

Multi-Camera Asynchronous Ball Localization and Trajectory Prediction with Factor Graphs and Human Poses

Qingyu Xiao¹, Zulfiqar Zaidi¹, and Matthew Gombolay¹

Abstract—The rapid and precise localization and prediction of a ball are critical for developing agile robots in ball sports, particularly in sports like tennis characterized by high-speed ball movements and powerful spins. The Magnus effect induced by spin adds complexity to trajectory prediction during flight and bounce dynamics upon contact with the ground. In this study, we introduce an innovative approach that combines a multi-camera system with factor graphs for real-time and asynchronous 3D tennis ball localization. Additionally, we estimate hidden states like velocity and spin for trajectory prediction. Furthermore, to enhance spin inference early in the ball’s flight, where limited observations are available, we integrate human pose data using a temporal convolutional network (TCN) to compute spin priors within the factor graph. This refinement provides more accurate spin priors at the beginning of the factor graph, leading to improved early-stage hidden state inference for prediction. Our results show the trained TCN can predict the spin priors with RMSE of 5.27 Hz. Integrating TCN into the factor graph reduces the prediction error of landing positions by over 63.6% compared to a baseline method that utilized an adaptive extended Kalman filter.

I. INTRODUCTION

Object tracking is a well-studied problem in the field of computer vision [1]. However, when it comes to tracking small and fast-moving objects, we find ourselves in a relatively unexplored territory. This presents a unique set of challenges, such as the limited time available for inference and the occurrence of motion blur due to the high speeds involved [2]. Development of algorithms capable of effectively tracking and predicting the trajectories of these fast-moving objects holds tremendous potential across various domains including robots throwing objects in factories [3], dynamic obstacle avoidance for autonomous driving [4], drone tracking [5], as well as estimating the trajectory of balls in sports, such as tennis [6], [7], table tennis [8], [9], [10], [11], [12], [13], [9], [14], soccer [15], cricket [16], baseball [17], and many others [18].

Tennis, as an example, is a sport where precise ball trajectory estimation and prediction are of extreme importance [19]. Human players develop these skills through years of dedicated training. However, for an autonomous robotic system to excel in tennis, it must accurately anticipate the trajectory of an incoming ball and swiftly position itself. Previous studies [6], [7] have demonstrated that an autonomous robotic system can return a ball effectively when faced with non-aggressive and cooperative human shots. Nevertheless,

¹ The authors are affiliated with the Georgia Institute of Technology, Atlanta, GA 30332 USA. Corresponding author: Matthew Gombolay (matthew.gombolay@cc.gatech.edu)




Fig. 1: Fusing time-series human pose data for ball spin estimation with asynchronous camera detections within factor graphs substantially improves performance of trajectory prediction.

when challenged by highly skilled players capable of delivering high-velocity shots with “heavy spin,” the vision system encounters difficulties in precisely determining the current and future ball positions with sufficient lead time for effective robot motion planning.

Prior work has employed a range of techniques for tracking balls [20], [21], [22]. However, these methods often rely on precise synchronization between paired cameras to triangulate the ball’s location, which is difficult to achieve. Attempts to compensate using time filters can result in either data loss or increased localization errors. Stereo cameras [6], [7] provide more precise synchronization but face challenges when the ball is at a distance, leading to localization errors. Some work has used a single camera and leveraged the known ball diameter for distance estimation [23]. However, this method has practical limitations, especially with smaller balls at a distance, such as in tennis.

Our motivation for this paper is to address the challenge of multi-camera asynchronous ball localization while simultaneously enhancing the accuracy of trajectory predictions. This paper explores 3D tennis ball localization and prediction by employing a factor graph framework with human poses to compute spin priors. Factor graphs have traditionally found their niche in the field of robotics, particularly in the context of simultaneous localization and mapping (SLAM). SLAM involves robots fusing data from IMU sensors, capturing their dynamic motions, and stationary landscape landmarks from cameras to estimate their positions [24]. Drawing a parallel,

we conceptualize the tennis ball as a dynamic landmark influenced by aerodynamic forces and restitutional bounces. We associate stationary camera detections with timestamps and connect them through dynamic factors derived from a physics-based dynamics model. This approach allows us to simultaneously estimate hidden states such as the ball’s velocity and spin, significantly refining prediction accuracy. To achieve precise early-stage hidden state estimation in the ball’s trajectory, we incorporate a temporal convolutional network (TCN) trained on human poses. In our final implementation, we employ GTSAM [25] in conjunction with the efficient incremental solver ISAM [26], [27] to handle the inherent nonlinearity of the factor graph model. The principal contributions of our work are summarized below:

- Developed an asynchronous approach utilizing an incremental factor graph achieving fast and accurate estimation of the tennis ball’s location, velocity, and spin.
- Derived potential functions of the factors based on the existing research findings and incorporated into our proposed factor graph.
- Incorporated human poses to calculate spin priors for the factor graph, facilitating an early and precise estimation of the ball’s trajectory. The integration of spin priors led to a substantial 63.6% reduction in the RMSE for the ball’s landing position vs. baseline methods.

II. RELATED WORK

A. Ball Localization

Various ball localization techniques have been reported in literature, typically using multi-camera systems and triangulation to estimate the ball’s location [20], [28], [29], [30]. These approaches distribute cameras around to maximize coverage of the target area but relies on synchronized image capture between paired cameras, which can lead to errors if the tracked object is moving rapidly. In contrast, our approach does not require synchronized frames from the paired cameras, avoiding this issue.

Other work has used stereo cameras for estimating the 3D coordinates of the ball [6], [7]. This method minimizes synchronization discrepancies, as the cameras are interconnected at a hardware level. However, when computing the locations of an object situated at considerable distances from the cameras—several meters or more—this technique encounters noise issues. The proximity of stereo cameras to each other (the small distance between two cameras) results in very small disparities between the captured images. Consequently, even minor noise at a one- or two-pixel level becomes significant compared to the small disparity, leading to substantial errors in the localization process [31]. In contrast, our approach places cameras at a significant distance from each other, which alleviates this issue.

Another approach [23] utilizes a single calibrated camera for ball localization in basketball. It detects the ball’s centroid and apparent diameter within the image, using the latter as a metric to approximate the ball-camera distance. While effective for larger balls and close-range scenarios in sports

like soccer or basketball, this method encounters difficulties with smaller spherical objects like tennis balls.

B. Ball Dynamics and Prediction

Accurate ball trajectory prediction hinges on two critical factors: aerodynamics during flight and bounce dynamics. Aerodynamics during ball flight have been well-studied, considering effects such as drag, Magnus, and lifting forces [32], [33]. Lift force, in particular, causes the ball’s path to curve, adding complexity to prediction. Ball bounce behavior is equally vital. Inbound velocity and spin at the moment of contact profoundly impact the ball’s subsequent path, leading to various outcomes. Prior studies [34] have formulated equations for inbound and outbound dynamics, aiding prediction. While prior robotics-focused works use simpler bounce models [6], [29], [29], they struggle to predict spin-heavy ball trajectories post-bounce. Our approach accounts for these complexities, enhancing ball trajectory prediction.

C. Stroke and Spin detection

Various methods have been documented that utilize ball trajectory to detect the strokes of table tennis players [35], as well as to infer the spin on the ball [36]. An alternative technique proposed in [37] involves applying non-regulation markings to the ball and then tracking these markings to precisely estimate the ball’s spin. Nevertheless, this method’s use of non-regulation markings renders it incompatible with standard equipment. Additionally, previous research has established a significant causal relationship between the executed stroke and the resulting spin on the ball [38]. Therefore, techniques employed for stroke detection can be repurposed to estimate spin dynamics as well. Several studies [39], [40], [41] have attempted to detect human strokes using IMU sensors mounted directly onto the racket. However, this approach might be perceived as cumbersome and aesthetically unattractive, potentially hindering its wide acceptance and usage. In a more recent study, innovative methods have emerged that leverage the analysis of human pose sequences to identify strokes with temporal convolutional networks (TCN) [42]. Inspired from this work, this paper uses a similar approach where we design a regression network to deduce spin priors directly from human poses from a camera.

III. FACTOR GRAPHS FOR BALL LOCALIZATION

Factor graphs for asynchronous ball localization can be viewed as a graph where the size grows as the number of camera detections increases. Once the connections of the factor graph are updated, the location, velocity, and spin of the ball can then be inferred. Since every single detection can update the factor graph and have a localization result from inference, this process is asynchronous without any pairing.

A. Variables and Factors

The variables for our factor graph are shown in Figure 2. Variables are labeled with circles, where X_i is the i^{th} camera pose, L_t is the 3D location of the tennis ball at time step t , V_t is the 3D velocity of the tennis ball at time step t , W_i is the 3D spin of the tennis ball before the i^{th} bounce.

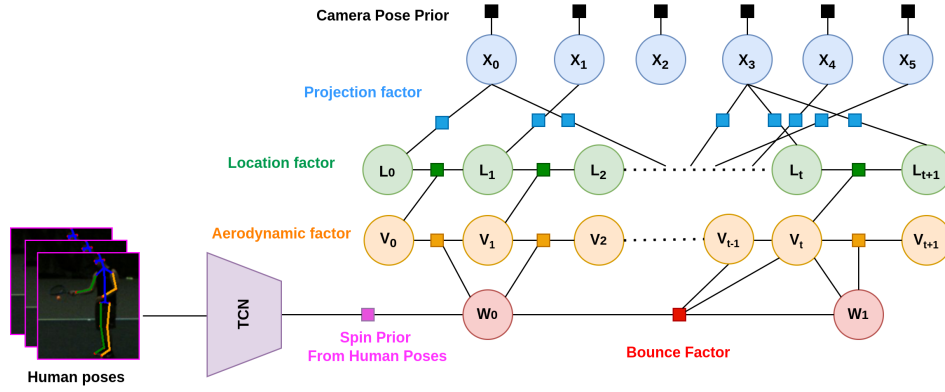


Fig. 2: Example of the factor graph for ball localization at the $t + 1$ time stamp. All the factors are colored squares and all the variables are labeled with circles, where X_i is the i^{th} camera pose, L_t is the 3D location of the tennis ball at time step t , V_t is the 3D velocity of the tennis ball at time step t , W_j is the 3D spin of the tennis ball before the j^{th} bounce. If detection is available in the queue, the factor graph will first expand.

The factors are colored with squares in Figure 2. Factors can be thought of as functions that take the values of their associated variables as input and produce a scalar value or probability that quantifies the compatibility of those values with the factor's constraint (usually associated with an error function). In the proposed factor graph, the camera poses, X , have camera pose priors (black squares) which are computed from extrinsic calibration. The 3D ball locations, L , and the observing camera poses, X , have projection factors (blue squares). The adjacent location changes from previous time stamp L_{t-1} to current time stamp L_t is determined by the previous velocity, V_{t-1} , which form the location factor (green squares). The current velocity, V_t , is affected by previous velocity, V_{t-1} , as well as the previous spin, W_0 , which is derived from aerodynamics. This factor is referred to as the aerodynamic factor (orange square). For the spin, W , we assume the spin of the ball only changes after bouncing from the ground. Therefore, if bouncing is detected, we will form a bounce factor (red square), connecting previous velocity, V_{t-1} , and spin, W_0 , with current velocity, V_t , and spin, W_1 . Finally, we also have our initial spin with a spin prior (pink square), which can be computed from human poses before hitting the ball. If human pose data is not available, we can set W_0 to zero with large noise.

The whole factor graph can be factorized using Equation (1), where η represents all variables in the graph, f_i is the potential function of the i^{th} factor, ξ_i is all the variables connected to the i^{th} factor, and M is the total number of factors in the graph.

$$G(\eta) = \prod_{i=0}^{M-1} f_i(\xi_i) \quad (1)$$

The potential function can be written as (2)

$$f_i(\xi_i) = \exp\left(-\frac{1}{2}\|err_i(\xi_i)\|_{\Sigma_i}^2\right) \quad (2)$$

where err_i is the error function for the i^{th} factor, Σ_i is the covariance matrix for the i^{th} factor. Based on geometric or

physics constraints, we derive our error function for each factor in the following sections.

1) *Projection factor*: The 3D to 2D projection on camera is shown in Equation (3), where K is camera intrinsics, R and T are camera extrinsics (or camera pose X).

$$\pi^{proj}(L) = K [R \ T] L \quad (3)$$

Thus, the error function for the projection factor can be written as shown in Equation (4).

$$err^{proj}(\xi) = u_{detect} - \pi(L) \quad (4)$$

2) *Location factor*: The location factor can be viewed as an integration of velocity at a small time interval. We use Euler formula to compute the next time stamp location, as shown in Equation (5), therefore, the error function for the location factor is given by Equation (6).

$$\pi^{loc}(L_s, V_s) = L_s + V_s [t(s+1) - t(s)] \quad (5)$$

$$err^{loc}(\xi) = L_{s+1} - \pi^{loc}(L_s, V_s) \quad (6)$$

3) *Aerodynamic factor*: We use the formula derived in [32] to compute the air drag and air lift [33], as shown in Equation (7) and (8), where U is the cross product $W \times V$, C_D and C_L are coefficients for air drag and air lift, respectively, A is the area of the ball, ρ_{air} is the air density.

$$F_D(V) = -\frac{1}{2}\rho_{air}AC_D\|V\|V \quad (7)$$

$$F_L(V, W) = \frac{1}{2}\rho_{air}AC_L\|V\|^2U \quad (8)$$

Using the Newton's second law, we can derive the acceleration with

$$acc(V, W) = (F_D(V) + F_L(V, W))/m + g \quad (9)$$

$$\pi^{aero}(V_s, W_j) = V_s + acc(V_s, W_j) [t(s+1) - t(s)] \quad (10)$$

Therefore, the error function for the velocity factor is

$$err^{aero}(\xi) = V_{s+1} - \pi^{aero}(V_s, W_j) \quad (11)$$

4) *Bounce factor*: For the bounce factor, we extend the formula derived in [34] to estimate the outbound velocity and spin in 3D space as described in Equation (12) - (18), where R is the outer radius of the tennis ball, R_1 is the average radius of the shell, and e_z is the restitution of the tennis ball.

$$\pi^{bc}(\xi) = [V_{x,t} \quad V_{y,t} \quad V_{z,t} \quad W_{x,j} \quad W_{y,j} \quad W_{z,j}]^T \quad (12)$$

where

$$V_{x,t} = \frac{RW_{y,j-1} + (1.5R^2/R_1^2)V_{x,t-1}}{1 + 1.5R^2/R_1^2} \quad (13)$$

$$V_{y,t} = \frac{-RW_{x,j-1} + (1.5R^2/R_1^2)V_{y,t-1}}{1 + 1.5R^2/R_1^2} \quad (14)$$

$$V_{z,t} = -e_z V_{z,t-1} \quad (15)$$

$$W_{x,j} = -V_{y,t}/R \quad (16)$$

$$W_{y,j} = V_{x,t}/R \quad (17)$$

$$W_{z,j} = W_{z,j-1} \quad (18)$$

Therefore, the error function for the bounce factor is

$$err^{bc}(\xi) = [V_{s+1}, W_{j+1}]^T - \pi^{bc}(V_s, W_j) \quad (19)$$

B. Incremental Inferences

As illustrated in Figure 2, the depth of the factor graph grows linearly with the number of camera detections. For each new detection, we will first add new connections to the factor graph, then update values for all the variables in the graph. This approach avoids the need for a synchronization process or pairing, which improves the flexibility of the system.

Factor graph inferences need to solve the maximize likelihood optimization problem, since the size of the factor graph will become very large as more detections come, which makes traditional optimizers impossible to run in real time. In this paper, we use the factor graph package, GTSAM, to solve the factor graph we build as in Figure 2. This package uses ISAM2 algorithm [26], [27], which can relinearize and reorder the factor graph so that the graph can be optimized incrementally, and hence achieve real time inference speed. The average inference speed for the proposed factor graph approach is 335 Hz. The linearization process is described in Equation (20), where J is the Jacobian of the error function whitened by noise model.

$$\operatorname{argmin}_{\eta} -\log(G(\eta)) \approx \operatorname{argmin}_{\eta} \|J\eta - b\|^2 \quad (20)$$

C. Prediction

The prediction of the ball will use the dynamic equations we derived in (9) and use Runge-Kutta 4th order to integrate. The ODE for ball during flight is written as

$$\begin{bmatrix} \dot{L} \\ \dot{V} \\ \dot{W} \end{bmatrix} = \begin{bmatrix} V \\ acc(V, W) \\ 0 \end{bmatrix} \quad (21)$$

The integration will also check if the ball bounce on the ground ($V_z < 0$ and $Lz < z_{\text{ground}}$). If the bounce is detected,

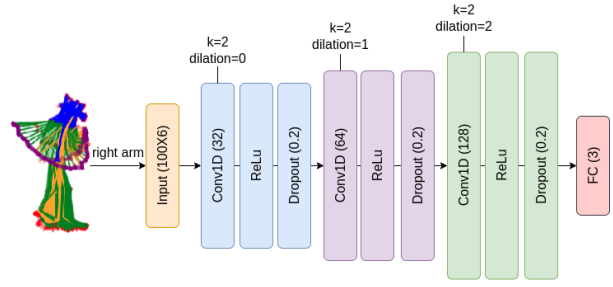


Fig. 3: TCN for spin prior regression with human pose sequences. k is the kernel size. The model input is the right arm position (wrist, elbow, and shoulder) in 100 frames, and the output is the spin vector in \mathbb{R}^3 .

we will terminate the current ODE integration and apply the bounce dynamics to the current state velocity and spin as derived in equation (12), followed by a new integration starting with the new velocity and spin.

IV. SPIN PRIORS FROM HUMAN POSES

A. Data Preparation

We collect 200 tennis trajectories from three different participants. AlphaPose [43] is employed to detect human poses in camera frames. Each trajectory is linked with 100 frames of human poses, emphasizing the key points on the right arm, shoulder, elbow, and wrist. These poses are centered at the “hip.”

We use an iterative approach to label the spins in each trajectory. Factor graph inference, known for its capacity to update variable values within the graph, gains confidence with accumulating detections. However, initially, our factor graph lacks knowledge about suitable spin priors. To tackle this, we start with initial spin estimations and conduct inference across all detections until the trajectory’s end. As the initial spin estimations evolve, we utilize the inference from the final detection as input for a new iteration. This iterative process continues until the spin estimate converges to a value that is no longer changing significantly, which is then used to label the spin for each trajectory.

B. Neural Network Model

For sequence modeling, we use temporal convolutional neural networks (TCNs) to estimate spin priors from sequential human poses using regression. Figure 3 illustrates the architecture of our model. We empirically found that TCNs outperformed recurrent neural networks (RNNs) and long short-term memory networks (LSTMs) for this task, which is supported by prior work [44].

V. EXPERIMENTS AND RESULTS

A. Hardware Setups

In order to track and detect the tennis ball on an indoor tennis court, we use a multi-cameras setup (Figure 4). The system consists of six Blackfly S cameras, each connected to a Jetson Xavier. The cameras are mounted on the front and back walls of the court. Note the proposed method is adaptable to configurations with two or more cameras. In order to




Fig. 4: Camera setup on an indoor tennis court.

achieve optimal coverage of the tennis court area, we have strategically selected six cameras. This specific arrangement guarantees that every spot within the area can be observed by at least two cameras, thus ensuring thorough coverage while also maintaining a manageable level of computational complexity. For extrinsic calibration, an AprilTag [45] is placed on the court at a known position which all six cameras can observe.

The Jetson Xaviers acquire color images from the cameras. We then run a YOLO tiny [46] model trained to detect tennis balls on the isolated parts of the image. The image acquisition and ball detection algorithm runs at ~ 140 Hz on each Jetson Xavier. The average inference speed for the factor graph is 335 Hz. Although the inference speed is fast, we admit that the factor graph method cannot process all the detections if all six cameras observe the tennis ball (at the center area of the court). In this case, the latest detection in the stack will be used.

We use Ethernet communication to share ball detections from the Jetson Xaviers to the router and central computer, which uses the detections to run our ball trajectory estimation algorithm. This results in an average latency of one millisecond between the Jetson Xaviers sending the detections and the central computer receiving them. This represents a substantial improvement compared to a prior decentralized multi-camera configuration that relied on WiFi-based communication, as outlined in the work by Zaidi et al. [6]. In their setup, they achieved a detection rate of 25 Hz per camera and exhibited a considerably higher average latency of 110 milliseconds between the Jetson Nano and the central computer.

B. Spin priors from Human Poses

We partitioned our dataset into training and validation sets at an 80:20 ratio, aiming to provide a robust training ground for our model while retaining a substantial portion for validation. The TCN was utilized as the backbone of our neural network, chosen for its prowess in handling sequential data effectively. To quantify the performance of our model, we utilized the Root Mean Square Error (RMSE) metric. This choice is grounded in its efficacy in measuring the average magnitude of error between predicted and actual values, which is particularly pertinent in our context of estimating spin priors. Our model demonstrated promising

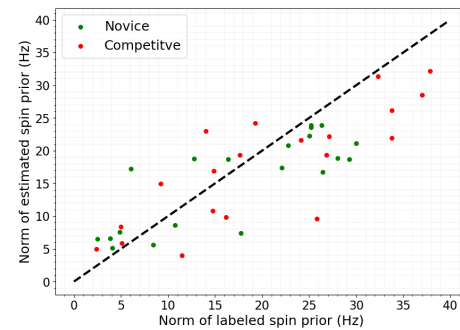


Fig. 5: Scatterplot of labeled vs. estimated spin priors in the validation dataset, showing a strong correlation.

results with a training RMSE of 2.95 Hz and a validation RMSE of 5.27 Hz. Figure 5 shows the scatter plot of the spin prior estimation across the validation dataset. The figure illustrates a linear distribution between the labeled spin prior and the estimated spin prior. Notably, both novice and competitive players exhibit strong linear relationships, with Pearson correlation coefficients (R) of 0.807 and 0.815, respectively, and highly significant $p < 0.001$ values for both scenarios.

C. Prediction Results

In this section, we focus on assessing the accuracy of predictions pertaining to both the initial and the second landing positions of the ball. These two landing positions are chosen because wheelchair tennis rules permit players to strike the ball either before its initial landing or prior to the second landing.

To assess the performance of our approach, we conduct several comparative evaluations. Firstly, we compare it with the method proposed in [6], which employed a similar multi-camera decentralized setup and utilized an Extended Kalman Filter (EKF) for tennis ball trajectory estimation and prediction. We also explore an alternative by using an adaptive EKF (AEKF) that models the spin as a hidden state instead of using the standard EKF. Furthermore, we conduct an ablation study by removing the spin priors from our approach. Subsequently, we compare our method’s performance when employing the estimated spin from the TCN against using the spin labels directly.

In order to analyze how prediction error changes with different number of observations from cameras, we bin each ball trajectory evenly into quartiles or “stages:” Stage 1 contains the first 25% of the observations; Stage 2 contains the next 25% of the observations, etc. We then analyze the mean and standard deviation of the RMSE for each stage (Figure 6). For example, Stage 1 results average across the accuracy of ball estimation from when there is only a single detection until a total of 25% of detections for that trajectory are available to the algorithm. We chose this binning procedure as ball detections happen at a non-uniform cadence (i.e., asynchronously) and the trajectories are heterogeneous in duration and in the timing of the bounces.

The prediction results in Figure 6 show that the prediction




Fig. 6: The average RMSE of 4 stages for predicting (a) first landing and (b) second landing position.

error reduces across time, with the biggest drop occurring after the first 25% of the trajectory has been completed (from Stage 1 to Stage 2). Furthermore, it’s notable that incorporating our TCN for spin prediction from human pose data (Figure 6, green) significantly enhances the accuracy of second bounce location prediction (Figure 6(b)) compared to the baseline approaches, AEKF (purple) and EKF (red). It also outperforms our approach without considering spin (blue) and closely approaches the trajectory prediction achieved using the labeled spin values (orange). Notice that on stage 4, the accuracy of our approach without spin nearly matches that of the method with spin. This is because the increased accumulation of camera observations enable the factor graph to more accurately infer the ball’s spin. Consequently, the results at this stage show only a slight improvement when incorporating spin.

The advantage of incorporating spin information becomes evident primarily in predicting the second bounce, as the factor graph can effectively handle in-flight spin dynamics but requires spin details to accurately model the impact of spin on the ball’s first bounce. This hypothesis gains support from the observation that the no-spin model (blue) shows a significant improvement in predicting the second bounce location (Figure 6(b)) in Stages 3-4, which typically occur after the initial bounce.

To delve deeper into the potential impact of Magnus and lift forces resulting from spin, especially prevalent among highly skilled tennis players renowned for employing powerful spins, we enlisted the expertise of a competitive tennis player to hit shots with substantial spin. Assessment of the performance of our approach (green) in accurately predicting the locations of the first and second bounces is illustrated in Figure 7, where we also compare it with baseline methods for a single trajectory. The most accurate predictions are achieved when our approach is combined with labeled spin values (orange). However, our neural network model faced challenges in accurately approximating spin priors, likely due to limited training data and the absence of information regarding the racket’s pose. Nonetheless, our approach consistently outperformed the baseline EKF (red) and AEKF (purple) methods, with an overall 63.6% improvement on prediction accuracy.

It is important to highlight that incorporating information about the racket’s pose could enhance spin estimation accuracy, complementing insights derived from the analysis of

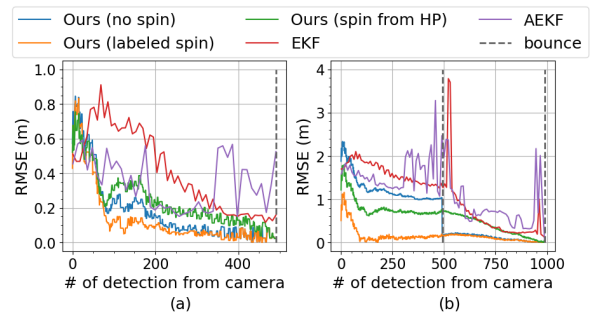


Fig. 7: RMSE results with respect to number of observation (a) initial landing position prediction errors (b) second landing position errors for a single shot with heavy spin played by a professional tennis player.

human poses. Overall, the substantial improvement achieved through the use of labeled spin priors underscores the immense potential of accurate spin estimation.

VI. LIMITATIONS AND FUTURE WORK

We note key limitations in our approach for ball trajectory prediction. When dealing with balls that have been struck with a substantial spin, our model encounters challenges in precisely estimating the spin prior from the human pose data. This limitation stems from the small amount of training data and the absence of racket pose details and the type of grip used (e.g., Eastern, Western, or Continental). Despite these challenges, the prediction using labeled spin priors elucidate a substantial potential to enhance the accuracy of landing position predictions if spin parameters can be precisely extrapolated from human postures, indicating promising avenues for further improvement. Further, we found that we had to perform ad hoc modifications for the coefficients associated with the bounce dynamics equations from prior work to be robust to high-spin dynamics (e.g., ball-ground slip). We also noted that the accuracy of the initial landing position prediction sometimes deteriorated due to human pose priors. This may be attributed to inaccuracies in the aerodynamic equations used to model the ball’s trajectory. In future work, we propose to augment our factor graphs to have a learnable factor for aerodynamics and bouncing.

VII. CONCLUSIONS

In this study, we introduce an innovative asynchronous approach to predicting the state of a tennis ball utilizing a factor graph. To facilitate precise inference, we developed a neural network utilizing Temporal Convolutional Network (TCN) to infer spin priors from human postures. Our findings demonstrate that the algorithm can proficiently and accurately estimate the ball’s location, yielding smooth trajectory plots.

ACKNOWLEDGMENT

This work was supported by Google under Award 004686. Special thanks to Dr. Pannag Sanketi and Dr. David D’Ambrosio for their invaluable support and insightful feedback throughout the duration of this work.

REFERENCES

- [1] A. Yilmaz, O. Javed, and M. Shah, "Object tracking: A survey," *ACM Comput. Surv.*, vol. 38, no. 4, p. 13–es, dec 2006. [Online]. Available: <https://doi.org/10.1145/1177352.1177355>
- [2] Z. Zhang, F. Wu, Y. Qiu, J. Liang, and S. Li, "Tracking small and fast moving objects: A benchmark," in *Computer Vision – ACCV 2022*, L. Wang, J. Gall, T.-J. Chin, I. Sato, and R. Chellappa, Eds. Cham: Springer Nature Switzerland, 2023, pp. 552–569.
- [3] T. Frank, U. Janoske, A. Mittnacht, and C. Schroedter, "Automated throwing and capturing of cylinder-shaped objects," in *2012 IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 5264–5270.
- [4] Y. Luo, P. Cai, A. Bera, D. Hsu, W. S. Lee, and D. Manocha, "Porca: Modeling and planning for autonomous driving among many pedestrians," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3418–3425, 2018.
- [5] I. Guvenc, F. Koohifar, S. Singh, M. L. Sicitiu, and D. Matolak, "Detection, tracking, and interdiction for amateur drones," *IEEE Communications Magazine*, vol. 56, no. 4, pp. 75–81, 2018.
- [6] Z. Zaidi, D. Martin, N. Belles, V. Zakharov, A. Krishna, K. M. Lee, P. Wagstaff, S. Naik, M. Sklar, S. Choi, et al., "Athletic mobile manipulator system for robotic wheelchair tennis," *IEEE Robotics and Automation Letters*, vol. 8, no. 4, pp. 2245–2252, 2023.
- [7] F. Yang, Z. Shi, S. Ye, J. Qian, W. Wang, and D. Xuan, "Varsm: Versatile autonomous racquet sports machine," in *2022 ACM/IEEE 13th International Conference on Cyber-Physical Systems (ICCPs)*. IEEE, 2022, pp. 203–214.
- [8] S. Gomez-Gonzalez, S. Prokudin, B. Schölkopf, and J. Peters, "Real time trajectory prediction using deep conditional generative models," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 970–976, 2020.
- [9] D. B. D'Ambrosio, J. Abelian, S. Abeyruwan, M. Ahn, A. Bewley, J. Boyd, K. Choromanski, O. Cortes, E. Coumans, T. Ding, et al., "Robotic table tennis: A case study into a high speed learning system," *arXiv preprint arXiv:2309.03315*, 2023.
- [10] S. W. Abeyruwan, L. Graesser, D. B. D'Ambrosio, A. Singh, A. Shankar, A. Bewley, D. Jain, K. M. Choromanski, and P. R. Sanketi, "i-sim2real: Reinforcement learning of robotic policies in tight human-robot interaction loops," in *Conference on Robot Learning*. PMLR, 2023, pp. 212–224.
- [11] L. Chen, R. Paleja, M. Ghuy, and M. Gombolay, "Joint goal and strategy inference across heterogeneous demonstrators via reward network distillation," in *Proceedings of the 2020 ACM/IEEE international conference on human-robot interaction*, 2020, pp. 659–668.
- [12] L. Chen, S. Jayanthi, R. R. Paleja, D. Martin, V. Zakharov, and M. Gombolay, "Fast lifelong adaptive inverse reinforcement learning from demonstrations," in *Conference on Robot Learning*. PMLR, 2023, pp. 2083–2094.
- [13] L. Chen, R. Paleja, and M. Gombolay, "Learning from suboptimal demonstration via self-supervised reward regression," in *Conference on robot learning*. PMLR, 2021, pp. 1262–1277.
- [14] K. Mülling, J. Kober, O. Kroemer, and J. Peters, "Learning to select and generalize striking movements in robot table tennis," *The International Journal of Robotics Research*, vol. 32, no. 3, pp. 263–279, 2013.
- [15] E. A. Nurrohmah, B. S. Bayu, M. M. Bachtiar, I. K. Wibowo, and R. Adryantoro, "Detecting features of middle size soccer field using omnidirectional camera for robot soccer ersow," in *2020 International Conference on Smart Technology and Applications (ICoSTA)*. IEEE, 2020, pp. 1–6.
- [16] C. J. Baker, "A calculation of cricket ball trajectories," *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, vol. 224, no. 9, pp. 1947–1958, 2010. [Online]. Available: <https://doi.org/10.1243/09544062JMES1973>
- [17] L. W. Alaways and M. Hubbard, "Experimental determination of baseball spin and lift," *Journal of Sports Sciences*, vol. 19, no. 5, pp. 349–358, 2001, pMID: 11354614. [Online]. Available: <https://doi.org/10.1080/02640410152006126>
- [18] J. R. Kensrud, "Determining aerodynamic properties of sports balls in situ," Ph.D. dissertation, Washington State University, 2010.
- [19] K. Weber, S. Pieper, and T. Exler, "Characteristics and significance of running speed at the australian open 2006 for training and injury prevention," *Journal of Medicine and Science in Tennis*, vol. 12, no. 1, pp. 15–18, 2007.
- [20] J. Ren, J. Orwell, G. A. Jones, and M. Xu, "Tracking the soccer ball using multiple fixed cameras," *Computer Vision and Image Understanding*, vol. 113, no. 5, pp. 633–642, 2009.
- [21] P. Wong, H. Myint, L. Dooley, and A. Hopgood, "A multi-view automatic table tennis umpiring framework," *Proceedings of the Institution of Mechanical Engineers, Part P: Journal of Sports Engineering and Technology*, p. 175433712311171433, 2023.
- [22] K. M. Lee, A. Krishna, Z. Zaidi, R. Paleja, L. Chen, E. Hedlund-Botti, M. Schrum, and M. Gombolay, "The effect of robot skill level and communication in rapid, proximate human-robot collaboration," in *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI '23. New York, NY, USA: Association for Computing Machinery, 2023, p. 261–270. [Online]. Available: <https://doi.org/10.1145/3568162.3577002>
- [23] G. Van Zandycke and C. De Vleeschouwer, "3d ball localization from a single calibrated image," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 3472–3480.
- [24] F. Dellaert, M. Kaess, et al., "Factor graphs for robot perception," *Foundations and Trends® in Robotics*, vol. 6, no. 1-2, pp. 1–139, 2017.
- [25] F. Dellaert, "Factor graphs and gtsam: A hands-on introduction," *Georgia Institute of Technology, Tech. Rep.*, vol. 2, p. 4, 2012.
- [26] M. Kaess, A. Ranganathan, and F. Dellaert, "isam: Incremental smoothing and mapping," *IEEE Transactions on Robotics*, vol. 24, no. 6, pp. 1365–1378, 2008.
- [27] M. Kaess, H. Johannsson, R. Roberts, V. Ila, J. J. Leonard, and F. Dellaert, "isam2: Incremental smoothing and mapping using the bayes tree," *The International Journal of Robotics Research*, vol. 31, no. 2, pp. 216–235, 2012.
- [28] M. Fazio, K. Fisher, and T. Fujinami, "Tennis ball tracking: 3-d trajectory estimation using smartphone videos," *Department of Electrical Engineering, Stanford University*, 2018.
- [29] J. Tebbe, Y. Gao, M. Sastre-Rienietz, and A. Zell, "A table tennis robot system using an industrial kuka robot arm," in *Pattern Recognition: 40th German Conference, GCPR 2018, Stuttgart, Germany, October 9-12, 2018, Proceedings 40*. Springer, 2019, pp. 33–45.
- [30] S. Gomez-Gonzalez, Y. Nemmour, B. Schölkopf, and J. Peters, "Reliable real-time ball tracking for robot table tennis," *Robotics*, vol. 8, no. 4, 2019.
- [31] C. F. Olson and H. Abi-Rached, "Wide-baseline stereo vision for terrain mapping," *Machine Vision and Applications*, vol. 21, no. 5, pp. 713–725, Feb. 2009. [Online]. Available: <https://doi.org/10.1007/s00138-009-0188-9>
- [32] A. J. Cooke, "An overview of tennis ball aerodynamics," *Sports Engineering*, vol. 3, no. 2, pp. 123–129, 2000.
- [33] R. Cross and C. Lindsey, "Measurements of drag and lift on tennis balls in flight," *Sports Engineering*, vol. 17, pp. 89–96, 2014.
- [34] R. Cross, "Measurements of the horizontal coefficient of restitution for a superball and a tennis ball," *American Journal of Physics*, vol. 70, no. 5, pp. 482–489, 2002.
- [35] K. M. Kulkarni, R. S. Jamadagni, J. A. Paul, and S. Shenoy, "Table tennis stroke detection and recognition using ball trajectory data," *arXiv preprint arXiv:2302.09657*, 2023.
- [36] J. Tebbe, L. Klamt, Y. Gao, and A. Zell, "Spin detection in robotic table tennis," in *2020 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2020, pp. 9694–9700.
- [37] T. Gossard, J. Tebbe, A. Ziegler, and A. Zell, "Spindoe: A ball spin estimation method for table tennis robot," 2023.
- [38] C. Rigozzi, J. Cox, G. A. Vio, W. L. Martens, and P. Poronnik, "The effect of spin level and ball exit speed on forearm muscle activity in the tennis forehand stroke," *International Journal of Sports Science & Coaching*, vol. 17, no. 1, pp. 123–133, 2022. [Online]. Available: <https://doi.org/10.1177/17479541211007611>
- [39] P. Blank, B. H. Groh, and B. M. Eskofier, "Ball speed and spin estimation in table tennis using a racket-mounted inertial sensor," in *Proceedings of the 2017 ACM International Symposium on Wearable Computers*, 2017, pp. 2–9.
- [40] H. Zhao, S. Wang, G. Zhou, and W. Jung, "Tenniseye: tennis ball speed estimation using a racket-mounted motion sensor," in *Proceedings of the 18th international Conference on information Processing in Sensor Networks*, 2019, pp. 241–252.
- [41] P. Blank, J. Hoßbach, D. Schuldhuis, and B. M. Eskofier, "Sensor-based stroke detection and stroke type classification in table tennis," in

Proceedings of the 2015 ACM International Symposium on Wearable Computers, 2015, pp. 93–100.

- [42] K. M. Kulkarni and S. Shenoy, “Table tennis stroke recognition using two-dimensional human pose estimation,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 4576–4584.
- [43] H.-S. Fang, J. Li, H. Tang, C. Xu, H. Zhu, Y. Xiu, Y.-L. Li, and C. Lu, “Alphapose: Whole-body regional multi-person pose estimation and tracking in real-time,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [44] S. Bai, J. Z. Kolter, and V. Koltun, “An empirical evaluation of generic convolutional and recurrent networks for sequence modeling,” *arXiv preprint arXiv:1803.01271*, 2018.
- [45] E. Olson, “Apriltag: A robust and flexible visual fiducial system,” in *2011 IEEE International Conference on Robotics and Automation*, 2011, pp. 3400–3407.
- [46] R. Huang, J. Pedoeem, and C. Chen, “Yolo-lite: a real-time object detection algorithm optimized for non-gpu computers,” in *2018 IEEE international conference on big data (big data)*. IEEE, 2018, pp. 2503–2510.