Athletic Mobile Manipulator System for Robotic Wheelchair Tennis

Zulfiqar Zaidi*, Daniel Martin*, Nathaniel Belles, Viacheslav Zakharov, Arjun Krishna, Kin Man Lee, Peter Wagstaff, Sumedh Naik, Matthew Sklar, Sugju Choi, Yoshiki Kakehi, Ruturaj Patil, Divya Mallemadugula, Florian Pesce, Peter Wilson, Wendell Hom, Matan Diamond, Bryan Zhao, Nina Moorman, Rohan Paleja, Letian Chen, Esmaeil Seraj, and Matthew Gombolay

Georgia Institute of Technology

Corresponding Email: zzaidi8@gatech.edu

***Abstract*— Athletics are a quintessential and universal expression of humanity. From French monks who in the 12^{th} century invented *jeu de paume*, the precursor to modern lawn tennis, back to the *K'iche'* people who played the Maya Ballgame as a form of religious expression over three thousand years ago, humans have sought to train their minds and bodies to excel in sporting contests. Advances in robotics are opening up the possibility of robots in sports. Yet, key challenges remain, as most prior works in robotics for sports are limited to pristine sensing environments, do not require significant force generation, or are on miniaturized scales unsuited for joint human-robot play. In this paper, we propose the first open-source, autonomous robot for playing regulation wheelchair tennis. We demonstrate the performance of our full-stack system in executing *ground strokes* and evaluate each of the system's hardware and software components. The goal of this paper is to (1) inspire more research in human-scale robot athletics and (2) establish the first baseline towards developing a robot in future work that can serve as a teammate for mixed, human-robot doubles play. Our paper contributes to the science of systems design and poses a set of key challenges for the robotics community to address in striving towards a vision of human-robot collaboration in sports.**

I. INTRODUCTION

Sports have been an integral part of human history as they have served as an important venue for humans to push their athletic and mental abilities. Sports transcend culture [1], [2] and provide strong physiological and social benefits [3]. The physical and mental training from sports is widely applicable to multiple aspects of life [4], [5]. Developing robotic systems for competitive sports can increase participation in sports and be used for sports training [6], [7], thus benefiting society by promoting a healthier lifestyle and increased economic activity. Furthermore, sports serve as a promising domain for developing new robotic systems through the exploration of high-speed athletic behaviors and human-robot collaboration.

Researchers in robotics have sought to develop autonomous systems for playing various sports such as soccer and table tennis. RoboGames [8] and RoboCup [9] have inspired the next generation of roboticists to solve challenges in sensing, navigation, and control. However, many of the advancements or techniques used in solving these obstacles have been restricted to overly miniaturized (e.g. Robo-Soccer [10], Robo-sumo [11]), relatively stationary [12], or highly idealized/controlled environments [11], [13], [14]. The hardware is often specifically engineered to the sport [15], making re-purposing designs and components for other applications difficult. We address these limitations of prior work by proposing an autonomous system for regulation wheelchair tennis by leveraging only commercial off-the-shelf (COTS) hardware. Our system named ESTHER (Experimental Sport Tennis wHEelchair Robot) after arguably the greatest wheelchair tennis player, Esther Vergeer.

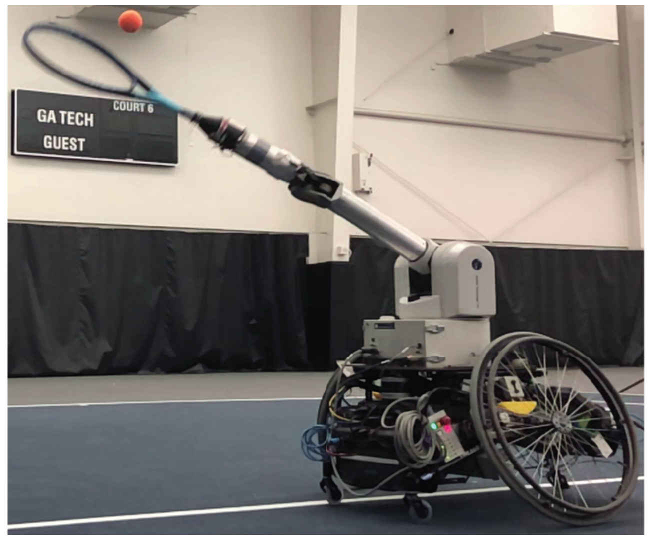

Fig. 1: ESTHER hitting a tennis ground stroke.

Tennis is a challenging sport requiring players to exhibit accurate trajectory estimation, strategic positioning, tactical shot selection, and dynamic racket swings. The design of a system that meets these demands comes with a multitude of complications to address: precise perception, fast planning, low-drift control, and highly-responsive actuation. These challenges must be resolved in a framework efficient enough to respond in fractions of a second [16].

In this paper, we present the system design and an empirical analysis to enable ESTHER (Fig. 1) to address these challenges. Our design adheres to ITF tennis regulations by utilizing a legal tennis wheelchair and an anthropomorphic robot arm, which were required to obtain permission to field ESTHER on an NCAA Division 1 indoor tennis facility. We contribute to the science of robotic systems by (1) Demonstrating a fully mobile setup that includes a decentralized low-latency vision system and a light-weight (190 lbs) mobile manipulator that can be easily set up on any tennis court within 30 minutes, (2) A design for a motorized wheelchair base that is capable of meeting the athletic demands of tennis

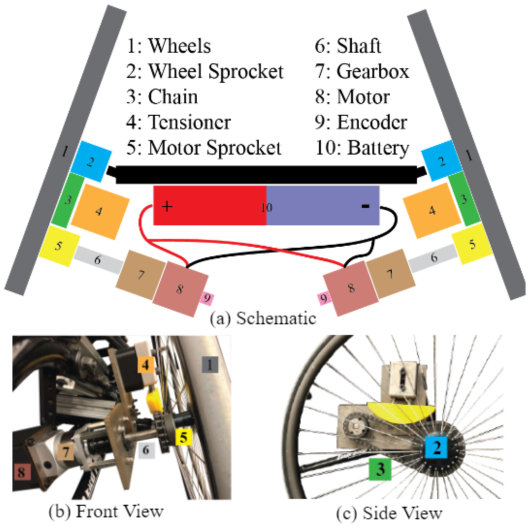

Fig. 2: The schematic (Fig. a) and images (Fig. b-c) of the mechanical assembly of the chain drive system.

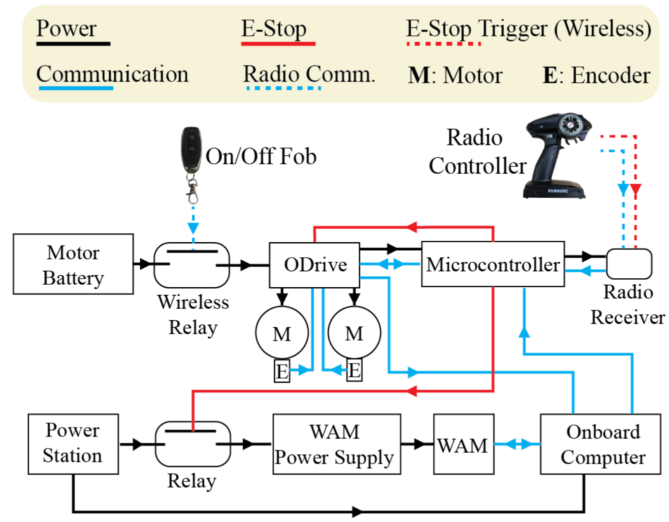

Fig. 3: Electrical and communication diagram for ESTHER's mobile base and robotic arm.

under the full-load of the manipulator, onboard computer, and batteries, and (3) Demonstrate that our planning and control algorithm is capable of hitting powerful ground strokes that traverse the court. We open-sourced our designs and software on our project website: `https://core-robotics-lab.github.io/Wheelchair-Tennis-Robot/`

II. Related Work

Systems that play tennis must be able to rapidly traverse the court, control the racket at high speeds with precision, and deliver sufficient force to withstand the impulse from ball contact. Existing systems that attempt to play other racket sports [15], [17] typically use pneumatically driven manipulators for quick and powerful control; however, these systems are stationary and are thus unsuitable. Other works [18]–[23] have mounted robotic manipulators on rails, but these systems do not support full maneuverability. Furthermore, these robot environments are heavily modified to suit the robot, and thus, are rigid and unsuitable for human proximity.

There have been several recent works in creating agile mobile manipulators [24], [25]. For example, [24] proposes a UR10 arm with 6 degrees of freedom (DoF) mounted on a 3-DoF omni-directional base for catching balls. This system works well for low-power, precision tasks, but lacks the force required for a tennis swing. VaRSM [25] is a system designed to play racket sports including tennis and is most similar to our work. The system proposed uses a custom 6-DoF arm and custom swerve-drive platform as its base. While this system appears effective, it is difficult to reproduce given that the hardware is custom-made. Further, the system was developed by a company that has not open-sourced its implementation, and experimental details are lacking to afford a proper baseline. In contrast, ESTHER is a human-scale robot constructed from a COTS 7-DoF robot manipulator arm and a regulation wheelchair with an open-sourced design and codebase. This enables our system to serve as a baseline for human-scale athletic robots while being reproducible and adaptable for performing mobile manipulation tasks in a variety of settings, such as healthcare [26]–[28], manufacturing [29]–[31], and more [32].

III. Methods

We detail the ESTHER's hardware (Sec. III-A), perception (Sec. III-B.1), planning (Sec. III-B.2), and control (Sec. III-B.3) components. ESTHER's design meets the athletic demands of tennis and is easy and quick to set up.

A. Hardware

We mounted a 7-DoF high-speed Barrett WAM on a motorized Top End Pro Tennis Wheelchair (Fig. 1).

Wheelchair – We designed a chain-drive system to deliver power from the motors to the wheels. The schematic of the system's mechanical assembly is illustrated in Fig. 2. Power from the battery is delivered through an ODrive motor controller to two ODrive D6374 motors. The motors for both wheels are coupled to a 1:10 ratio speed-reducer planetary gearbox. The gearbox output shaft is coupled with the wheel through a chain and sprocket system that provides an additional 1:2 speed reduction to give a total reduction of 1:20. With the motors rotating at maximum speed, the wheelchair can achieve linear velocities of up to $10\,\mathrm{m\,s^{-1}}$ and in-place angular yaw velocity of up to $20\,\mathrm{rad\,s^{-1}}$. A chain drive system has higher durability, torque capacity, increased tolerances, and a simpler design versus belt drive or friction drive systems.

The electrical system to power and control the mobile wheelchair base consists of a DC battery, motors, a motor controller, a wireless radio controller and receiver pair, and a microcontroller as illustrated in Fig. 3. The radio controller allows the user to remotely swap between idle, manual, and autonomous modes. In idle mode, motors are de-energized. In manual mode, a human can remotely control the motion of the wheelchair. In autonomous mode, the wheelchair's motion is controlled by the onboard computer that communicates the desired wheel velocities to the microcontroller. The radio

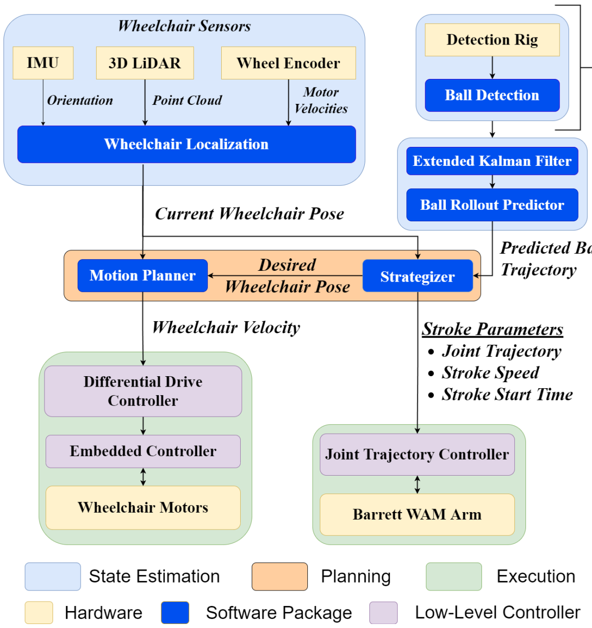

Fig. 4: ESTHER's System Architecture.

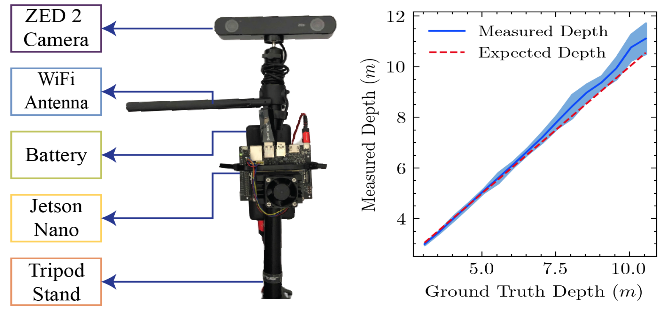

(a) (b)

Fig. 5: Fig. 5a depicts the ball detection rig, and Fig. 5b depicts the measured vs. actual distance of a ball.

controller also has an emergency stop (E-stop) button. If the E-stop button is pressed, the microcontroller commands both the motor speeds to be zero and turns off the power to the arm. Maintaining a remote E-stop enables the whole system to be tether-free and safely operable from a distance.

Barrett WAM – We mount a "HEAD Graphene Instinct Power" tennis racket at the end of the WAM arm using a 3D-printed connector (Fig. 1). The design and visualization of potential failure modes are available on the project website. The arm is powered by a 500 W mobile power station that outputs AC current as illustrated in Fig. 3. The power station is used to power the onboard computer, the LIDAR, and the WAM. The computer requires up to 300 W when charging and running the full code stack. The LIDAR requires 30 W when running at maximum frequency. The WAM requires 50 W when static and up to 150 W while swinging.

B. System Design

We decompose the ESTHER into three main components: Sensing and State Estimation, Planning, and Low-Level Controllers (Fig. 4). We leverage ROS [33], [34] to build an interconnected modular software stack using a combination of open-source packages [34] and custom packages.

1) Sensing and State Estimation: To track a tennis ball as it moves across the court, we built six ball detection rigs (Fig. 5a). Each rig consists of a Stereolabs ZED2 stereo camera, which is connected to an NVIDIA Jetson Nano that communicates to a central onboard computer over Wi-Fi. Each rig is mounted on a 13-foot tripod stand, is placed to maximize coverage of the court (Fig. 8), and is calibrated using an AprilTag [35] box.

Leveraging our camera rigs, we then employ multiple computer vision techniques to detect an airborne tennis ball: color thresholding, background subtraction, and noise removal to locate the pixel center of the largest moving colored tennis ball. Orange tennis balls were used since orange provided a strong contrast against the green/blue tennis courts. The ball coordinates are computed relative to the ball detection rig's reference frame by applying epipolar geometry on the pixel centers from the rig. The ball's pose was then mapped to the world frame. The positional error covariance was modeled through a quadratic relationship (i.e., aD^2+bD+c) where D is the ball's distance to the camera and a, b, and c are tuned via experiments of ball rolling along a linear rail (Fig. 5b). A quadratic expression is utilized as it provides a simple and accurate approximation of the measurement error. By pursuing a distributed vision approach, image data could be efficiently processed locally on each Jetson Nano and sent to the onboard computer via Wi-Fi. These detection modules are able to process the stereo camera's 1080p images, produce a positional estimate, and transmit the measurement to the central computer within 100 ms at a frequency of 25 Hz. The vision system is able to achieve up to $150\,\text{estimates/s}$.

a) Ball EFK and Roll-out Prediction: The detection rigs' ball position estimates are fused through a continuous-time Extended Kalman Filter (EKF) [36] to produce a single pose estimate. The EKF is based on the robot_localization package [37], which is augmented to incorporate ballistic trajectory, inelastic bouncing, court friction, and ball-air interactions [38] into the state-estimate prediction. Our localization approach allows the ball detection errors to be merged into a single covariance estimate, informing how much confidence should be placed on the ball's predicted trajectory. To enhance the EKF predictions, the state estimation is reset when a tennis ball is first detected, and any measurement delays are handled by reverting the EKF to a specified lag time and re-applying all measurements to the present. By rolling out the EKF's state predictions forward in time, it is possible to estimate the tennis ball's future trajectory on the court, enabling the wheelchair to move to the appropriate position. If the confidence in the predicted trajectory is high, ESTHER acts upon the prediction. Otherwise, the trajectory estimate is ignored until the EKF converges to higher confidence. The EKF and rollout trajectory estimator run at 100 Hz to ensure the high-level planners can incorporate timely estimates.

Our decentralized, modular approach to ball state estimation and trajectory prediction enables us to create a low-cost

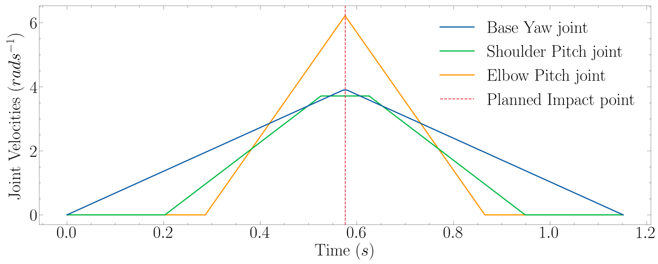

Fig. 6: Trapezoidal velocity profiles for individual joints.

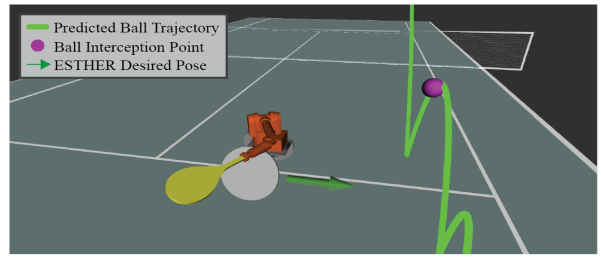

Fig. 7: ESTHER simulation in RViz.

vision system that can be quickly set up on any tennis court. While we can achieve high-speed ball detection at up to 150 Hz, unfortunately, this rate is not enough to infer precisely the spin on the tennis ball (i.e., Magnus effect forces). We attempt to infer this effect by incorporating indirect estimation methods such as Adaptive EKF (AEKF) [39], and trajectory fitting [40] into our EKF.

b) Wheelchair Localization: The motion of ESTHER's wheelchair base can be modeled as a differential drive base that is equipped with three different sensors to determine the robot's state in the world as displayed in Fig. 4. 8192 Counts Per Revolution (CPR) motor encoders provide the velocity and position of each wheel at 250 Hz. A ZED2 Inertial Measurement Unit (IMU) is used to obtain the linear velocity, angular velocity, and orientation of the motion base at 400 Hz. A Velodyne Puck LiDAR sensor is used to obtain an egocentric 360° point cloud with a 30° vertical field of view at 20 Hz. For localizing the wheelchair, we first create a map of the environment offline by recording IMU data and 3D point cloud data while manually driving the wheelchair around at slow speeds. We use the hdl_graph_slam [41] package to create a Point Cloud Data (PCD) map from this recorded data for use online. Afterward, we use a point cloud scan matching algorithm [41] to obtain an odometry estimate based on LIDAR and IMU readings. A differential drive controller gives a second odometry estimate using encoder data from wheels and IMU readings. We follow guidance from [37] for fusing these estimates to get the current state of the wheelchair.

2) Planning: Given the predicted ball trajectory and current position of the wheelchair, the "strategizer", a behavior orchestrator, selects the desired interception point. Then, the strategizer determines the required wheelchair pose, approximate stroke trajectory, and stroke timing. The strategizer locates the point where the ball crosses a hitting plane chosen such that the ball is at a hittable height when passing the robot. The strategizer ensures that the interception point is not too close to the ground and that the corresponding joint positions are within limits.

Next, the strategizer finds the stroke parameters and the corresponding wheelchair placement for that stroke. The stroke is parameterized by three points: a start point (i.e., the joints at the swing's beginning), a contact point (i.e., the joints at ball contact), and an end point (i.e., the joints at the swing's completion). These points are determined such that each joint is traveling at the maximum possible speed at the contact point while staying within each individual joint's position, velocity, and acceleration limits. Using the contact point and the robot's kinematics, the wheelchair's desired position is geometrically determined. Lastly, the strategizer identifies the exact time to trigger the stroke by combining the stroke duration and the time the ball takes to reach the interception point.

The strategizer continuously updates the interception point as the vision system updates its estimate of the ball's predicted trajectory, allowing us to adjust the stroke parameters and the wheelchair position until a few milliseconds before the ball crosses the interception plane, thus increasing the chances of a successful hit.

3) Low-Level Control/Execution: We next describe the interaction between the strategizer and the low-level controllers.

a) Wheelchair Control: Given the robot's position and state in the world, we use ROS's move_base package [34] to command the robot to go to the desired position. We use the move_base's default global planner (i.e., Dijkstra's algorithm) as our global planner and the Timed Elastic Band (TEB) planner [42], [43] as our local planner. The global planner finds a path between the wheelchair's current pose p_{curr} and the desired pose p_{dest} given a map and obstacles detected by the LiDAR. The local planner tries to follow this path as closely as possible while performing real-time collision checking and obstacle avoidance. The plan from the local planner is translated to wheel velocities through a differential drive controller. The velocities are then communicated from the onboard computer to the microcontroller over serial and are executed on the wheelchair. The motor controller uses PID control to command the wheels at desired velocities.

b) Arm Control: The arm-control subsystem deals with generating joint trajectories for the joints and executing them on the actual arm. The Barrett high-speed WAM has 7 joints in total: base yaw, shoulder pitch, shoulder yaw, elbow pitch, wrist pitch, wrist yaw, and palm yaw. The swing utilizes base yaw, shoulder pitch, and elbow pitch joints to generate speed, and the rest of the joints are held at constant positions during the swing so the racket is in the right orientation when making contact with the ball.

To hit a successful return, humans generate high racket head speeds by merging contributions from multiple body joints [44]. Inspired by this "summation of speed principle" [45], we created a "Fully-Extended" Ground stroke (FEG) that maximizes the racket speed when it makes contact with the ball. The base yaw and elbow pitch joint provide the

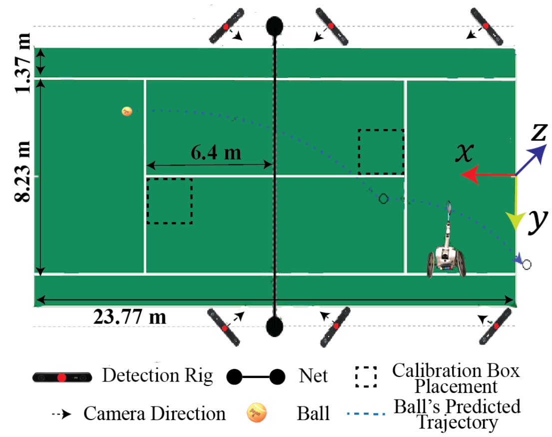

Fig. 8: Schematic showing the court setup

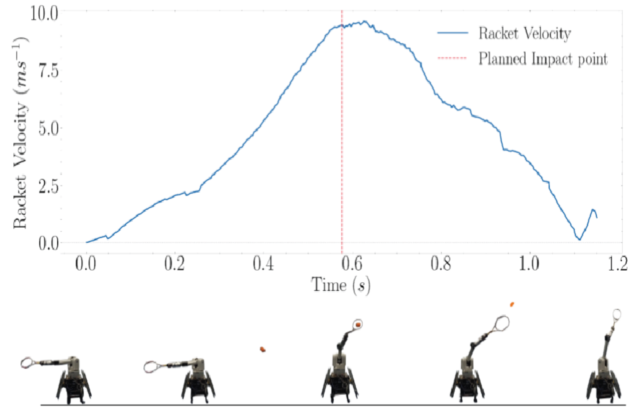

Fig. 9: Racket head speed during a FEG stroke.

velocity component in the direction we want to return the ball, and the shoulder pitch joint is used to adjust for the ball height and hit the ball upward (Fig. 9). Each of these joints follows a trapezoidal velocity profile (Fig. 6). The joints' trajectories are sent from the onboard computer to the WAM computer which uses a PID controller to execute the trajectory on the arm.

C. *Simulation*

Playing tennis involves high-speed motion, which can be hazardous. As such, we set up a kinematic simulator in RViz [46] (Fig. 7) for testing behavior interactions.

IV. Experiments and Results

We evaluate ESTHER and its components to serve as a baseline for future research on athletic robots for tennis. In Section IV-A, we describe the details of our experimental setup. In Section IV-B, we evaluate the capabilities of individual subsystems and the system as a whole. Finally, we report the results of our experiments in Section IV-C.

A. *Experimental Setup*

We have two evaluation settings for our system: one at an NCAA Division 1 indoor tennis court facility and the other within a lab space. The tennis court setup is equal to the size of a regulation tennis court i.e. $23.77\,\mathrm{m}$ by $10.97\,\mathrm{m}$, and the lab setup is $11\,\mathrm{m}$ by $4.57\,\mathrm{m}$. Fig. 8 displays the world coordinate frame and on-court setup. The ball is launched by a ball launcher (Lobster Sports - Elite Two Tennis Ball Launcher) or a human towards the robot from the other side of the court ($11.9\,\mathrm{m} < x < 23.8\,\mathrm{m}$). The lab setup is similar to the court setup except the lab covers a smaller area, has a different surface texture affecting ball bounce and wheel traction, does not have court lines, and has different lighting conditions. Overall, the lab setup is easier for hitting balls.

B. *Sub-System Experiments*

In this section, we overview the capabilities of our current system and the results of experiments conducted in the tennis court and lab based on the setup described in Section IV-A

1) Wheelchair: In manual mode, we safely drove the wheelchair with load at a linear speed of $4.34\,\mathrm{m\,s^{-1}}$ and at an angular speed of $5.8\,\mathrm{rad\,s^{-1}}$, less than half of the maximum possible speeds achievable with our system. In autonomous mode testing with the whole stack, we achieve accelerations of $1.42\,\mathrm{m\,s^{-2}}$ and deceleration of $1.60\,\mathrm{m\,s^{-2}}$. These acceleration and deceleration values are more than the average side-to-side acceleration and deceleration values of $\approx 1.00\,\mathrm{m\,s^{-2}}$ achieved by professional human players [47].

2) Stroke Speed: ESTHER reaches pre-impact racket head speeds (Fig. 9) of $10\,\mathrm{m\,s^{-1}}$. This speed is on the same order of magnitude of professional players ($17 - 36\,\mathrm{m\,s^{-1}}$ [48]).

3) Sensing and State Estimation: An important consideration for playing a sport like tennis is to anticipate the ball trajectory early so the robot can get in position to return the ball. To measure the performance of our ball trajectory prediction system, we measure the error between the ball's predicted x, y, z position and its ground-truth position as it crosses the interception plane as a function of the fraction of time passed to reach the interception plane. We report the average interception point prediction error over 10 trajectories in our court setup in Fig. 10. Each trajectory is roughly $2\,\mathrm{s}$ long, and the ball was launched at roughly $8\,\mathrm{m\,s^{-1}}$. In the court setup, the ball travels along the x-axis, thereby the prediction error in x affects the timing of the stroke. Similarly, prediction error in y affects wheelchair positioning, and z affects the height at which the arm swings. As the FEG stroke takes about $0.5\,\mathrm{s}$ to go from the start to the interception point, we must be certain about the x coordinate of the interception point before 75% of the total trajectory has elapsed. Defining acceptable error margins for racket positions at the interception plane to be equal to the racket head width, we can observe from Fig. 10 that the y and z predictions converge to be within the acceptable error margin after 50% of the trajectory is seen. Given that the wheelchair is able to accelerate at $1.42\,\mathrm{m\,s^{-2}}$, we assess that ESTHER is capable of moving up to $0.71\,\mathrm{m}$ from a standing start after we establish confidence in the predicted interception point. For larger distances, we need more accurate predictions quicker.

4) Reachability: Researchers found that in professional tennis, 80% of all strokes require players to move less than

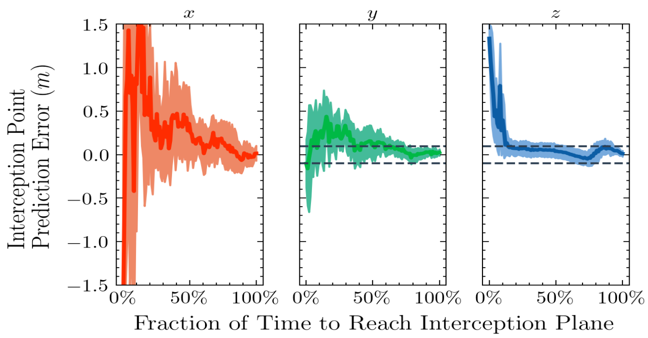

Fig. 10: Error in ball's predicted vs. actual position at desired intercept. Dashed lines represent range of acceptable error.

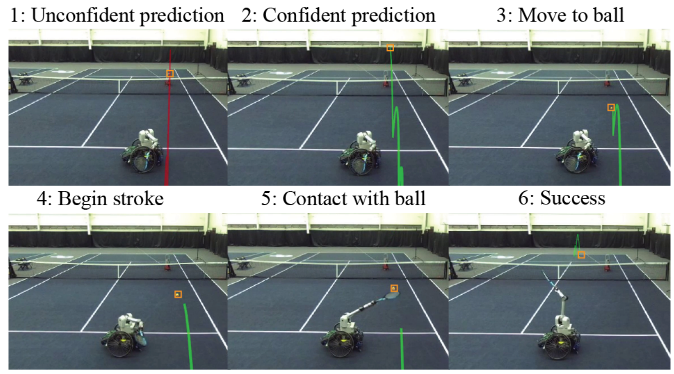

Fig. 11: Outline of ESTHER in action. Orange box highlights ball in image.

$2.5\,\mathrm{m}$ [49]. To benchmark our system's ability to reach balls that are in this range, we conduct a reachability study. We started launching balls toward the robot in our court setup and recorded the results. After every 10 balls, we increased the distance between the wheelchair and the ball's average position at the interception plane by $0.3\,\mathrm{m}$. In this test, to mitigate the latency introduced by the time it requires the vision system to converge to an accurate estimate, we start moving the wheelchair towards the ball's average detected y position as soon as the ball is detected for the first time after launch. The position of the wheelchair is later fine-tuned once the prediction from the vision system has converged. We can see the histogram of the success rate as a function of the distance needed to be moved by the wheelchair in Fig. 12a. As expected, the success rate decreases as the distance needed to be moved by wheelchair increases. However, we still see success even if the wheelchair movement approaches $2\,\mathrm{m}$, demonstrating our system's agility for tennis.

C. Whole-System Experiments

We perform 15 trials in each scenario and report the details of the hit and return rate in Table I. A returned ball over the net landing inside the singles lines of the court is marked as successful. Fig. 11 illustrates a frame-by-frame depiction of our system in action on the court. For tests inside the lab, a ball that goes over the net height ($1.07\,\mathrm{m}$) while crossing the position from which the ball was launched is marked successful.

To further evaluate the variance in ball launches, we conducted a contiguous trial of 50 launches with a ball launcher in the court setting and plotted the histogram of the time it takes to reach the pre-specified interception plane $8\,\mathrm{m}$ away from the launcher in Fig. 12b. In Fig. 12c, we visualize the position of the balls at the interception plane and the configurations that were reachable without wheelchair movement. 65% of the launched balls were in a configuration that required wheelchair movement, showcasing the importance of a mobile system to play a sport like tennis. The experiment was repeated in the lab setting as well. The results are marginally better in the lab setting due to the closer proximity of the cameras and more controlled conditions. These results along with the data in Table I serve as a good baseline for future works as beating them would require faster perception, increased agility, and improved planning and control algorithms.

V. Discussion and Key Challenges

We have experimentally demonstrated that ESTHER is fast, agile, and powerful enough to successfully hit ground strokes on a regulation tennis court. Our vision for the system is to be able to rally and play with a human player. To realize that goal, we lay down a set of key challenges to address.

Section IV-B.4 showcased that ESTHER is able to successfully return balls that require the wheelchair to move up to $2\,\mathrm{m}$ given that the movement starts right after the ball is launched. However, as discussed in Section IV-B.3, our vision system needs some time to converge to an accurate prediction. To get an accurate, early estimate of the ball's pose, we will need a higher fps vision setup that utilizes cameras to triangulate a ball from multiple angles [40]. To make the ball's predicted trajectory more precise, the system should also be able to estimate the spin of the ball. Directly estimating spin on the ball is a hard problem [40]. However, the ball's spin can be estimated indirectly using methods such as trajectory fitting [40] or an AEKF [39]. An interesting research direction would be to accurately estimate the ball's state [50] by observing how the ball was hit on the previous shot for e.g., a racket going from low to high indicates top spin, and a racket going from high to low is indicative of backspin. The ultimate challenge is to perform this sensing with only onboard mechanisms.

We currently plan the motion of the wheelchair and the arm separately. However, the two components are dynamically coupled, and an ideal planner should consider both simultaneously. In future work, our system serves as an ideal test-bed for the development of kinodynamic motion planners for athletic mobile manipulators as current dynamical planners are not good enough for agile movement [51]. Such a planner would also help us reach higher racket head velocities as we will be able to utilize the wheelchair's potential to rotate at high angular speeds as established in the same way that professional human players rotate their trunk while hitting a ground stroke. Moreover, the planner will need to be safe and fast for live human-robot play.

Setup	Ball Launch		Interception Point Lateral (y) spread		Success Metrics	
	Distance to Interception Plane (m)	Average Launch Speed ($\mathrm{m\,s^{-1}}$)	IQR (m)	Std. Dev (m)	Hit Rate	Success Rate
Court (Ball Launcher)	7.9	8.01	0.31	0.23	73%	66%
Court (Ball Launcher)	12.8	12.64	0.26	0.29	60%	53%
Lab (Ball Launcher)	7.5	6.79	0.28	0.20	93%	80%
Lab (Human)	7.5	6.56	0.54	0.52	40%	33%

TABLE I: Experiment results for 15 consecutive trials in different scenarios

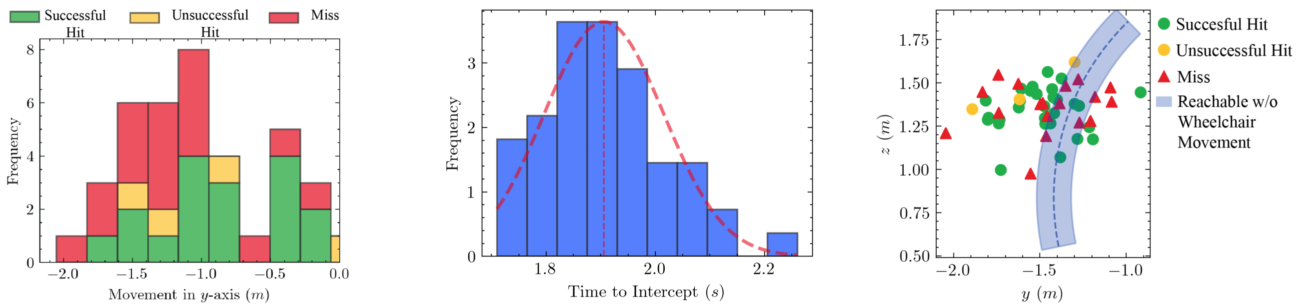

(a) (b) (c)

Fig. 12: Fig. 12a shows a histogram of successful returns as a function of the distance the wheelchair needed to move to hit the ball. Fig. 12b shows a histogram for times of ball arrival at interception point from the time of the first detection with the court setup. Fig. 12c shows a scatter plot of ball positions in the interception plane with the court setup.

VI. Applications and Future Work

The system opens up numerous exciting opportunities for future research such as learning strokes from expert demonstrations via imitation & reinforcement learning, incorporating game knowledge to return at strategic points, and effective human-robot teaming for playing doubles. We briefly describe our vision for some of these exciting future directions:

Interactive Robot Learning – Interactive Robot Learning (i.e., Learning from Demonstration (LfD)) focuses on extracting skills from expert demonstrations or instructions for how to perform a task. It is quite challenging to manually encode various strategies, tactics, and low-level stroke styles for a tennis-playing robot, and, ESTHER serves as a great platform to develop and deploy LfD algorithms that learn from expert gameplay. Prior work has shown how kinesthetically teaching can be employed to teach robots various styles of playing Table-Tennis [52], [53], but these techniques are demonstrated on arms that are mounted to a stationary base and require much lower racket head speeds. Therefore, we challenge the LfD community to leverage our platform to explore a more challenging robotics domain.

Human-Robot Teaming – Having agile robots that collaborate with humans to perform well on a shared task is a challenging problem as it requires an accurate perception of human intention, clear communication of robot intention, and collaborative planning and execution. We believe that our system is a great benchmark to study these problems as doubles play in tennis often has many collaborative strategies that are based on a good rapport that players develop over many practice sessions. Having robots similarly collaborate with humans and build trust and personalized strategies over time in a highly competitive and dynamic environment is an important research problem more broadly [54].

VII. Conclusion

In this paper, we present ESTHER, a fully autonomous system built with general-purpose robotic hardware. Our work contributes to the existing research in three main ways: introduction of a low-cost, fast, decentralized perception system that can be easily set up anywhere for accurate ball tracking; design of a chain drive system that can be used to motorize a regulation sports wheelchair into an agile mobile base for any robotic manipulator; and planning and control of an agile mobile manipulator to exhibit athletic behaviors. We experimentally determined the capabilities of individual subsystems and verified that the system is able to perform at a human scale. We also demonstrated that our system can successfully hit ground strokes on a regulation tennis court. Avenues of future research include improvements to the perception system, development of kinodynamic planners for athletic mobile manipulators, learning game strategies, court positioning, and strokes from human play, as well as collaboration with a human partner to successfully play doubles tennis.

References

[1] M. Taylor, "Editorial – sport, transnationalism, and global history," *Journal of Global History*, vol. 8, no. 2, p. 199–208, 2013.

[2] Ørnulf Seippel, "Do sports matter to people? A cross-national multilevel study," *Sport in Society*, vol. 22, no. 3, pp. 327–341, 2019.

[3] C. Malm, J. Jakobsson, and A. Isaksson, "Physical activity and sports—real health benefits: a review with insight into the public health of sweden," *Sports*, vol. 7, no. 5, p. 127, 2019.

[4] R. M. Eime, J. A. Young, J. T. Harvey, M. J. Charity, and W. R. Payne, "A systematic review of the psychological and social benefits of participation in sport for children and adolescents: informing development of a conceptual model of health through sport," *International Journal ofBehavioral Nutrition and Phy. Activity*, vol. 10, no. 1, p. 98, 2013.

[5] C. Malm, J. Jakobsson, and A. Isaksson, "Physical activity and Sports-Real health benefits: A review with insight into the public health of sweden," *Sports (Basel)*, vol. 7, no. 5, May 2019.

[6] D. J. Reinkensmeyer and J. L. Patton, "Can robots help the learning of skilled actions?" *Exerc Sport Sci Rev*, vol. 37, no. 1, pp. 43–51, 2009.

[7] W. S. Erdmann, "Problems of sport biomechanics and robotics," *International Journal of Advanced Robotic Sys.*, vol. 10, no. 2, p. 123, 2013.

[8] D. Calkins, "An overview of robogames [competitions]," *IEEE robotics & automation magazine*, vol. 18, no. 1, pp. 14–15, 2011.

[9] H. Kitano, M. Asada, Y. Kuniyoshi, I. Noda, E. Osawai, and H. Matsubara, "RoboCup: A challenge problem for AI and robotics," in *RoboCup-97: Robot Soccer World Cup I*, J. G. Carbonell, J. Siekmann, G. Goos, J. Hartmanis, J. van Leeuwen, and H. Kitano, Eds. Berlin, Heidelberg: Springer, 1998, vol. 1395, pp. 1–19, lecture Notes in CS.

[10] S. Cass, "Robosoccer [mobile robotics experiment]," *IEEE spectrum*, vol. 38, no. 5, pp. 75–77, 2001.

[11] R. Balogh, "I am a robot – competitor: A survey of robotic competitions," *International Journal of Advanced Robotic Sys.*, vol. 2, no. 2, p. 17, 2005.

[12] K. Muelling, J. Kober, and J. Peters, "Learning table tennis with a mixture of motor primitives," in *Proc. Humanoids*, 2010, pp. 411–416.

[13] M. Hattori, K. Kojima, S. Noda, F. Sugai, Y. Kakiuchi, K. Okada, and M. Inaba, "Fast tennis swing motion by ball trajectory prediction and joint trajectory modification in standalone humanoid robot real-time system," in *Proc. IROS*, 2020, pp. 3612–3619.

[14] R. Terasawa, S. Noda, K. Kojima, R. Koyama, F. Sugai, S. Nozawa, Y. Kakiuchi, K. Okada, and M. Inaba, "Achievement of dynamic tennis swing motion by offline motion planning and online trajectory modification based on optimization with a humanoid robot," in *Proc. Humanoids*, 2016, pp. 1094–1100.

[15] D. Büchler, S. Guist, R. Calandra, V. Berenz, B. Schölkopf, and J. Peters, "Learning to play table tennis from scratch using muscular robots," *IEEE Transactions on Robotics*, 2022.

[16] 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.

[17] S. Mori, K. Tanaka, S. Nishikawa, R. Niiyama, and Y. Kuniyoshi, "High-speed humanoid robot arm for badminton using pneumatic-electric hybrid actuators," *IEEE RA-L*, vol. 4, no. 4, pp. 3601–3608, 2019.

[18] P. Janssens, G. Pipeleers, M. Diehl, G. Pinte, and J. Swevers, "Energy-optimal time allocation of a series of point-to-point motions: Demonstration on a badminton robot," in *International Workshop on Research and Education in Mechatronics*, 2012, pp. 27–31.

[19] M. Liu, B. Depraetere, G. Pinte, I. Grondman, and R. Babuška, "Model-free and model-based time-optimal control of a badminton robot," in *Proc. ASCC*, 2013, pp. 1–6.

[20] W. Gao, L. Graesser, K. Choromanski, X. Song, N. Lazic, P. Sanketi, V. Sindhwani, and N. Jaitly, "Robotic table tennis with model-free reinforcement learning," in *Proc. IROS*, 2020, pp. 5556–5563.

[21] W. Zhang, J. Li, Q. Huang, Z. Yu, X. Chen, G. Ma, L. Meng, Y. Liu, S. Zhang, F. Meng, W. Zhang, and J. Gao, "System design of a 9-dof robot capable of fast and flexible rally task," in *2014 IEEE International Conference on Mechatronics and Automation*, 2014, pp. 1428–1433.

[22] F. Miyazaki, M. Matsushima, and M. Takeuchi, *Learning to Dynamically Manipulate: A Table Tennis Robot Controls a Ball and Rallies with a Human Being*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006, pp. 317–341. [Online]. Available: https://doi.org/10.1007/978-3-540-37347-6_15

[23] Y. Ji, X. Hu, Y. Chen, Y. Mao, G. Wang, Q. Li, and J. Zhang, "Model-based trajectory prediction and hitting velocity control for a new table tennis robot," in *Proc. IROS*, 2021, pp. 2728–2734.

[24] K. Dong, K. Pereida, F. Shkurti, and A. P. Schoellig, "Catch the Ball: Accurate High-Speed Motions for Mobile Manipulators via Inverse Dynamics Learning," in *Proc. IROS*, 2020.

[25] F. Yang, Z. Shi, S. Ye, J. Qian, W. Wang, and D. Xuan, "Varsm: Versatile autonomous racquet sports machine," in *Proc. ICCPS*, 2022, pp. 203–214.

[26] Z. Li, P. Moran, Q. Dong, R. J. Shaw, and K. Hauser, "Development of a tele-nursing mobile manipulator for remote care-giving in quarantine areas," in *Proc. ICRA*, 2017, pp. 3581–3586.

[27] S. Caselli, E. Fantini, F. Monica, P. Occhi, and M. Reggiani, "Toward a mobile manipulator service robot for human assistance," in *Proc. 1st Robocare Workshop*, 2003.

[28] C.-H. King, T. L. Chen, Z. Fan, J. D. Glass, and C. C. Kemp, "Dusty: an assistive mobile manipulator that retrieves dropped objects for people with motor impairments," *Disabil Rehabil Assist Technol*, vol. 7, no. 2, pp. 168–179, Oct. 2011.

[29] M. Hvilshøj, S. Bøgh, O. Rosenlund, and O. Madsen, "Autonomous industrial mobile manipulation (aimm): Past, present and future," *Industrial Robot*, vol. 39, pp. 120 – 135, 03 2012.

[30] H. Engemann, S. Du, S. Kallweit, P. Cönen, and H. Dawar, "Omnivil—an autonomous mobile manipulator for flexible production," *Sensors*, vol. 20, no. 24, 2020.

[31] R. Bostelman, T. Hong, and J. Marvel, "Survey of research for performance measurement of mobile manipulators," *Journal of Research of the NIST*, vol. 121, p. 342, 06 2016.

[32] H. Moradi, K. Kawamura, E. Prassler, G. Muscato, P. Fiorini, T. Sato, and R. Rusu, "Service robotics (the rise and bloom of service robots)," *IEEE Robotics & Automation Magazine*, vol. 20, pp. 22–24, 09 2013.

[33] Stanford Artificial Intelligence Laboratory et al. Robotic operating system. [Online]. Available: https://www.ros.org

[34] A. Koubaa, *Robot Operating System (ROS): The Complete Reference (Volume 2)*, 1st ed. Springer Publishing Company, Incorporated, 2017.

[35] J. Wang and E. Olson, "AprilTag 2: Efficient and robust fiducial detection," in *Proc. IROS*. IEEE, oct 2016, pp. 4193–4198.

[36] H. Sorenson, *Kalman Filtering: Theory and Application*, ser. IEEE Press selected reprint series. IEEE Press, 1985.

[37] T. Moore and D. Stouch, "A generalized extended kalman filter implementation for the robot operating system," in *Proc. IAS*, 2014.

[38] R. Mehta, F. Alam, and A. Subic, "Review of tennis ball aerodynamics," *Sports Technology*, vol. 1, no. 1, pp. 7–16, 2008.

[39] L. Jetto, S. Longhi, and G. Venturini, "Development and experimental validation of an adaptive extended kalman filter for the localization of mobile robots," *IEEE Transactions on Robotics and Automation*, vol. 15, no. 2, pp. 219–229, 1999.

[40] J. Tebbe, L. Klamt, Y. Gao, and A. Zell, "Spin detection in robotic table tennis," in *Proc. ICRA*, 2020, pp. 9694–9700.

[41] K. Koide, J. Miura, and E. Menegatti, "A portable three-dimensional lidar-based system for long-term and wide-area people behavior measurement," *International Journal of Advanced Robotic Sys.*, vol. 16, no. 2, 2019.

[42] C. Rösmann, F. Hoffmann, and T. Bertram, "Integrated online trajectory planning and optimization in distinctive topologies," *Robotics and Autonomous Systems*, vol. 88, pp. 142–153, 2017.

[43] C. Rösmann, W. Feiten, T. Wösch, F. Hoffmann, and T. Bertram, "Trajectory modification considering dynamic constraints of autonomous robots," in *ROBOTIK*, 2012.

[44] B. Pedro, F. João, J. P. R. Lara, S. Cabral, J. Carvalho, and A. P. Veloso, "Evaluation of upper limb joint contribution to racket head speed in elite tennis players using imu sensors: Comparison between the cross-court and inside-out attacking forehand drive," *Sensors*, vol. 22, no. 3, 2022.

[45] J. Landlinger, S. Lindinger, T. Stöggl, H. Wagner, and E. Müller, "Key factors and timing patterns in the tennis forehand of different skill levels," *J Sports Sci Med*, vol. 9, no. 4, pp. 643–651, Dec. 2010.

[46] H. R. Kam, S.-H. Lee, T. Park, and C.-H. Kim, "Rviz: A toolkit for real domain data visualization," *Telecommun. Syst.*, vol. 60, no. 2, p. 337–345, oct 2015.

[47] A. Filipčič, B. Leskošek, G. Munivrana, G. Ochiana, and T. Filipčič, "Differences in movement speed before and after a Split-Step between professional and junior tennis players," *Journal ofHum. Kinet.*, vol. 55, pp. 117–125, Jan. 2017.

[48] T. Allen, S. Choppin, and D. Knudson, "A review of tennis racket performance parameters," *Sports Eng.*, vol. 19, no. 1, pp. 1–11, 2016.

[49] M. S. Kovacs, "Movement for tennis: The importance of lateral training," *Strength & Conditioning Journal*, vol. 31, no. 4, 2009.

[50] E. Wu and H. Koike, "Futurepong: Real-time table tennis trajectory forecasting using pose prediction network," in *CHI, Extended Abstracts*. New York, NY, USA: ACM, 2020, p. 1–8.

[51] J. Maria Porta Pleite and L. Ros Giralt. (2021, Sep) Kinodyn+: Synthesis of optimally agile and graceful robot motions. [Online]. Available: https://www.iri.upc.edu/project/show/271

[52] L. Chen, R. Paleja, M. Ghuy, and M. Gombolay, "Joint goal and strategy inference across heterogeneous demonstrators via reward network distillation," in *Proc. HRI*, 2020, pp. 659–668.

[53] S. Gomez-Gonzalez, G. Neumann, B. Schölkopf, and J. Peters, "Using probabilistic movement primitives for striking movements," in *Proc. Humanoids*, 2016, pp. 502–508.

[54] A. Bauer, D. Wollherr, and M. Buss, "Human–robot collaboration: A survey," *International Journal ofHumanoid Robotics*, vol. 05, no. 01, pp. 47–66, 2008.