

# A Solution to Adaptive Mobile Manipulator Throwing

Yang Liu Aradhana Nayak Aude Billard

**Abstract**—Mobile manipulator throwing is a promising method to increase the flexibility and efficiency of dynamic manipulation in factories. Its major challenge is to efficiently plan a feasible throw under a wide set of task specifications. We show that the mobile manipulator throwing problem can be simplified to a planar problem, hence greatly reducing the computational costs. Using machine learning approaches, we build a model of the object’s inverted flying dynamics and the robot’s kinematic feasibility, which enables throwing motion generation within 1 ms for given query of target position. Thanks to the computational efficiency of our method, we show that the system is adaptive under disturbance, via replanning on the fly for alternative solutions, instead of sticking to the original throwing plan.

## I. INTRODUCTION

Throwing, as one form of dynamic manipulation [1], [2], can not only augment the feasible work space of the robot, but also increase the efficiency of object manipulation. This is highly desirable in applications such as logistics and handling of goods. However, automated robot manipulation with human-level speed is generally difficult to achieve and good solutions need to be hand-tuned for specific tasks [3], which is a common drawback of model-based methods. On the other hand, with the emergence of machine learning, there is a growing popularity on using model-free methods that distill skills from offline training. With sufficient data coverage, the robot is able to handle a wide set of tasks once deployed. Despite impressive empirical result [4], model-free methods in the literature lead to algorithms which are difficult to generalize to new tasks and can not ensure successful task execution.

We take an approach that combines the strength of model-based methods and model-free methods: model-free learning to model the complex non-linear object flying dynamics [5] and model-based approach to formulate throwing as a feasibility problem. To design an architecture that combines these two modules in harmony, we start by asking the following questions:

- What is an appropriate task representation for throwing?
- How to separate mundane computation *offline* while ensuring reliable and efficient *online* solution generation?

To answer the above questions, we conduct an analysis on the structure of mobile manipulator-throwing problem, and design a data structure that enables efficient online query for throwing configurations. As a result, we propose a framework for mobile manipulator throwing that is able to generate

The authors are with Learning Algorithms and Systems Laboratory (LASA), EPFL, Switzerland. Email: {yang.liuu, aradhana.nayak, aude.billard}@epfl.ch. Corresponding author: Yang Liu.




Fig. 1: Snapshots of a mobile manipulator throwing.

TABLE I: Classification of approaches on robot throwing.

	ad-hoc throwing	generic throwing
model-based	[6]–[8], [10], [11], [13]	Optimization [14], [15] Sampling [16] Ours
model-free	[9]	End-to-end Learning [4]

throwing motion within 1 ms<sup>1</sup>. We demonstrate a sample throw using a mobile manipulator in Fig. 1.

### Related Work

Previously, several approaches have been proposed to solve robot throwing problem [4], [6]–[12]. We notice that besides the dichotomy of model-based methods [6]–[8], [10]–[12] and model-free methods [4], [9] in the literature, there also exists a large difference on Operational Design Domain (ODD) among different works. We roughly categorize them into two types of ODDs: ad-hoc throwing with narrow ODD and generic throwing with wide ODD, see Table I.

Most early works on robot throwing have been designed on a very narrow ODD with ad-hoc throwing solutions [6]–[11], [13]. They either worked on simple robots with low degrees of freedom [6]–[9], [13] or had a specific object to throw [10], [11].

Among the works aiming for generic throwing, many tried to accurately model the flying dynamics of objects as well as the robot dynamics and determine the optimal robot throwing trajectory via numerical optimization [14], [15]. However, as we shall see later in the paper, the optimization problem for throwing is hard to solve because of the nonconvex constraints. In [15], it takes 0.5s to generate feasible throwing trajectory for a 3–DOF robot. Therefore, it presumably needs at least several seconds for our mobile manipulator

<sup>1</sup>The source code for our algorithms and simulations is available at: <https://github.com/epfl-lasa/mobile-throwing>




Fig. 2: The Backward Reachable Tube models the set of valid throwing configurations, given object’s flying dynamics.

with 7-DOF arm, which is not suitable for online motion generation.

Another approach for generic throwing is using sampling-based methods [16] such as Probabilistic Roadmap (PRM) [17] or Rapidly-exploring Random Trees (RRT) [18]. Although sampling-based methods provide probabilistic completeness, Zhang et al. [16] reported that it takes up to 1 minute to get a throwing trajectory, which is far from human-level speed.

Recently, TossingBot [4] offered an end-to-end learning approach for throwing. The throwing configuration is generated from throwing velocity predicted by trained model given throwing query and object image. The solution is generic and can handle a large set of objects and multiple target positions in the training set. However, the performance degrades for unseen query during training. While one could retrain the robot with new data, it is not clear how fast the robot can learn to throw new objects.

### Summary of Contribution

Our approach complements the state-of-the-art by presenting a method for efficiently computing feasible throwing configurations and motions for mobile manipulators. In particular, we offer a comprehensive analysis on the problem of mobile manipulator throwing, including:

- A geometric analysis, which exposes the independent variables and hence the task manifold of mobile manipulator throwing;
- A data structure for offline robot kinematic analysis that enables fast online query of throwing configurations.

With the above analysis, our framework enables the system to:

- Throw a rigid-body object given its flying dynamics;
- Throw towards an arbitrary target position;
- Guarantee robot-feasibility and throwing-validity;
- Generate solutions within 1 ms and achieve adaptive throwing.

## II. METHOD

Given a throwing task specified by object flying dynamics, target landing position and possibly landing velocity range, there is a set of valid throwing configurations (throwing position and throwing velocity) varied in landing velocity and




Fig. 3: Visualization of throwing triangle and throwing plane.

TABLE II: Notations for geometric modeling in Fig. 3.

$A$	Manipulator base
$E$	End effector
$B$	Target box
$\triangle AEB$	Throwing triangle
$\phi$	Throwing yaw angle
$r$	Throwing range
$z$	Throwing height
$\dot{r}$	Horizontal throwing velocity
$\dot{z}$	Vertical throwing velocity
$v$	Throwing speed
$\gamma$	Throwing pitch angle

flying time. We call the set of valid throwing configurations Backward Reachable Tube (BRT). As shown in Fig. 2, our idea is to bring the object towards the BRT and release the object once the robot end-effector enters the BRT.

### A. Geometric modeling

The geometric relationship between robot mobile manipulator and task throwing is shown in Fig. 3a and Fig. 3b. The notations are introduced in Table II.

We observe that there exist geometric equivariance structures in mobile manipulator throwing, namely:

- Mobile manipulator alone is equivariant in horizontal position because of the mobile base with omnidirectional wheels.
- Throwing alone is equivariant in incident direction.
- Mobile manipulator throwing can be described by throwing triangle  $\triangle AEB$ , which is equivariant in rotation around Z-axis at target box B.

Above analysis showed that to find out the joint configuration for throwing, we only have to determine  $\triangle AEB$ . Furthermore, the triangle family around B could also be




Fig. 4: Tradeoff between small yaw angle and large throwing velocity. For configurations with large feasible Cartesian velocity (large maximum singular value), the maximum velocities tend to be perpendicular to arm stretching direction (small dot product value). See Subsection II-B for detailed explanation.

useful to avoid obstacle or communicate robot intent. These insights are crucial to our efficient and reliable solver.

### B. Consideration of yaw angle

The essence of throwing is to move the object to somewhere not reachable by the robot end-effector. A typical choice in the literature [4], [7], [9], [14] is to assume that the throwing yaw angle  $\phi$  is zero, i.e., throwing velocity direction is aligned with end-effector horizontal direction  $\overrightarrow{AE}_{xy}$  and hence the throwing triangle  $\triangle AEB$  is degenerated to a line. However, we observe that throwing yaw angle  $\phi = 0$  is not necessarily a good choice. Intuitively, if the arm is horizontally stretched, it is hard to generate velocity along the arm yet easy to generate velocity perpendicular to the arm. This hypothesis is validated in the following manipulability analysis.

We sample 1 million joint configurations for a 7 DoF manipulator. For each sampled joint configuration  $q$ , we calculate their corresponding end-effector horizontal positions  $\overrightarrow{AE}_{xy}(q)$  and Jacobian in  $XY$  plane  $J_{xy}(q)$ . Then we group the 1 million data according to two attributes:

- Maximum singular value of  $J_{xy}$ , denoted as  $s_1(J_{xy})$ , approximates the maximum horizontal velocity of the end-effector with unit joint velocity limit;
- The dot product between  $\overrightarrow{AE}_{xy}$  and the singular vector associated with  $s_1(J_{xy})$ , approximates the “alignment” between maximum horizontal velocity direction and arm stretching direction  $\overrightarrow{AE}_{xy}$ , note that these two directions are aligned when the dot product is 1 and are perpendicular when the dot product is 0.

The resulting 3D histogram is shown in Fig. 4.

One can observe that indeed there exists a tradeoff between small yaw angle and large throwing velocity, and if we predetermine  $\phi = 0$ , a lot of appropriate throwing configurations are filtered out. Therefore, in this paper, we leave the throwing yaw angle as a decision variable and let the solver handle this tradeoff automatically.

### C. Backward Reachable Tube

In the throwing plane, we denote the object flying state as  $x = [r, z, \dot{r}, \dot{z}]^T$ . The flying dynamics is described by a first order differential equation  $\dot{x} = f(x)$ ,  $x \in \mathbb{R}^4$ . We denote the flying trajectories of  $f$  starting from state  $x_0$  as  $\zeta_{f,x_0}(t) : [0, +\infty) \rightarrow \mathbb{R}^4$ . We assume that a user is giving the robot a landing target set  $\mathcal{X}_l \subseteq \mathbb{R}^4$ , describing the allowed landing position slack and allowed range of landing velocities.

If there exist a position along the object flying trajectory that enters the landing target set, then any state on the trajectory segment before the entering state is a valid throwing configuration. Hence, if we aggregate all the trajectories that eventually enter the target landing set, we get the set of valid throwing configurations, which we call the Backward Reachable Tube (BRT).

Mathematically, the BRT is defined as:

$$BRT_{f,\mathcal{X}_l} = \{x_0 \mid \exists t \geq 0, \zeta_{f,x_0}(t) \in \mathcal{X}_l\}.$$

The BRT can be generated from object flying dynamics as follows:

1) *Choose landing target set  $\mathcal{X}_l$* : the landing target set can be simply specified as “position box” and “velocity box”. Note that the horizontal velocity bound can be regarded as a limit to avoid flip out of the box and vertical velocity limit can be regarded as a limit to avoid object damage.

2) *Generating data inside BRT*: exact reachability analysis is generally computationally difficult for nonlinear dynamics. Our idea is to approximate BRT by sampling data inside it. Our observation is that flying dynamics is not memoryless and hence its corresponding Ordinary Differential Equation (ODE) can be solved backwards in time. Therefore, we sample landing configurations from landing target set  $\mathcal{X}_l$ , and solve the flying dynamics backwards in time as a Initial Value Problem (IVP) and then aggregate the ODE solutions. This is much more efficient than otherwise naive Boundary Value Problem (BVP) and hence can generate the BRT dataset quickly.

In this paper, the flying dynamics of an object is considered to be known apriori. This is assumed to be derived from state of the art system identification methods [19] or machine learning techniques [5].

3) *Learning-based implicit representation of BRT*: with gathered data inside BRT, we use a Neural Network (NN) to represent the BRT as a level-set function  $f_{BRT}(x) : \mathbb{R}^4 \rightarrow \mathbb{R}$ , where:

- $f_{BRT}(x) > 0 \Leftrightarrow x \notin BRT$ ;
- $f_{BRT}(x) \leq 0 \Leftrightarrow x \in BRT$ ;
- $f_{BRT}(x) = 0 \Leftrightarrow x \in \partial BRT$ .

#### D. Feasibility problem formulation

The BRT only provides valid throwing configurations inside throwing plane in Euclidean space. To find throwing configurations in robot joint space, there are two types of nonlinear constraints to be considered:

- Equality constraint: the throwing triangle  $\triangle AEB$  defines a task manifold, i.e. for fixed  $\overrightarrow{AB}$ , the end-effector position  $\overrightarrow{AE}$  must lie in the throwing plane defined by  $\overrightarrow{EB}$ , which imposes a nonlinear equality constraint in the joint configuration  $q$ .
- Inequality constraint: joint position limit and joint velocity limit.

These constraints can be expressed into the following Throwing Feasibility Problem (TFP):

$$\begin{aligned} & \text{Find } \{r, z, \dot{r}, \dot{z}, \overrightarrow{EB}, \overrightarrow{AB}, q, \dot{q}\} \\ & \text{such that: } \left\{ \begin{array}{l} r = \|\overrightarrow{EB}_{xy}\| \\ z = -\overrightarrow{EB}_z \\ \overrightarrow{EB} = \overrightarrow{AB} - \overrightarrow{AE} \\ \overrightarrow{AE} = \text{fkine}(q) \\ \vec{v} = \begin{bmatrix} v_x \\ v_y \\ v_z \end{bmatrix} = \begin{bmatrix} \dot{r} \frac{\overrightarrow{EB}_{xy}}{\|\overrightarrow{EB}_{xy}\|} \\ \dot{z} \end{bmatrix} \\ \dot{q} = J^\dagger(q)\vec{v} \\ q_{\min} \leq q \leq q_{\max} \\ \dot{q}_{\min} \leq \dot{q} \leq \dot{q}_{\max} \\ f_{BRT}(r, z, \dot{r}, \dot{z}) \leq 0 \end{array} \right. \end{aligned} \quad (\text{TFP})$$

where  $\text{fkine}(\cdot)$  is the robot forward kinematics, and  $\overrightarrow{EB}_z$  is projection of  $\overrightarrow{EB}$  along vertical direction  $Z$ . TFP is difficult to solve because of the nonlinear equality constraint. However, geometric analysis provides the following reasoning:

- Once  $(\overrightarrow{AB}, q)$  is fixed:
  - $\triangle AEB$  is determined,
  - $(\overrightarrow{EB}, r, z)$  is determined,
  - $(\dot{r}, \dot{z})$  is to be determined;
- Once  $(\dot{r}, \dot{z})$  is further determined:
  - $\dot{q} = J^\dagger(q)\vec{v}$ ;
- Independent variables:  $(\overrightarrow{AB}, q, \dot{r}, \dot{z})$ .

From the above analysis, we expose the ‘‘task manifold’’ parameterized explicitly by  $(\overrightarrow{AB}, q, \dot{r}, \dot{z})$  and hence eliminated the manifold constraint in the naive formulation. Now the problem reduces to the following formulation with only

inequality constrains:

$$\begin{aligned} & \text{Find } \{\overrightarrow{AB}, q, \dot{r}, \dot{z}\} \\ & \text{such that: } \left\{ \begin{array}{l} q_{\min} \leq q \leq q_{\max} \\ \dot{q}_{\min} \leq J^\dagger(q)\vec{v}(\overrightarrow{AB}, q, \dot{r}, \dot{z}) \leq \dot{q}_{\max} \\ f_{BRT}(r(\overrightarrow{AB}, q), z(\overrightarrow{AB}, q), \dot{r}, \dot{z}) \leq 0 \end{array} \right. \end{aligned} \quad (\text{TFP-Reduce})$$

#### E. Robot kinematics analysis

The reduced feasibility problem still suffers from the nonconvex inequality constraint:

$$\dot{q}_{\min} \leq J^\dagger(q)\vec{v}(\overrightarrow{AB}, q, \dot{r}, \dot{z}) \leq \dot{q}_{\max}$$

The constraint is linear in  $(\overrightarrow{AB}, \dot{r}, \dot{z})$  but nonlinear in  $q$ , and it reflects different velocity capabilities at different arm configurations of the robot. As  $q$  also determines end-effector position, this makes the robot kinematics analysis intertwined with throwing task.

To mitigate the difficulty of finding a suitable  $q$  to generate desired velocity at desired position, we propose *velocity hedgehog*<sup>2</sup>, a data structure that enables efficient online query of robot configurations. The velocity hedgehog discretizes the end-effector height  $z$  and throwing direction  $(\phi, \gamma)$  into cells. At each cell, it stores,

- The maximum feasible end-effector speed  $v_{max}$  at height  $z$  along direction  $(\phi, \gamma)$ ;
- The robot configuration  $q$  that enables  $v_{max}$ .

Later, velocity hedgehog is used to generate batches of initial guesses for robot joint configurations in 1 millisecond. The algorithm to generate the velocity hedgehog is shown in Algorithm 1, where a Linear Program (LP) is solved as a subroutine to get maximum speed along direction  $(\phi, \gamma)$  at robot configuration  $q$  in the following manner:

$$\begin{aligned} & \max_s s \\ & \text{subject to: } \left\{ \begin{array}{l} \dot{q}_{\min} \leq J^\dagger(q)\vec{v} \leq \dot{q}_{\max} \\ \vec{v} = s \begin{bmatrix} \cos \gamma \cos \left( \frac{\overrightarrow{AE}_y}{\overrightarrow{AE}_x} + \phi \right) \\ \cos \gamma \sin \left( \frac{\overrightarrow{AE}_y}{\overrightarrow{AE}_x} + \phi \right) \\ \sin \gamma \end{bmatrix} \end{array} \right. \end{aligned} \quad (\text{LP})$$

A typical velocity hedgehog is shown in Fig. 5, indicating that the feasible velocity set is not convex but appears as needles with different lengths along different directions and resembles a hedgehog.

#### F. Combining robot velocity hedgehog with BRT to determine throwing configurations

Here we introduce a simple method to combine robot velocity hedgehog and BRT, resulting in an efficient method to get qualitatively different initial guesses for  $(q, \phi, x)$  quickly. We first group the BRT data according to height

<sup>2</sup>The name comes from the spiky shape of the velocity distribution.

---

**Algorithm 1:** Algorithm to get velocity hedgehog

---

**Input** : robot model with forward kinematics and differential forward kinematics *robot*  
**Output** :  $\max\_z\_phi\_gamma$ ,  $q\_z\_phi\_gamma$   
**Data** : robot joint position limit  $(q_{min}, q_{max})$ , robot joint velocity limit  $(\dot{q}_{min}, \dot{q}_{max})$ , joint grid size  $\Delta q$ , velocity hedgehog grids  $Z, \Phi, \Gamma$

```
/* Build robot dataset  $\mathcal{X}$  */
1  $\mathcal{Q} \leftarrow \text{ComputeMesh}(q_{min}, q_{max}, \Delta q)$ 
/* Filter out  $q$  with small singular value */
2  $\mathcal{X} \leftarrow \text{FilterBySingularValue}(\mathcal{Q})$ 
/* Group data by  $Z$  */
3  $\{\mathcal{X}_z\} \leftarrow \text{GroupBy}(\mathcal{X}, Z)$ 
/* Initialize velocity hedgehog */
4  $\max\_z\_phi\_gamma \leftarrow \text{zeros}([\#Z, \#\Phi, \#\Gamma])$ 
5  $q\_z\_phi\_gamma \leftarrow \text{arrays}([\#Z, \#\Phi, \#\Gamma, \#\text{joints}])$ 
/* Build velocity hedgehog */
6 for  $[z, \phi, \gamma, data] \in Z \times \Phi \times \Gamma \times \mathcal{X}_z$  do
/* Get max. speed along  $(\phi, \gamma)$  at joint configuration  $q$  */
7  $res \leftarrow LP(\phi, \gamma, data.q, robot, \dot{q}_{min}, \dot{q}_{max})$ 
8 if  $res > \max\_z\_phi\_gamma(z, \phi, \gamma)$  then
9 |  $\max\_z\_phi\_gamma(z, \phi, \gamma) \leftarrow res$ 
10 |  $q\_z\_phi\_gamma(z, \phi, \gamma) \leftarrow data.q$ 
11 end
12 end
13 return  $\max\_z\_phi\_gamma, q\_z\_phi\_gamma$ 
```

---

---

**Algorithm 2:** Algorithm to get initial guesses

---

**Input** : robot velocity hedgehog, BRT dataset  
**Output** : Set of initial guess for joint configuration and throwing configuration  $(q, \phi, x)$

```
1  $\mathcal{X}_{z,\gamma} \leftarrow \text{GroupBy}(\mathcal{X}_{BRT}, Z, \Gamma)$ 
2  $q\_guess, \phi\_guess, x\_guess \leftarrow \text{empty}$ 
3  $idxs \leftarrow \text{where}(\max\_z\_phi\_gamma > \mathcal{X}_{z,\gamma})$ 
4  $q\_guess = q\_z\_phi\_gamma[idxs]$ 
5  $\phi\_guess = \Phi[idxs]$ 
6  $x\_candidates = \mathcal{X}_{z,\gamma}[idxs]$ 
7 return  $q\_guess, \phi\_guess, x\_guess$ 
```

---

$z$  and throwing pitch angle  $\gamma$ . Then for BRT data in bin indexed by  $\hat{z}$  and  $\hat{\gamma}$ , the initial guesses for throwing state are the BRT data whose flying speed is smaller than the maximum feasible velocity at a certain  $(\hat{z}, \hat{\gamma})$  stored in velocity hedgehog. From this matching operation, we can also read out the corresponding joint configurations and throwing yaw angles from velocity hedgehog, resulting in initial guesses of  $(q, \phi, x)$ . The detailed algorithm is shown in Algorithm 2. The obtained feasible initial guesses can be fed into Problem TFP-Reduce, to be further refined. With  $(q, \phi, x)$  determined, all the other decision variables can be written in closed-form as shown in Problem TFP. As a result,




Fig. 5: Illustration of a typical velocity hedgehog. Red needles represents the feasible velocities along different throwing pitch angles with fixed throwing height  $z$  and throwing yaw angle  $\phi$ , which are stored in velocity hedgehog. Given certain throwing height and throwing direction  $(z^*, \phi^*, \gamma^*)$  query, velocity hedgehog gives maximum feasible velocity  $v_{max}^*$  and the joint configuration  $q^*$  that enables  $v_{max}^*$ .

the overall architecture to solve TFP efficiently and reliably is shown in Fig. 6.

### III. EXPERIMENT

We verify our method in the simulated environment Py-Bullet [20] with following setups:

- Known accurate flying dynamics;
- Perfect grasping (object won't fall off while on the way to throwing configuration);
- Instant gripper opening;
- Perfect throwing trajectory tracking;
- Disabled collision checking between robot and box.

The object is a ball with radius 5 cm and the target box dimensions are 5 cm  $\times$  25 cm  $\times$  25 cm. This paper's emphasis is on finding valid joint space throwing configuration in a reliable and computational efficient manner. Therefore, the above assumptions are mild.

#### A. BRT data generation

We describe the data generation procedure for object flying states inside BRT as defined in Subsection II-C. For the sake of simplicity, we use pure projectile flying dynamics with gravity  $g = 9.81 \text{ m/s}^2$ , i.e.:

$$\dot{x} = f(x) = \frac{d}{dt} \begin{bmatrix} r \\ z \\ \dot{r} \\ \dot{z} \end{bmatrix} = \begin{bmatrix} \dot{r} \\ \dot{z} \\ 0 \\ -g \end{bmatrix},$$

with following target set  $\mathcal{X}_l$ :

$$\mathcal{X}_l = \left\{ \begin{array}{l} r = 0 \\ z = 0 \\ 0.2 \leq \dot{r} \leq 2.0 \\ -5.0 \leq \dot{z} \leq -2.0 \end{array} \right\}.$$

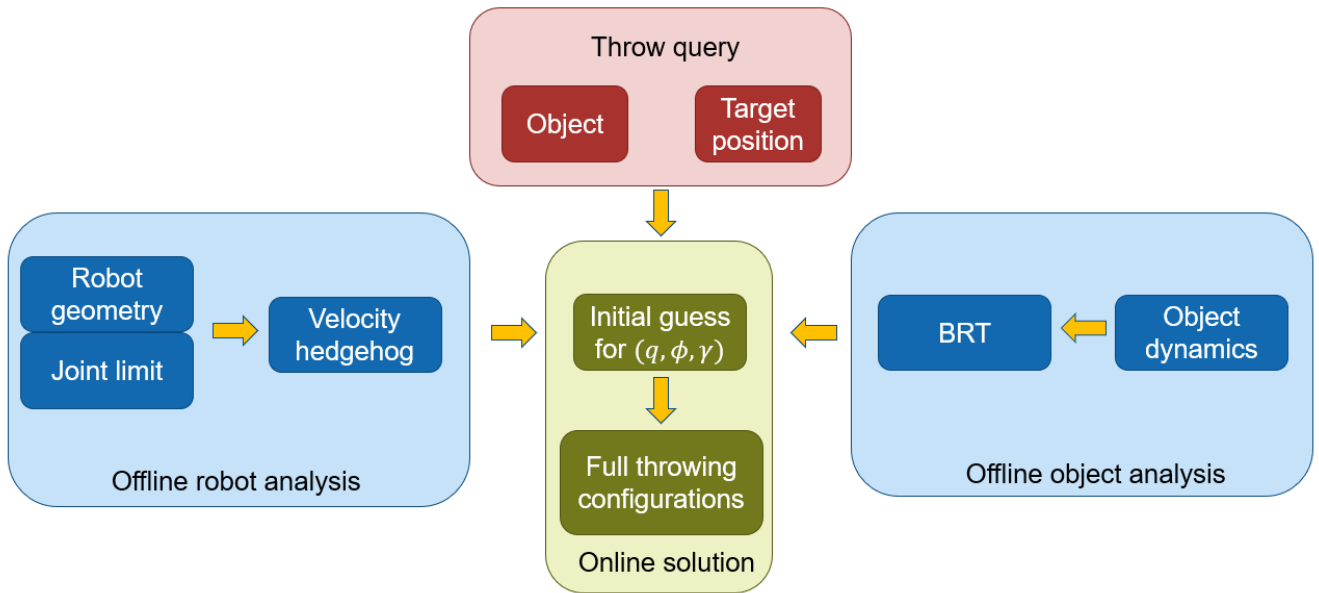


Fig. 6: Framework to obtain throwing configurations. The velocity hedgehog and BRT are obtained as described in Sections II-E and II-C, respectively. The initial guess is obtained by combining them as outlined in Section II-F.

To approximate the BRT, we sample 2160 landing states inside  $\mathcal{X}_l$ , solve the flying dynamics backwards in time for 1 second to get the 2160 flying trajectories, then aggregate all the data points on the trajectories, filter out the data with high velocities that is for sure not feasible be the robot ( $|\dot{r}| > 5.0$ ,  $|\dot{z}| > 5.0$ ), finally yield 75000 throwing configurations in throwing plane coordinates.

### B. Neural network-based BRT representation

From the pipeline of BRT data generation, we also obtain 80000 throwing configurations outside BRT via sample 2160 landing configurations outside target set  $\mathcal{X}_l$ . These is treated as negative data, together with the above 75000 positive data as the dataset to train a neural network-based BRT representation. To make sure the BRT is a dense 4D volume, we augment the dataset by copying and shifting the data in  $r$  and  $z$  directions. We split the dataset to 70% training data and 30% testing data. The neural network contains 4 layers with size 4-64-64-2 and sigmoid activation function. We use Adam optimizer for training, and the training takes 2 minutes for 10 epochs. The trained neural network achieved 97% classification accuracy on the test set. The decision boundary  $f_{BRT}(x) = 0$  of the neural network is illustrated in Fig. 7.

### C. Robot velocity hedgehog generation

We denote a scalar array consists of values from  $a$  to  $b$  with equal interval  $c$  as  $[a : c : b]$ . We discretize the throwing height with grid  $Z = [0 : 0.05 : 1.2]$ , the throwing yaw angle with grid  $\Phi = [-90 : 15 : 90]$  and the throwing pitch angle with grid  $\Gamma = [20 : 5 : 70]$ , yield 3289 cells of robot throwing candidates. The robot velocity hedgehog is generated using Algorithm 1 with 1 million joint state samples.




Fig. 7: Object flying trajectory family (rainbow) and boundary of learned neural network-based BRT representation (pink). Decision boundary strictly encompasses the flying trajectories because neural network was trained with data augmented in  $r$  and  $z$  directions.

### D. Trajectory generation towards throwing configuration

Given initial state of the robot  $(q_0, \dot{q}_0)$  and feasible throwing configuration  $(q_d, \dot{q}_d)$ , we use Ruckig [21] to generate jerk-limited time-optimal trajectory to move the robot towards throwing configuration. Then the robot opens the gripper to release the object at the end of the trajectory.

## IV. EVALUATION AND DISCUSSION

### A. Computational Efficiency

Here we summarize the computation time of offline stage in Table III and online stage in Table IV.

TABLE III: Offline computation time.

Algorithm stage	Running time
Velocity hedgehog generation	7 hours
BRT data generation	1 second
Training BRT neural network	2 minutes

TABLE IV: Online computation time.

Algorithm stage	Running time per solution
Initial guess	20 $\mu s$
Full throwing configuration	250 $\mu s$
Trajectory generation	200 $\mu s$
Overall	500 $\mu s$

Although the algorithm takes a long time to generate velocity hedgehog, it only need to run once for every robot with the same geometry. In terms of new object, given its flying dynamics, the algorithm is efficient (2 minutes) to be ready to throw.

With our proper separation of offline computation and on-line computation, the algorithm is quick to generate throwing motion given throwing query.

### B. Reliability of solver

Our throwing geometry analysis and robot velocity analysis results in highly reliable solver, which provides batches of throwing configurations. In comparison, generic solvers have a hard time to find one feasible configuration, and they have to be carefully tuned and be provided with good initial guess by human. Table V demonstrates that our solver provides ample of throwing configurations for various target heights. In Fig. 8, we show three throws towards targets with different heights. Furthermore, for the same target box, the solver is able to find out qualitatively different solutions, as shown in Fig. 9.

### C. Throw success rate in simulation

We systematically assessed our algorithm’s ability to generate a valid throw for a set of target height queries in the range of  $Z = [-1.2 : 0.1 : 0.9]$ . The algorithm generated in total 10366 throwing trajectories from the same initial configuration of the robot. We simulated the generated throwing trajectories for validation. As a result, 10306 (99.4%) throws were successful throw. As we observe the failure in 0.6% of the cases, they are all due to collision: ball collided either with robot or box while flying. Considering the fact that we do not handle collision avoidance explicitly, this is anticipated.

### D. On throwing configuration selection

The proposed algorithm outputs batches of throwing configurations, which are valid (object will fly into the target box) and feasible (throwing state satisfies robot joint position and velocity constraints). However, at this moment, we haven’t determine the final throwing configuration from the solution set, which is to be executed by the robot. Selecting

TABLE V: Solver feasibility.

Target height	$z = -0.2$	$z = 0.0$	$z = 0.2$	$z = 0.5$
Number of solutions	11955	10504	7118	2422




Fig. 8: Throw to three different target heights.

throwing configuration should consider broader context or higher level specifications of the use cases. For example, if one seeks for the time-optimal solution, the best throwing configuration would depend on the initial configuration of the robot. If one seeks for a solution with lower risk, the throwing configuration selection procedure would bias towards small-range throws. While here in this paper, we emphasize on the guaranteed validity and feasibility of the solver, throwing configuration selection is out of the scope of this work.

### E. Adaptive throwing

Our method also enables adaptive throwing. As shown in Fig. 10, while the robot was moving towards the originally planned time-optimal throwing configuration (red right), it suddenly got disturbed to a different configuration during execution (from red left to green left). There are two strategies to handle this disturbance:

- Recalculate the trajectory from the disturbed configuration to the original throwing configuration;
- Discard the originally planned throwing configuration, sample 100 initial guesses obtained before execution from Algorithm 2, find a new throwing configuration that is closest in time to the disturbed configuration.

We summarize the computation time and duration of obtained throwing trajectory in Table VI. As shown in Table VI, strategy (b) trades 15 ms computation time for 670 ms task execution time, which is preferable. As a result, the robot switched to another valid throwing configuration (green right) that is actually closer to the disturbed configuration.

## V. CONCLUSION

In this paper, we propose a solution to mobile manipulator throwing. The merits of our method are:

- Ability to handle arbitrary rigid object given its flying dynamics;
- Ability to handle arbitrary target position within the robot kinematic limits;
- Capacity for adaptive throwing.

We validate our framework in simulation and the algorithm is shown to be efficient and reliable. In the future, we




Fig. 9: Two qualitatively different throwing configurations, note the different heights of the end-effectors.




Fig. 10: Adaptive throw as the reaction to disturbance. See Subsection IV-E for detailed explanation.

will implement our algorithm on the real robot platform to test its effectiveness. Moreover, we plan to investigate the influence of modeling error in flying dynamics and of trajectory tracking accuracy on the real throwing set-up.

#### ACKNOWLEDGMENT

The authors acknowledge the H2020 ICT-46-2020 EU project DARKO for supporting this work.

#### REFERENCES

- [1] M. T. Mason and K. M. Lynch, "Dynamic manipulation," in *Proceedings of 1993 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'93)*, vol. 1. IEEE, 1993, pp. 152–159.
- [2] H. Khurana, M. Bombile, and A. Billard, "Learning to hit: A statistical dynamical system based approach," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 9415–9421.
- [3] N. Correll, K. E. Bekris, D. Berenson, O. Brock, A. Causo, K. Hauser, K. Okada, A. Rodriguez, J. M. Romano, and P. R. Wurman, "Analysis and observations from the first amazon picking challenge," *IEEE Transactions on Automation Science and Engineering*, vol. 15, no. 1, pp. 172–188, 2016.
- [4] A. Zeng, S. Song, J. Lee, A. Rodriguez, and T. Funkhouser, "Tossing-bot: Learning to throw arbitrary objects with residual physics," *IEEE Transactions on Robotics*, vol. 36, no. 4, pp. 1307–1319, 2020.
- [5] S. Kim and A. Billard, "Estimating the non-linear dynamics of free-flying objects," *Robotics and Autonomous Systems*, vol. 60, no. 9, pp. 1108–1122, 2012.

TABLE VI: Computation time of different strategies to handle disturbance.

Strategy	Computation time	Trajectory duration	Total
(a)	0.2 ms	1800 ms	~1800 ms
(b)	15 ms	1134 ms	~1150 ms

- [6] K. M. Lynch and M. T. Mason, "Dynamic nonprehensile manipulation: Controllability, planning, and experiments," *The International Journal of Robotics Research*, vol. 18, no. 1, pp. 64–92, 1999.
- [7] T. Senoo, A. Namiki, and M. Ishikawa, "High-speed throwing motion based on kinetic chain approach," in *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2008, pp. 3206–3211.
- [8] W. Mori, J. Ueda, and T. Ogasawara, "1-dof dynamic pitching robot that independently controls velocity, angular velocity, and direction of a ball: Contact models and motion planning," in *2009 IEEE International Conference on Robotics and Automation*. IEEE, 2009, pp. 1655–1661.
- [9] J. Kober, K. Muelling, and J. Peters, "Learning throwing and catching skills," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 5167–5168.
- [10] Y. Gai, Y. Kobayashi, Y. Hoshino, and T. Emaru, "Motion control of a ball throwing robot with a flexible robotic arm," *International Journal of Computer and Information Engineering*, vol. 7, no. 7, pp. 937–945, 2013.
- [11] A. Pekarovskiy and M. Buss, "Optimal control goal manifolds for planar nonprehensile throwing," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 4518–4524.
- [12] O. Taylor and A. Rodriguez, "Optimal shape and motion planning for dynamic planar manipulation," *Autonomous Robots*, vol. 43, no. 2, pp. 327–344, 2019.
- [13] H. Miyashita, T. Yamawaki, and M. Yashima, "Control for throwing manipulation by one joint robot," in *2009 IEEE International Conference on Robotics and Automation*. IEEE, 2009, pp. 1273–1278.
- [14] F. Lombai and G. Szederkényi, "Throwing motion generation using nonlinear optimization on a 6-degree-of-freedom robot manipulator," in *2009 IEEE International Conference on Mechatronics*. IEEE, 2009, pp. 1–6.
- [15] A. Sintov and A. Shapiro, "A stochastic dynamic motion planning algorithm for object-throwing," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 2475–2480.
- [16] Y. Zhang, J. Luo, and K. Hauser, "Sampling-based motion planning with dynamic intermediate state objectives: Application to throwing," in *2012 IEEE International Conference on Robotics and Automation*. IEEE, 2012, pp. 2551–2556.
- [17] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE transactions on Robotics and Automation*, vol. 12, no. 4, pp. 566–580, 1996.
- [18] J. J. Kuffner and S. M. LaValle, "Rrt-connect: An efficient approach to single-query path planning," in *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065)*, vol. 2. IEEE, 2000, pp. 995–1001.
- [19] T. B. Schön, A. Wills, and B. Ninness, "System identification of nonlinear state-space models," *Automatica*, vol. 47, no. 1, pp. 39–49, 2011.
- [20] E. Coumans and Y. Bai, "Pybullet, a python module for physics simulation for games, robotics and machine learning," <http://pybullet.org>, 2016–2021.
- [21] L. Berscheid and T. Kröger, "Jerk-limited real-time trajectory generation with arbitrary target states," *Robotics: Science and Systems XVII*, 2021.