecond estimation. Therefore, we established the Ground Truth through manual inspection of the event stream. To minimize subjective bias and ensure reliability, three human annotators independently pinpointed the exact timing of contact. The final ground truth was defined as the average of these three annotations.

Evaluation Protocol. We adopted a two-stage strategy to evaluate the generalization capability of our method. First, for the low-light datasets (Scenarios 1–3), we employed a Leave-One-Scenario-Out (LOSO) cross-validation. In each fold, one scenario was used for testing, while the remaining two were combined for training. Second, to assess robustness against unseen pitch speeds and types, Scenarios 4–7 were used exclusively for testing. For these scenarios, the inference was performed using the model trained on the union of Scenarios 1–3. In all phases, the dataset set was split into training and validation subsets with an 8:2 ratio to monitor convergence.

4.2. Baseline

To rigorously evaluate the contributions of each component in our framework, we compare our method against a baseline. For a fair comparison, all baseline methods utilize the same stack of event frames $\mathcal{E}$ as input to estimate the impact timing.

Table 1. Detailed configuration and statistics of our dataset.

Scen.	Lux	Net	Ball	Type	Speed	# Clips	# Frames
1	Low	No	Soft	Fast	100	5	2,080
2	Low	Yes	Soft	Fast	100	4	1,440
3	Low	Yes	Hard	Fast	100	6	2,270
4	Mod	Yes	Soft	Fast	145	9	7,290
5	Mod	Yes	Soft	Slider	130	6	4,510
6	High	Yes	Soft	Fast	145	5	3,235
7	High	Yes	Soft	Slider	130	5	2,954
Total						40	23,779

Table 2. Quantitative evaluation results.

Method	MAE (ms) ↓	SR (< 1 σ) (%) ↑	SR (< 2 σ) (%) ↑
E2E-Spot	17.2	10.0	15.0
SAM 3*	3.09	30.0	52.5
YOLO 11*	2.06	32.5	55.0
Ours	0.767	52.5	77.5

* Methods adapted using geometric centers of masks/boxes.

E2E-Spot [3]: A representative RGB-based action spotting method originally designed for standard video frames. Since this model is not optimized for event data, we fine-tune it on the generated event frames $\mathcal{E}$ using the same training protocol as our proposed method to adapt it to the event domain.

SAM 3 [14]: As there is no standard event-based method for this specific task, we utilize SAM 3, a state-of-the-art foundation model for image segmentation. We apply SAM 3 directly to the event frames $\mathcal{E}$ to generate segmentation masks. To ensure a fair comparison in spatial localization, we compute the weighted centroids from the output probability maps of SAM 3, following the same mechanism as our approach.

YOLO 11 [23]: We also employ YOLO 11, a state-of-the-art object detection framework. Similar to E2E-Spot, we fine-tune the model on the event frames $\mathcal{E}$. Since YOLO outputs bounding boxes, we select the detection with the highest confidence score and utilize the geometric center of this bounding box as the predicted spotting location.

4.3. Evaluation Metrics

We evaluate the temporal precision using Mean Absolute Error (MAE) and Success Rate (SR). For SR, we define the success thresholds based on human annotation variability to account for the inherent ambiguity of the impact timing. Specifically, we calculate the standard deviation of the three annotators for each clip and average it across the dataset, resulting in a human baseline noise level of $\sigma = 0.513$ ms. We report SR using tolerance thresholds of 1σ and 2σ (1.026 ms).

5. RESULTS

5.1. Quantitative Evaluation

As shown in Table 2, our method outperforms all baselines, achieving the best MAE of 0.767 and SR of 52.5% (< 1σ). These results correspond to an approximately 63% reduction in MAE and a 20.0% increase in SR compared to the second-best method, YOLO 11. The ablation study in Table 3 validates

Table 3. Ablation study on the proposed components.

Refine				MAE (ms) ↓	SR (%) ↑	
(fwd)	(bwd)	+ $\mathcal{L}_{\text{circ}}$	+ $\mathcal{L}_{\text{smooth}}$		< 1σ	< 2σ
				3.09	30.0	52.5
✓				2.665	37.5	62.5
✓	✓			1.787	42.5	72.5
✓	✓	✓		1.696	47.5	72.5
✓	✓	✓	✓	0.767	52.5	77.5




Fig. 3. Qualitative comparison of estimated impact frames. Frames with **yellow borders** indicate failure cases where the error exceeds the 2σ threshold.

our design choices. The bidirectional refinement drastically drops the MAE from 2.665 to 1.787. Additionally, the geometric constraints ($\mathcal{L}_{\text{circ}}$ and $\mathcal{L}_{\text{smooth}}$) further refine the estimation, yielding the best performance across all metrics. (see supplementary material for per-scene MAE).

5.2. Qualitative Evaluation

Figure 3 visualizes the estimated impact frames to demonstrate the temporal precision. Consistent with the success rate threshold defined in Sec. 4.3, we mark predictions that deviate from the ground truth by more than 2σ with yellow borders. The visualization highlights that while baseline methods frequently produce outliers indicated by yellow borders, our proposed method consistently estimates the impact timing within the tolerance range, aligning closely with the ground truth.

6. CONCLUSION

In this paper, we proposed the first high-temporal-resolution framework for estimating baseball impact timing using event-based cameras. We constructed a novel dataset and introduced a pipeline that refines segmentation masks to calculate the weighted centroids of the ball and bat, enabling precise timing estimation. Experimental results demonstrated that our method significantly outperforms baseline approaches. Although our current approach assumes full visibility of the ball and bat and relies on a modestly sized dataset, we plan to address these limitations in future work by expanding the dataset and diversifying recording conditions.

Acknowledgement. This work was partially supported by JST Presto JPMJPR22C1, JSPS Grant-in-Aid for Challenging Research (Exploratory) 24K22296. Wataru Ikeda was supported by JST BOOST, Japan Grant Number JPMJBS2409.

7. REFERENCES

- [1] Paresh R. Kamble, Avinash G. Keskar, and Kishor M. Bhurchandi, “Ball tracking in sports: a survey,” *Artificial Intelligence Review*, vol. 52, no. 3, pp. 1655–1705, 2019.
- [2] A. Kobayashi, D. Nasu, Y. Morimoto, M. Kashino, and T. Kimura, “Detecting and sonifying temporal patterns of body segments when batting,” *MDPI*, vol. 2, no. 6, 2018.
- [3] James Hong, Haotian Zhang, Michaël Gharbi, Matthew Fisher, and Kayvon Fatahalian, “Spotting temporally precise, fine-grained events in video,” in *ECCV*, 2022.
- [4] Artur Xarles, Sergio Escalera, Thomas B. Moeslund, and Albert Clapés, “T-deed: Temporal-discriminability enhancer encoder-decoder for precise event spotting in sports videos,” in *CVPRW*, 2024.
- [5] Lucio Caprioli, Francesca Campoli, Saeid Edriss, Cristiana Frontuto, Amani Najlaoui, Elvira Padua, Cristian Romagnoli, Giuseppe Annino, and Vincenzo Bonaiuto, “Assessment of tennis timing using an acoustic detection system,” in *STAR*, 2024.
- [6] Andrew Owens, Phillip Isola, Josh McDermott, Antonio Torralba, Edward H. Adelson, and William T. Freeman, “Visually indicated sounds,” in *CVPR*, 2016.
- [7] Niroshan G. Punchihewa, Hideki Arakawa, Etsuo Chosa, and Go Yamako, “A hand-worn inertial measurement unit for detection of bat–ball impact during baseball hitting,” *Sensors*, vol. 21, no. 9, pp. 3002, 2021.
- [8] Guillermo Gallego, Tobi Delbruck, Garrick Orchard, Chiara Bartolozzi, Brian Taba, Andrea Censi, Stefan Leutenegger, Andrew J. Davison, Jorg Conradt, Kostas Daniilidis, and Davide Scaramuzza, “Event-based vision: A survey,” *IEEE TPAMI*, vol. 44, no. 1, pp. 154–180, 2022.
- [9] Mathias Gehrig and Davide Scaramuzza, “Recurrent vision transformers for object detection with event cameras,” in *CVPR*, 2023.
- [10] Zhiwen Chen, Zhiyu Zhu, Yifan Zhang, Junhui Hou, Guangming Shi, and Jinjian Wu, “Segment any event streams via weighted adaptation of pivotal tokens,” in *CVPR*, 2024.
- [11] Stepan Tulyakov, Daniel Gehrig, Stamatios Georgoulis, Julius Erbach, Mathias Gehrig, Yuanyou Li, and Davide Scaramuzza, “Time lens: Event-based video frame interpolation,” in *CVPR*, 2021.
- [12] Taewoo Kim, Yujeong Chae, Hyun-Kurl Jang, and Kuk-Jin Yoon, “Event-based video frame interpolation with cross-modal asymmetric bidirectional motion fields,” in *CVPR*, 2023.
- [13] Yuto Kase, Kai Ishibe, Ryoma Yasuda, Yudai Washida, and Sakiko Hashimoto, “Locating tennis ball impact on the racket in real time using an event camera,” in *Sports Analytics*, 2025.
- [14] Nicolas Carion, Laura Gustafson, Yuan-Ting Hu, Shoubhik Debnath, Ronghang Hu, Didac Suris, Chaitanya Ryali, Kalyan Vasudev Alwala, Haitham Khedr, Andrew Huang, Jie Lei, Tengyu Ma, Baishan Guo, Arpit Kalla, Markus Marks, Joseph Greer, Meng Wang, Peize Sun, Roman Rädle, Triantafyllos Afouras, Effrosyni Mavroudi, Katherine Xu, Tsung-Han Wu, Yu Zhou, Liliane Momeni, Rishi Hazra, Shuangrui Ding, Sagar Vaze, Francois Porcher, Feng Li, Siyuan Li, Aishwarya Kamath, Ho Kei Cheng, Piotr Dollár, Nikhila Ravi, Kate Saenko, Pengchuan Zhang, and Christoph Feichtenhofer, “Sam 3: Segment anything with concepts,” in *arXiv*, 2025.
- [15] Ho Kei Cheng, Jihoon Chung, Yu-Wing Tai, and Chi-Keung Tang, “Cascadeps: Toward class-agnostic and very high-resolution segmentation via global and local refinement,” in *CVPR*, 2020.
- [16] Lei Ke, Mingqiao Ye, Martin Danelljan, Yifan Liu, Yu-Wing Tai, Chi-Keung Tang, and Fisher Yu, “Segment anything in high quality,” in *NeurIPS*, 2023.
- [17] Ivan Alberico, Marco Cannici, Giovanni Cioffi, and Davide Scaramuzza, “Egocentric event-based vision for ping pong ball trajectory prediction,” in *CVPRW*, 2025.
- [18] AJ Piergiovanni and Michael S. Ryoo, “Fine-grained activity recognition in baseball videos,” in *CVPRW*, 2018.
- [19] Olaf Ronneberger, Philipp Fischer, and Thomas Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *MICCAI*, 2015.
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, “Deep residual learning for image recognition,” in *CVPR*, 2016.
- [21] Carole H. Sudre, Wenqi Li, Tom Vercauteren, Sebastien Ourselin, and M. Jorge Cardoso, “Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations,” in *MICCAIW*, 2017.
- [22] iniVation, “Davis346,” <https://inivation.com/wp-content/uploads/2019/08/DAVIS346.pdf>, 2019.
- [23] Glenn Jocher and Jing Qiu, “Ultralytics yolo11,” <https://github.com/ultralytics/ultralytics>, 2024.