Catch the Ball: Accurate High-Speed Motions for Mobile Manipulators via Inverse Dynamics Learning

Ke Dong¹, Karime Pereida¹, Florian Shkurti² and Angela P. Schoellig¹

Abstract-Mobile manipulators consist of a mobile platform equipped with one or more robot arms and are of interest for a wide array of challenging tasks because of their extended workspace and dexterity. Typically, mobile manipulators are deployed in slow-motion collaborative robot scenarios. In this paper, we consider scenarios where accurate high-speed motions are required. We introduce a framework for this regime of tasks including two main components: (i) a bi-level motion optimization algorithm for real-time trajectory generation, which relies on Sequential Quadratic Programming (SOP) and Quadratic Programming (QP), respectively; and (ii) a learningbased controller optimized for precise tracking of high-speed motions via a learned inverse dynamics model. We evaluate our framework with a mobile manipulator platform through numerous high-speed ball catching experiments, where we show a success rate of 85.33%. To the best of our knowledge, this success rate exceeds the reported performance of existing related systems [1], [2] and sets a new state of the art.

I. INTRODUCTION

Mobile manipulators combine mobile platforms and robot arms. Unlike fixed-base industrial arms, mobile manipulators have a larger workspace and can deal with more flexible and complex tasks such as housekeeping, construction, and airplane assembly [3]. Current research on mobile manipulators mainly focuses on techniques in slow-motion scenarios, such as kinematics redundancy resolution [4], optimal path planning [3] and control [5]. However, mobile manipulators are capable of high-speed motions, and thus have the potential to be used in applications such as industrial highspeed pick-and-place, or other challenging tasks such as table tennis [6]. In this paper we focus on high-speed trajectory generation and accurate tracking for mobile manipulators, with an emphasis on not compromising accuracy at higher speeds.

We focus on a benchmark task for accurate high-speed motions: object catching, specifically, ball catching. Within a fraction of a second, the robot needs to (*i*) predict the ball's trajectory, (*ii*) generate and update a trajectory for graceful catching, and (*iii*) accurately track it with the mobile manipulator. There are numerous papers [1], [7], [8], [9] that address the ball catching problem with a *fixed-base* manipulator. However, using a *mobile* manipulator introduces additional challenges, due to the interplay between the following three characteristics:



Fig. 1: Ball catching experimental setup. On average, the ball is thrown by a human operator from a distance of about 4 m towards the robot with a speed of typically 6 m/s, resulting in a flight time of about 1 s. A video of the experiment results is available at http://tiny.cc/ball_catch.

Controlling wheel-driven bases: The mobile bases that are driven by Mecanum wheels [10], as used in this paper, usually have lower position accuracy than manipulators. One example is the position accuracy of the Kuka KMR Quantec driven by Mecanum wheels of $\pm 5 mm$ [11] while the accuracy of the UR10 arm is $\pm 0.05 mm$ [12]. Thus extra efforts are needed to accurately control the base.

High precision in position and time: Precise end effector trajectory tracking, both spatially and temporally and in position and orientation, is a pre-requisite for successful catching. A precision of 16 *mm* in position and 3 *ms* in time must be achieved in this paper. These parameters vary depending on the gripper [2].

Real-time computation: On average, the duration of the ball's free flight trajectory is one second, which puts a hard constraint on the computation time of ball prediction, motion planning and control. Motion planning is very challenging since complex constraints such as the nonlinear robot dynamics and forward kinematics must be considered.

Given the three challenges above, accurate high-speed motion generation for mobile manipulator ball catching has not been widely explored, with a few very notable exceptions [2]. In this paper we introduce: (i) a bi-level motion planning scheme that generates smooth trajectories despite nonlinear constraints in real time, and (ii) a learning-based controller improving tracking accuracy and enabling the mobile manipulator to follow trajectories sufficiently accurately. Our proposed bi-level optimization scheme for motion planning uses sequential quadratic programming (SQP) to select the catching pose, and solves a low-level quadratic program

¹ Ke Dong, Karime Pereida and Angela P. Schoellig are with the Dynamic Systems Lab (www.dynsyslab.org) at the University of Toronto Institute for Aerospace Studies (UTIAS), and the Vector Institute for artificial Intelligence, Canada. Email: ke.dong, karime.pereida@robotics.utias.utoronto.ca, schoellig@utias.utoronto.ca

² Florian Shkurti is with Department of Computer Science at the University of Toronto, Canada. Email: florian@cs.toronto.edu

(QP) that generates the trajectory to the catching pose. Our proposed learning-based controller learns the inverse dynamics of the system [13] to improve the control accuracy of the mobile manipulator. We use a modified Kalman Filter for improved ball state estimation, which incorporates the influence of aerodynamic drag in the state prediction step.

We present a thorough analysis and evaluation of major components of the system. Experiment data shows that our bi-level optimization scheme has an average run time of 5.1 ms, and the inverse dynamics learning method can reduce tracking error up to 79.1%. A success rate of 85.33% of real robot ball catching task is achieved, which is higher than the 80% success rate reported in [1], [2]. A video with experiments showing accurate high-speed motions can be found at http://tiny.cc/ball_catch.

The contributions of this paper are two-fold. First, we propose an algorithmic framework to generate and accurately follow high-speed trajectories with a mobile manipulator. The proposed framework has a bi-level optimization scheme for motion planning, which can generate smooth joint trajectories under complex constraints in real-time. Additionally, the proposed framework has a learning-based controller to improve the tracking performance of the mobile manipulator. Our second contribution is the extensive evaluation of the performance of the proposed framework on the mobile manipulator, which includes a thorough error analysis of major components. This framework is tested on real robot catching experiments and achieves a success rate of 85.33%. This success requires a tight interplay between different system components and hardware constraints. Implementation details of such a complex system are presented, and should be a resource for similar system designs.

II. RELATED WORK

A large body of work has been devoted to ball catching with a fixed-base manipulator [1], [7], [8], [9], with special focus on motion planning. Early pioneering work in this field used heuristic methods for selecting the catching point and generated the joint trajectory via interpolation [7]. Such methods are simple, as they restrict catching points to a certain area, but may not find the optimal robot motion for catching. Recent work uses nonlinear optimization methods for selecting a catching pose, and parameterization methods, like trapezoidal velocity profile and third-order polynomials, for path generation for simplicity. In [1], a ball catching system consisting of a fixed-base 7 Degree-of-Freedom (DoF) arm and a four-finger hand is presented. It uses SQP for high-level goal selection, and restricts joint trajectories to trapezoidal velocity profiles. In contrast, in our framework, we do not parameterize the trajectory. Instead, we use QP to generate intermediate waypoints which define smooth and flexible trajectories. In the realm of high-speed manipulation, the work in [8] proposes a control law based on dynamical systems that enables the system to catch fast-moving objects. However, only fixed-base manipulators are used.

The work most similar to ours is the mobile humanoid system presented in [2], which consists of two light-weight 7

DoF arms and an omnidirectional mobile base. In their work, the mobile base is restricted to motion in the *x* axis only. Further, they use customized hardware capable of high-speed motions. In contrast, in our work, we achieve high-speed motions in the x - y plane with the mobile base, which brings the motion capacity of omnidirectional platforms into full play. Moreover, we use an off-the-shelf mobile manipulator and a learning-based controller for improved tracking accuracy, which makes our approach easy to generalize to other robotic platforms without the need for customized hardware and low-level access.

There are numerous works that try to improve the tracking performance of high-speed and aggressive trajectories. Non-learning-based methods [5], [14] rely on accurate analytical system dynamics models, which are hard to obtain. Learning-based methods mainly focus on learning the system's dynamics or inverse dynamics models [15], [16]. In [16], a Deep Neural Network (DNN) is used for direct inverse control of a quadrotor. However, guaranteeing stability with a DNN in the control loop can be very difficult. In contrast, the inverse dynamics learning method from [13] uses a DNN as an add-on module outside the original closed-loop system, which is safer in terms of stability.. This method's ability for improved, impromptu trajectory tracking has been demonstrated on quadrotors [13], [17], but not on mobile manipulators.

III. SYSTEM OVERVIEW AND PROBLEM STATEMENT

This section presents a system overview of the mobile manipulator and the flying ball. A formal statement of the ball catching problem is also given.

A. System Overview

The mobile manipulator system considered in this paper consists of a 6-DoF manipulator rigidly mounted on a 3-DoF omnidirectional mobile base. The yaw rotation of the mobile base is not used in this work since it is too slow for the application at hand. Thus the system has 8 DoF in total. We refer to these DoF as robot joints, for brevity. The joint configuration vector of the robot is $\mathbf{q} = [\mathbf{q}_a^T, \mathbf{q}_b^T]^\top$, where $\mathbf{q}_a = [\theta_1, \theta_2, ..., \theta_6]^\top$ is a vector containing the arm joint angles and $\mathbf{q}_b = [x_b, y_b]^\top$ represents the Cartesian position of the mobile base. The forward kinematics equations of the mobile manipulator shown in Fig. 2 are assumed to be known [18].

The flying ball's motion is modeled as a point mass under the influence of aerodynamic drag and gravity. The aerodynamic drag is proportional to the square of the speed [19]. The ball's equation of motion can be written as follows:

$$\dot{\mathbf{b}} = -\mathbf{g} - K_D ||\dot{\mathbf{b}}||\dot{\mathbf{b}}, \qquad (1)$$

where $\mathbf{b} = [b_x, b_y, b_z]^{\top}$ is the ball's Cartesian position in the world frame and $\dot{\mathbf{b}}$, $\ddot{\mathbf{b}}$ are the corresponding velocity and acceleration, $\mathbf{g} = [0, 0, g]^{\top}$ ($g = 9.81 \text{ m/s}^2$) is the gravitational acceleration, K_D is the aerodynamic drag coefficient, and $|| \cdot ||$ is the Euclidean norm.



Fig. 2: The mobile manipulator with the world frame, ${}_{w}F$, base frame ${}_{b}F$, arm frame ${}_{a}F$ and end effector frame ${}_{ee}F$ is shown. Rotation axes of the six arm joints are also presented. Adapted from [20].

B. Problem Statement

Given initial robot configuration \mathbf{q}_0 , ball state $[\mathbf{b}_0^T, \mathbf{b}_0^T]^T$, and time t_0 , the goal is to catch the ball with the robot's end effector at a certain catch time $t_f > t_0$. To achieve this, two conditions need to be satisfied: 1) at catch time t_f , the origin of the end effector frame should coincide with the ball's position; and 2) the z axis of the end effector frame should be aligned with the ball's velocity vector. Formally, these constraints can be written as:

$${}_{w}\mathbf{P}_{ee}(\mathbf{q}_{f}) = \mathbf{b}(t_{f})$$

$$\cos(\langle_{w} \mathbf{z}_{ee}(\mathbf{q}_{f}), -\dot{\mathbf{b}}(t_{f}) \rangle) = 1,$$
(2)

where \mathbf{q}_f is the robot joint configuration at catch time t_f , ${}_w \mathbf{P}_{ee}$ is the position of the end effector frame's origin in the world frame, ${}_w \mathbf{z}_{ee}$ is the vector of the *z* axis of the end effector in the world frame, and $\langle \cdot, \cdot \rangle$ refers to the angle between two vectors.

IV. METHODOLOGY

This section presents the framework we propose, which includes three main components: (*i*) ball estimation and prediction, (*ii*) bi-level motion planning, and (*iii*) low-level control with inverse dynamics learning, as shown in Fig. 3.

A. Ball Estimation and Prediction

The first component estimates the ball state and uses it to generate trajectory predictions within a finite horizon. To estimate the ball state $\mathbf{s}_b = [\mathbf{b}^T, \dot{\mathbf{b}}^T]^T$, we use a modified discrete-time Kalman filter that considers the influence of air drag in the state prediction step, as presented in [19]. In the state prediction step at the discrete time index k, we predict the ball state $\mathbf{s}_b[k+1|k]$ at time step k+1 given measurements up to time k according to the estimated state at the previous time step $\mathbf{s}_b[k|k]$ and motion model in (1) as follows:

$$\dot{\mathbf{b}}[k|k] = -\mathbf{g} - K_D ||\dot{\mathbf{b}}[k|k]||\dot{\mathbf{b}}[k|k]$$

$$\dot{\mathbf{b}}[k+1|k] = \dot{\mathbf{b}}[k|k] + \delta_k \dot{\mathbf{b}}[k|k]$$

$$\mathbf{b}[k+1|k] = \mathbf{b}[k|k] + \delta_k \dot{\mathbf{b}}[k|k] + \frac{1}{2} \delta_k^2 \ddot{\mathbf{b}}[k|k],$$
(3)

where δ_k is the time interval between time index k+1 and k. The update step and variance propagation and update are the same as the vanilla Kalman filter [21]. The aerodynamic drag coefficient K_D is estimated via a Recursive Least-square estimator [19] in an off-line fashion.

For the trajectory prediction, given the estimated ball state \hat{s}_b at time t_0 , future ball trajectory predictions are made



Fig. 3: Diagram of the overall framework for accurate high-speed ball catching. Note that the estimation module shown here is for ball state estimation, with separate, but not shown estimation modules for arm and base controllers.

via numerically integrating (1) for a finite horizon. Hence, ball trajectory predictions will change when a new state measurement comes, and a new motion plan should be made. This requires that the motion planner can make new plans in a very short time (typically 20 *ms* [9]) so that the information of the latest ball trajectory prediction can be used in the latest robot motion plan.

B. Motion Planning

In order to achieve better performance, the mobile platform trajectory needs to be updated when a new ball trajectory prediction is available. The motion planner is bi-level: 1) the *high-level* planner calculates a feasible and optimal catching configuration \mathbf{q}_f and catch time t_f that enable the end effector to intercept the ball's trajectory at time t_f , and 2) the *low-level* planner generates a smooth joint trajectory $\mathbf{q}_f(t)$ that takes the robot from its current to the desired catch configuration. Moreover, the motion planner must take into account kinematic redundancy, collision avoidance, and robot kinematic and dynamic constraints (e.g. position, velocity and torque constraints). This problem is a nonlinear and non-convex optimization problem [9]. Solving such an optimization problem even off-line is still an open question in the field [1], [9].

To tackle this planning problem in an on-line fashion, we need to make the following three simplifications. Similar simplifications are also made in the ball catching literature and work well in experiments [1], [2], [9], [22].

Simplification 1 (Arm-base collision): There are no nearby obstacles and no arm self-collision check; hence, no obstacle or self-collision check is needed. The only collision that may happen is the collision between the arm end effector and the mobile base.

This simplification is reasonable because during robot catching experiments, usually the arm will extend to reach out for a catching. In such a scenario, arm self-collision is rare but collisions between the arm and mobile base are possible. Specifically, to prevent such a collision, we require the arm end effector to stay inside a semi-cylinder centered on the arm's base:

$${}_{a}P^{2}_{ee,x}(\mathbf{q}) + {}_{a}P^{2}_{ee,y}(\mathbf{q}) \le R^{2}_{coll}$$

$$0 \le {}_{a}P_{ee,z}(\mathbf{q}) \le H_{coll},$$
(4)

where $[_{a}P_{ee,x}, _{a}P_{ee,y}, _{a}P_{ee,z}]$ is the end effector's position in the arm frame which can be calculated via forward kinematics

functions given the robot configuration \mathbf{q} , R_{coll} is the cylinder's radius and H_{coll} is the cylinder's height.

Through experiments, we noticed that as long as the end effector is inside this semi-cylinder space at catch time, the whole generated end effector trajectory usually will stay in that space. Thus constraints in (4) is only checked for \mathbf{q}_f instead of every $\mathbf{q}(t)$ for $t_0 \le t \le t_f$.

Simplification 2 (Kinematic planning): We use kinematics models to depict the joint motion, specify objective functions and constraints on the joint accelerations, instead of torques.

This simplification works in experiments as long as maximum joint acceleration values are chosen conservatively and low-level joint controls work properly. We use a double integrator as the kinematics model of a single joint and take joint acceleration as the control input:

$$\begin{bmatrix} q_i[k+1] \\ \dot{q}_i[k+1] \end{bmatrix} = \begin{bmatrix} 1 & \gamma \\ 0 & 1 \end{bmatrix} \begin{bmatrix} q_i[k] \\ \dot{q}_i[k] \end{bmatrix} + \begin{bmatrix} \frac{1}{2}\gamma^2 \\ \gamma \end{bmatrix} u_i[k], \qquad (5)$$

where index *i* represents the *i*-th joint, γ is the step size¹, and $u_i[k]$ is the control input. To ensure the trajectory feasibility, we have the following box constraints:

$$\begin{aligned} \mathbf{q}_{min} &\leq \mathbf{q}[k] \leq \mathbf{q}_{max}, \\ \dot{\mathbf{q}}_{min} &\leq \dot{\mathbf{q}}[k] \leq \dot{\mathbf{q}}_{max}, \\ \ddot{\mathbf{q}}_{min} &\leq \mathbf{u}[k] \leq \ddot{\mathbf{q}}_{max}, \end{aligned} \tag{6}$$

where $\mathbf{q}_{min}, \dot{\mathbf{q}}_{min}, \ddot{\mathbf{q}}_{max}, \dot{\mathbf{q}}_{max}, \ddot{\mathbf{q}}_{max}$ are minimum and maximum value of joint positions, velocities and accelerations, respectively.

Simplification 3 (Quadratic cost function): The cost functions in the high- and low-level planners are quadratic. We use quadratic cost functions because they are popular in the optimization field and usually easy to deal with.

With these three simplifications, we solve the motion planning problem in a bi-level fashion. The high-level planner selects the final catching configuration \mathbf{q}_f and catch time t_f , while including the nonlinear non-convex constraints in (2) and (4). Given the solution \mathbf{q}_f^*, t_f^* , the low-level path planner determines the intermediate trajectory waypoints $\mathbf{q}(t_k), t_0 < t_k < t_f$, but is restricted only by the linear constraints in (5) and (6). This is a quadratic program that can be solved very efficiently. To ensure that the optimal values (\mathbf{q}_f^*, t_f^*) calculated by the high-level planner are feasible for the low-level planning problem, a box constraint on \mathbf{q}_f which depends on t_f and initial robot conditions is defined.

Compared to the parameterization methods used in literature for the low-level planning problem, like trapezoidal velocity profiles [1], [2] and third-order polynomials [6], the QP technique we use allows more flexible trajectories due to the increased number of decision variables. Furthermore, via constraints of kinematic models in the quadratic program formulation, the generated trajectories are smoother and should be closer to the robot's actual behaviors when compared to trapezoidal velocity profiles and third-order polynomials. Details on the high- and low-level planners follow.

1) High-level goal planner: Given the initial robot configuration and velocity $(\mathbf{q}_0, \dot{\mathbf{q}}_0)$, ball position and velocity trajectory prediction $(\mathbf{b}(t), \dot{\mathbf{b}}(t))$ and initial time t_0 , the highlevel planner aims to find the optimal final catching configuration and catch time (\mathbf{q}_f^*, t_f^*) . Formally, this can be written as the following optimization problem:

$$\begin{aligned} \mathbf{q}_{f}^{*}, t_{f}^{*}) &= & \operatorname*{arg\,min}_{\mathbf{q}_{f}, t_{f}} \Gamma_{h}(\mathbf{q}_{f}, t_{f}) \\ \text{s.t.} & aP_{ee,x}^{2}(\mathbf{q}_{f}) + aP_{ee,y}^{2}(\mathbf{q}_{f}) \leq R_{coll}^{2}, \\ & 0 \leq aP_{ee,z}(\mathbf{q}_{f}) \leq H_{coll}, \\ & w\mathbf{P}_{ee}(\mathbf{q}_{f}) = \mathbf{b}(t_{f}) \\ & \cos(\langle_{w} \mathbf{z}_{ee}(\mathbf{q}_{f}), -\frac{\dot{\mathbf{b}}(t_{f})}{||\dot{\mathbf{b}}(t_{f})||} \rangle) = 1 \\ & |\mathbf{q}_{f} - \mathbf{q}_{0}| \leq \Delta \mathbf{q}(\mathbf{q}_{0}, \dot{\mathbf{q}}_{0}, t_{f} - t_{0}), \end{aligned}$$
(7)

where Γ_h is a user-defined quadratic cost function for the high-level planner, and $\Delta \mathbf{q}(\mathbf{q}_0, \dot{\mathbf{q}}_0, t_f - t_0)$ is the box constraint on \mathbf{q}_f given initial joint position, velocity and duration time.

Accurately calculating $\Delta \mathbf{q}(\mathbf{q}_0, \dot{\mathbf{q}}_0, t_f - t_0)$ requires calculating the optimal intermediate waypoints $\mathbf{q}(t)$ between $\mathbf{q}(t_0)$ and $\mathbf{q}(t_f)$ under kinematic constraints (7) for every catch time guess t_f need to be calculated. This is too time-consuming and an approximation is needed. In this paper, we use the following approximation:

$$\Delta q_i(q_{0,i}, \dot{q}_{0,i}, \Delta t) = \lambda \left(\dot{q}_{max,i} \left(\Delta t - \frac{\dot{q}_{max,i} - \dot{q}_{0,i}}{\ddot{q}_{max,i}} \right) + \frac{\dot{q}_{max,i}^2 - \dot{q}_{0,i}^2}{2\ddot{q}_{max,i}} \right), \quad (8)$$

where $\Delta t = t_f - t_0$ is the duration time, $\dot{q}_{max,i}$, $\ddot{q}_{max,i}$ are the maximum velocity and acceleration of the *i*-th joint, $\lambda \in (0,1)$ is a proportion parameter that needs to be tuned. The underlying assumption is that the maximum joint movement is proportional to the distance that the joint can traverse under the trapezoidal velocity profile [1]. Namely the joint first accelerates to its maximum speed and holds that speed until the end. We choose this approximation because of its simplicity and few corner cases needed to be considered in implementations. Other trajectory profiles like third-order polynomial can also be used for approximation [6]. The SQP algorithm is used to solve the problem specified in (7).

2) Low-level path planner: Given the desired configuration \mathbf{q}_{f}^{*} and catch time t_{f}^{*} from the high-level goal planner, the task of the low-level path planner is to generate a sequence of intermediate trajectory waypoints for each joint.

¹The step size γ used for joint motion equations is not necessarily the same as the step size δ_k used for the Kalman filter

For the *i*-th joint, this problem can be formulated as:

$$\mathbf{U}_{i}^{*} = \underset{\mathbf{U}_{i}}{\operatorname{argmin}} \Gamma_{l}(\mathbf{P}_{i}, \mathbf{U}_{i}) \\
\text{s.t.} \begin{bmatrix} q_{i}[k+1] \\ \dot{q}_{i}[k+1] \end{bmatrix} = \begin{bmatrix} 1 & \gamma \\ 0 & 1 \end{bmatrix} \begin{bmatrix} q_{i}[k] \\ \dot{q}_{i}[k] \end{bmatrix} + \begin{bmatrix} \frac{1}{2}\gamma^{2} \\ \gamma \end{bmatrix} u_{i}[k] \\
q_{min,i} \leq q_{i}[k] \leq q_{max,i} \\
\dot{q}_{min,i} \leq \dot{q}_{i}[k] \leq \dot{q}_{max,i} \\
\ddot{q}_{min,i} \leq u_{i}[k] \leq \ddot{q}_{max,i} \\
q_{i}[0] = q_{0,i}, \dot{q}_{i}[0] = \dot{q}_{0,i} \\
q_{i}[K] = q_{f,i}^{*}, \dot{q}_{i}[K] = 0,
\end{cases}$$
(9)

where Γ_l is a user-defined quadratic cost function for the low-level planner, $K = floor((t_f^* - t_0)/\gamma)$ is the finite horizon and floor(x) returns the greatest integer less than or equal to x, $\mathbf{P}_i = [q_i[1], \dots, q_i[K]]^T$ and $\mathbf{U}_i = [u_i[1], \dots, u_i[K]]^T$ are the position and control input sequence of the *i*-th joint.

This is an optimal control problem that consists of quadratic cost functions and linear constraints. Standard quadratic programming methods can be used to solve this problem efficiently [23]. With the solution of optimal control input sequence \mathbf{U}_i and initial conditions $q_i[0] = q_{0,i}, \dot{q}_i[0] = \dot{q}_{0,i}$, the optimal position sequence \mathbf{P}_i can be constructed via the kinematic model in (5).

C. Low-level Joint Control

Given the desired position trajectories generated by the motion planners, the task of the low-level joint controllers is to accurately track the trajectory. Accurate trajectory tracking is challenging due to the high-speed joint trajectories required for successful ball catching. Moreover, joint trajectory replanning may introduce discontinuities into the desired trajectories, which may be harder to track. To overcome these problems, we use the inverse dynamics learning technique developed in [13], shown in Fig. 4. It uses DNNs as an addon block to baseline feedback control systems to achieve a unity map between desired and actual outputs.

This method learns an inverse dynamics model of the baseline closed-loop system: $\hat{y}[k] = f(x[k], x[k+r])$, where x[k] is the system's discrete state at time index k, $\hat{y}[k]$ is the reference signal to the baseline control system, and r is the relative degree [13]. For discrete-time systems, the vector relative degree can be intuitively understood as the number of time steps before a change in the input is reflected in the output, which can be easily determined experimentally from the system's step response [13].

Specifically, for a single robot joint of the mobile manipulator, the DNN's inputs and outputs are:

$$\hat{y}_{i}^{d}[k+1] = NN(q[k], y_{i}^{d}[k+r]), \qquad (10)$$

where $y_j^d[k]$ is the reference signal sampled from desired trajectories for the *j*-th joint at time index *k*, and $\hat{y}_j^d[k+1]$ is the DNN's modified reference signal. As mentioned in [13], subtracting $y_j^d[k+1]$ from both the inputs and outputs of the neural network can improve the DNN's generalization ability since only the reference signal offsets $\tilde{y}_i^d[k+1] =$



Fig. 4: A diagram of the inverse dynamics learning framework and the training phase. The DNN module modifies the reference signal to the baseline system for trajectory tracking performance improvement. Adapted from [13].

 $\hat{y}_j^d[k+1] - y_j^d[k+1]$ instead of the reference signal itself $\hat{y}_j^d[k+1]$ are learned:

$$\tilde{y}_{j}^{d}[k+1] = NN(q[k] - y_{j}^{d}[k+1], y_{j}^{d}[k+r] - y_{j}^{d}[k+1]).$$
(11)

The modified position reference signal can then be sent to the baseline control system to actuate the robot joint.

V. FRAMEWORK IMPLEMENTATION AND EVALUATION

This section presents implementation details and evaluation of each of the three framework components.

A. Ball Estimation and Trajectory Prediction

First, we describe the implementation details of the ball state estimation and trajectory prediction. The modified Kalman filter, stated in Section IV-A, is used to reduce noise in VICON position and velocity measurements. In experiments, VICON position measurements are very precise. Therefore, ball state estimation consists of position measurements from the VICON system and velocity estimation from the Kalman filter. Trajectory prediction is done through numerical integration with a step size of 0.1 *ms*.

Both the Kalman filter (3) and trajectory prediction (1) need an estimate of the aerodynamic drag coefficient K_D . To estimate it, we record ball state data from 20 random ball tosses and use a recursive least-squares estimator. The estimated value of K_D is 0.0238 m^{-1} .

B. Motion Planning

We first describe the implementation details of the highlevel and low-level motion planner. The goal of the high-level planner is to find a catch configuration \mathbf{q}_f and catch time t_f while minimizing the robot movement:

$$\Gamma_h(\mathbf{q}_f, t_f) = w_a \|\mathbf{q}_{a,0} - \mathbf{q}_{a,f}\|_2^2 + w_b \|\mathbf{q}_{b,0} - \mathbf{q}_{b,f}\|_2^2, \quad (12)$$

where w_a, w_b are weight parameters that penalize movement of the arm joints and the mobile base, respectively. As mentioned in Section II, control accuracy of the mobile base is usually lower than the arm [11], [12]. Therefore, we select $w_a = 1, w_b = 5$ to reduce the mobile base motion. The proportion parameter used in the high-level planner (8) is set to be $\lambda = 0.4$. The high-level planner uses the SQP algorithm which we implemented using the NLopt c++ library [24]. The initial guess for the optimization variables (\mathbf{q}_f, t_f) are the robot's current configuration and a hard-coded catch time guess $t_f = 0.5 \ s$.

The goal of the low-level planner is to generate a smooth trajectory from current configuration \mathbf{q}_0 to catch configuration \mathbf{q}_f . This is achieved by penalizing joint acceleration change from one time step to the next during the trajectory, which can be written as follows for the *i*-th joint:

$$\Gamma_l(\mathbf{P}_i, \mathbf{U}_i) = \sum_{k=0}^{K-1} ||u_i[k+1] - u_i[k]||_2^2, \quad (13)$$

where *K* is calculated according to Section IV-B.2, and the discretization step size is $\gamma = 0.05$ *ms*. We solve the QP in the low-level planner with qpOASES c++ library [25].

We use numerical experiments to evaluate the motion planner's performance since we cannot control the ball initial state in real experiments when the ball is thrown by a human operator. We generated 2320 initial ball state samples, whose positions are uniformly distributed in a circle around the robot. In order to ensure that the resulting ball trajectories are feasible given the robot's motion constraints, we first note that the end effector can move 0.65 m in 0.86 s, as described in Section VI-B. We restrict the initial ball states to be such that after 0.7 s the ball is predicted, via numerical integration of (1), to be located inside a 0.50 m diameter sphere centered around the end effector of the given robot configuration.

The proposed motion planner can find a trajectory for 89.39% of all initial ball samples, and this success rate drops to 58.36% if we keep the base static, which shows the necessity of base movement. The average computation time is 5.10 *ms*, with 4.80 *ms* used for the high-level non-linear planner, and 0.3 *ms* used for the low-level (quadratic program) planner on an Intel Xeon CPU i7-8850H 2.60GHz computer. No sophisticated techniques like multi-threaded programming are used to speed up computation time.

C. Low-level Control

We now describe the implementation of the learning-based low-level controller. We first estimate experimentally the relative degree r of each joint, which is 8 for arm joints and 14 for mobile base joints.

To learn the mapping function in (11), we train 8 DNNs, one for each joint. Each DNN consists of two hidden layers and each layer contains 10 neurons. The ReLU function is used as the activation function. To create the training and testing dataset, we use the low-level path planner (described in Section IV-B.2) to generate trajectories similar to those used in real ball catching experiments. We characterize these trajectories by bounding the maximum joint displacement $\Delta \mathbf{q}$ in a given amount of time Δt . In our experiments we set $\Delta t = 2s, 3s, 4s$, to represent trajectories with different speeds and $\Delta \mathbf{q} = [1.36, 1.36, 1.36, 1.75, 1.75, 2.53, 0.76, 0.76]^T$, where the first six joint displacements correspond to the arm and are specified in radians, and the last two joint displacements.

The trajectories are executed using the robot's baseline control system. It consists of off-board PD controllers that generate desired joint velocities and on-board controllers of the UR10 arm and mobile base that turn desired velocities into joint torques (UR10 arm) or wheel rotation velocities (mobile base). In this way we collect a dataset with 1000 data pairs for each joint. The dataset is divided with 80% is used for training and the rest is used for validation. We use Adam optimizer implemented in *Tensorflow* to train the neural network with 50 training iterations.

To test the influence of inverse dynamics learning on tracking performance, we generate testing trajectories by setting $\Delta \mathbf{q} = [0.8, 0.8, 0.8, 1.1, 1.1, 1.1, 0.20, 0.20]^T$ and $\Delta t = 1 s$, which are different from training and validation trajectories in the dataset. The testing trajectories are executed ten times each when the robot uses the learning technique and ten times each when no learning technique is used. The trajectory tracking rooted-mean-squared error (RMSE) over the ten trajectories is shown in Table I. It can be seen that the inverse dynamics learning method can reduce tracking errors between $35\% \sim 80\%$ with the highest error reduction corresponding to the mobile base.

TABLE I: Joint trajectory tracking error comparison

Joint	With DNN [°]	Without DNN [°]	Reduction
arm 1	0.0218	0.0508	57.2%
2	0.0300	0.0508	40.9%
3	0.0264	0.0506	47.8%
4	0.0309	0.0561	45.0%
5	0.035	0.0562	36.5%
6	0.034	0.0561	37.1%
Joint	With DNN [m]	Without DNN [m]	Reduction
base x	0.0041	0.0196	79.1%
У	0.0047	0.0183	74.0%

VI. EXPERIMENTS

This section presents the experimental setup and ball catching experimental results. A video demonstrating the real robot experiments can be found at http://tiny.cc/ball_catch.

A. Experimental Setup

The mobile manipulator used in this experiment consists of a 6-DOF UR10 arm and a 3-DOF Ridgeback omnidirectional mobile base driven by four Mecanum wheels. We use symmetric joint motion boundary values, i.e. $\mathbf{q}_{min} = -\mathbf{q}_{max}$, $\dot{\mathbf{q}}_{min} = -\dot{\mathbf{q}}_{max}$. The six arm joints have the same maximum joint positions of 180° and accelerations of $458^{\circ}/s^2$. Their maximum joint velocities are $103^{\circ}/s$, $103^{\circ}/s$, $103^{\circ}/s$, $126^{\circ}/s$, $126^{\circ}/s$, and $180^{\circ}/s$ for arm joint $1 \sim 6$ respectively. The maximum positions, velocities and accelerations of the mobile base in the *x* and *y* axes are 3m, 1.0m/s, $2.5m/s^2$, respectively.

The three-finger end effector grasps a 94.90 *mm* diameter metal cup that is used to catch the ball, as shown in Fig. 5. The proposed framework is implemented via Robot Operation System (ROS) and runs on a ThinkPad P52 laptop.



Fig. 5: The 61.76 *mm* diameter ball and 94.90 *mm* diameter cup used in this paper. The position accuracy requirement in the *x*, *y*, *z* axes is approximately (94.90-61.76)/2 = 16.57 *mm*. The typical speed at the catch time is around 5 *m/s*, leading to the 16.57/5 \approx 3 *ms* time precision requirement.

The ball used in these experiments is a standard 61.76 mm tennis ball with a weight of 81.76 g. The ball is wrapped with retro-reflective tape in order to be visible to the VICON motion capture system, which provides position and velocity measurements to be used in the ball state estimation module.

B. Ball Catching Experiment

We conducted experiments to test three main characteristics of the proposed framework: (i) ball catching success rate, (ii) impact of the learning-based controller, and (iii) impact of QP-based low-level planner.

1) Experiment A: Ball Catching Success Rate: In this experiment, we use the framework proposed in Section IV and throw balls from different locations while the robot has different initial configurations. Three initial robot configurations are chosen. For each initial robot configuration, we throw the ball from five different locations, five throws for each location, for a total of 75 ball catching experiments. All ball trajectories are shown in Fig. 6. Of the 75 ball throws, 65 balls are caught successfully by the robot, which is a 85.33% success rate. The end effector position trajectory tracking RMSE over the 75 experiments are 11.03 mm, 14.01 mm, 13.25 mm for x, y, z axes respectively. Among the successful catches, the maximum end effector displacement in a catch is 0.65 m in 0.86 s with a mobile base displacement of 0.26 m in that catch. The corresponding arm and base position trajectories for this catch are shown in Fig. 7a and 7b. It can be seen that our mobile base is able to move accurately in both x and y directions while the humanoid in [2] can only move in one direction.

2) Experiment B: Impact of Learning-based Controller: In this experiment, we are interested in showing that the learning-based controller improves tracking performance, therefore improving the success rate of the overall framework. It is hard to repeat a ball trajectory twice because a human throws it. To do a fair comparison in this experiment we throw balls from the same location (2.5, 2.0, 0) while the robot has the same initial configuration, as shown in Fig. 1. We make 40 throws when the robot uses inverse dynamics learning and 40 when the robot does not use inverse dynamics learning for a total of 80 throws. The success rate for this particular configuration is 75% when



Fig. 6: Ball trajectories of Experiment A. Starting from the left corner and in the anti-clockwise direction, the five ball throwing positions in the *xy* plane in meters are (2.8, -3.1), (2.9, 0.0), (2.5, 2.0), (0.0, 2.8), (-2.0, 2.0). Ball trajectories are colored according to its throwing location.



Fig. 7: An example of (*a*) arm joint and (*b*) mobile base tracking performance. Dashed lines represent desired trajectories, while solid lines represent actual trajectories. For a better illustration effect, we subtract initial values from trajectory waypoints, and thus joint angles relative to their initial values are shown here.

using inverse dynamics learning and 65% when not using it. The ten percent decrease shows the effectiveness of inverse dynamics learning. Note that the success rate for this particular configuration is lower than the success rate in Experiment A. Visual inspection of the data and video show that the geometrical setup, i.e. ball throwing location and robot initial configuration, used in this experiment requires more robot maneuvers than other geometrical setups used in Experiment A.

3) Experiment C: Impact of QP-based Low-level Planner: In this experiment, we are interested in showing the impact of the QP-based low-level planner in the ball catching success rate. To do this, we implemented the trapezoidal velocity profile method in [1] for the low-level planner, and tested it by throwing 40 balls with the geometrical setup as in Experiment B. Note that the inverse dynamics learning is still used. The robot achieves a success rate of 72.5%, similar to the success rate of 75% when using the QP method.

The QP method does not have a significant impact on ball catching success rate. One possible reason is that the way the robot accelerates and decelerates does not have a large impact on the ball catching success due to the short duration of the trajectory [1]. However, the QP method does allow for more flexible and smooth trajectories, which may improve the performance of tasks with longer duration. Finally, similar success rates with the two different low-level planners may be due to the high accuracy of the learning-based controller when tracking arbitrary trajectories.

C. Failure analysis

The success of ball catching tasks requires a tight integration of every system component. The malfunction of any component will result in a failure. Among all throws in Experiment A, B, C, failure has occurred for the following reasons: inaccurate ball trajectory prediction, large tracking errors, and failure to find feasible solutions of the highlevel planning problem. Large tracking errors occur when maneuvers are too aggressive and overshooting happens. We could improve this by penalizing aggressive trajectories.

VII. CONCLUSION

We presented a framework for accurate high-speed motion generation and execution on mobile manipulators with the application scenario of ball catching. A modified Kalman filter that considers aerodynamics is used for estimating the ball's velocity, and ball trajectory prediction is done via numerical integration. A bi-level optimization scheme is proposed to handle complex trajectory requirements in real-time, with an average run time of 5.1 ms. The highlevel goal planning problem is formulated as a non-linear optimization problem and solved by SQP. The low-level path planning problem is solved by QP. The joint control is driven by an inverse dynamics learning method. This learning-based controller reduces joint tracking errors by up to 79.1%, which translate into a 10% increase in success rate compared to a non-learning-based joint controller. The proposed framework is validated via extensive real robot catching experiments under different configurations and achieves a success rate of 85.33%, setting a new state of the art. In future work we aim to focus on dynamic planning instead of kinematic planning, which should improve the feasibility of joint trajecotries.

REFERENCES

- B. Bäuml, T. Wimböck, and G. Hirzinger, "Kinematically optimal catching a flying ball with a hand-arm-system," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems* (*IROS*), 2010, pp. 2592–2599.
- [2] B. Bäuml, O. Birbach, T. Wimböck, U. Frese, A. Dietrich, and G. Hirzinger, "Catching flying balls with a mobile humanoid: System overview and design considerations," in *Proc. of the IEEE-RAS International Conference on Humanoid Robots*. IEEE, 2011, pp. 513–520.

- [3] D. M. Bodily, T. F. Allen, and M. D. Killpack, "Motion planning for mobile robots using inverse kinematics branching," in 2017 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2017, pp. 5043–5050.
- [4] A. De Luca, G. Oriolo, and P. R. Giordano, "Kinematic modeling and redundancy resolution for nonholonomic mobile manipulators," in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2006, pp. 1867–1873.
 [5] L.-P. Ellekilde and H. I. Christensen, "Control of mobile manipulator
- [5] L.-P. Ellekilde and H. I. Christensen, "Control of mobile manipulator using the dynamical systems approach," in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2009, pp. 1370–1376.
- [6] O. Koç, G. Maeda, and J. Peters, "Online optimal trajectory generation for robot table tennis," *Robotics and Autonomous Systems*, vol. 105, pp. 121–137, 2018.
- [7] U. Frese, B. Bauml, S. Haidacher, G. Schreiber, I. Schäfer, M. Hahnle, and G. Hirzinger, "Off-the-shelf vision for a robotic ball catcher," in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots* and Systems (IROS), vol. 3, 2001, pp. 1623–1629.
- [8] S. S. Mirrazavi Salehian, M. Khoramshahi, and A. Billard, "A dynamical system approach for catching softly a flying object: Theory and experiment," *IEEE Transaction on Robotics*, 2016.
- [9] R. Lampariello, D. Nguyen-Tuong, C. Castellini, G. Hirzinger, and J. Peters, "Trajectory planning for optimal robot catching in realtime," in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2011, pp. 3719–3726.
- [10] P. A. Flores Álvarez, "Modelling, simulation and experimental verification of a wheeled-locomotion system based on omnidirectional wheels," Master's thesis, Pontificia Universidad Católica del Perú, Ilmenau, Germany, 2016.
- [11] KUKA, "Mobile robotics kmr quantec," https://www.kuka.com/en-ca/ products/mobility/mobile-robots/kmr-quantec, [Online; accessed 20-February-2020].
- [12] U. Robots, "Technical specifications ur10," https://www. universal-robots.com/media/50880/ur10{_}bz.pdf, [Online; accessed 20-February-2020].
- [13] S. Zhou, M. K. Helwa, and A. P. Schoellig, "Design of deep neural networks as add-on blocks for improving impromptu trajectory tracking," in *Proc. of the IEEE Conference on Decision and Control (CDC)*, 2017, pp. 5201–5207.
- [14] H. Seraji, "A unified approach to motion control of mobile manipulators," *The International Journal of Robotics Research*, vol. 17, no. 2, pp. 107–118, 1998.
- [15] A. Punjani and P. Abbeel, "Deep learning helicopter dynamics models," in Proc. of the IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2015, pp. 3223–3230.
- [16] M. T. Frye and R. S. Provence, "Direct inverse control using an artificial neural network for the autonomous hover of a helicopter," in *Prof of the IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 2014, pp. 4121–4122.
- [17] Q. Li, J. Qian, Z. Zhu, X. Bao, M. K. Helwa, and A. P. Schoellig, "Deep neural networks for improved, impromptu trajectory tracking of quadrotors," in *Proc. of the IEEE International Conference on Robotics* and Automation (ICRA). IEEE, 2017, pp. 5183–5189.
- [18] K. P. Hawkins, "Analytic inverse kinematics for the universal robots ur-5/ur-10 arms," Georgia Institute of Technology, Tech. Rep., 2013.
- [19] M. Müller, S. Lupashin, and R. D'Andrea, "Quadrocopter ball juggling," in *Proc. of the IEEE/RSJ international conference on Intelligent Robots and Systems*. IEEE, 2011, pp. 5113–5120.
- [20] M. Jakob, "Position and force control with redundancy resolution for mobile manipulators," Master's thesis, University of Stuttgart, Stuttgart, Germany, 2018.
- [21] T. D. Barfoot, *State estimation for robotics*. Cambridge University Press, 2017.
- [22] S. Kim, A. Shukla, and A. Billard, "Catching objects in flight," *IEEE Transactions on Robotics*, vol. 30, no. 5, pp. 1049–1065, Oct 2014.
- [23] S. Boyd, S. P. Boyd, and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [24] S. G. Johnson, "The nlopt nonlinear-optimization package," http: //github.com/stevengj/nlopt, [Online; accessed 20-February-2020].
- [25] H. Ferreau, C. Kirches, A. Potschka, H. Bock, and M. Diehl, "qpOASES: A parametric active-set algorithm for quadratic programming," *Mathematical Programming Computation*, vol. 6, no. 4, pp. 327–363, 2014.