Clutter-Aware Spill-Free Liquid Transport via Learned Dynamics

Ava Abderezaei*, Anuj Pasricha, Alex Klausenstock, and Alessandro Roncone

Abstract— In this work, we present a novel algorithm to perform spill-free handling of open-top liquid-filled containers that operates in cluttered environments. By allowing liquid-filled containers to be tilted at higher angles and enabling motion along all axes of end-effector orientation, our work extends the reachable space and enhances maneuverability around obstacles, broadening the range of feasible scenarios. Our key contributions include: i) generating spill-free paths through the use of RRT* with an informed sampler that leverages container properties to avoid spill-inducing states (such as an upside-down container), ii) parameterizing the resulting path to generate spill-free trajectories through the implementation of a time parameterization algorithm, coupled with a transformer-based machine-learning model capable of classifying trajectories as spill-free or not. We validate our approach in real-world, obstacle-rich task settings using containers of various shapes and fill levels and demonstrate an extended solution space that is at least 3x larger than an existing approach.

I. INTRODUCTION

The integration of robots is rapidly expanding across various domains, spanning from industrial automation to healthcare, agriculture, and domestic services. The ability to manipulate objects is fundamental for robots to perform a variety of tasks in these diverse domains [1], [2]. Among the different categories of objects that robots encounter, liquids present a unique challenge due to their complex behavior. Despite the importance of robotic fluid manipulation in various applications such as handling hazardous materials in laboratories to facilitating assistive feeding technologies designed for differently-abled individuals, research in robotic fluid manipulation is under-explored and challenging for several reasons.

Fluid manipulation is challenging due to the intricate and nonlinear nature of fluid dynamics. Accurately modeling fluid dynamics is heavily task-dependent and requires substantial computational resources, hindering real-time control [6], [7]. Existing approaches tackle these challenges by simplifying the complex fluid dynamics into more computationally efficient linearized pendulum or mass-spring-damper models. However, these methods struggle in cluttered scenes due to several reasons. These linearized models restrict the motions to pendulum-like or mass-spring-damper-like behavior [3]–[5], [8]–[10]. Additionally, these approaches impose restrictions on the container's orientation by setting fixed or limited Euler angle ranges throughout the trajectory [3], [5]. Moreover, these methods rely on constant monitoring of the

* Corresponding author.

All authors are with the Department of Computer Science, University of Colorado Boulder, 1111 Engineering Drive, Boulder, CO USA. This work is partly supported by NSF #2222952/2953. `firstname.lastname@colorado.edu`
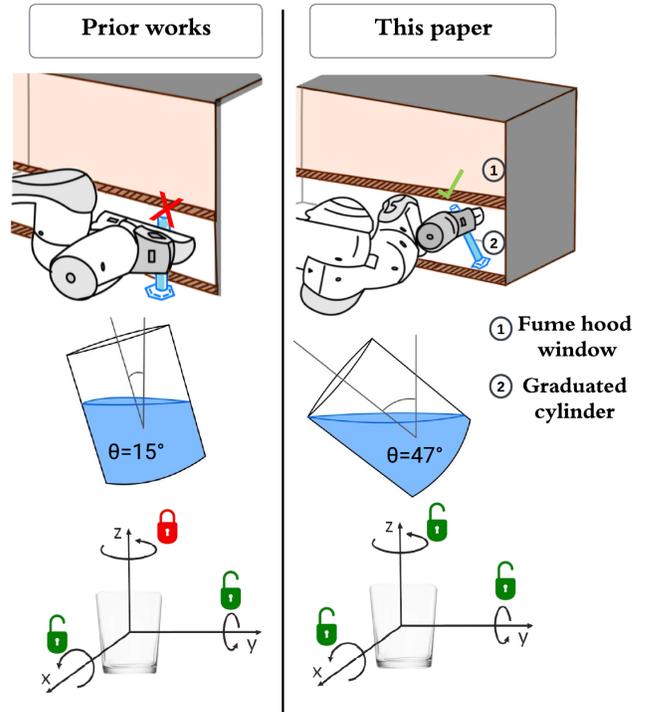
Fig. 1: Motivated by the need for robots to perform in cluttered environments such as a chemistry setup where robots must transport chemistry vessels inside a fume hood (*top row*), this paper targets the challenge of enabling spill-free liquid transport while avoiding obstacles. Our approach introduces two key enhancements over existing methods: (i) increasing the allowable tilt angle to the maximum quasi-static spill-free tilt angle, rather than the limited angles used previously [3], [4] (*second row*); (ii) enabling rotation about all axes throughout the trajectory, instead of maintaining fixed yaw or orientation (*third row*) [3], [5].

liquid surface to execute trajectories, which can be obstructed by opaque fumes produced by the liquid or by obstacles present in cluttered scenes [4], [11], [12]. Consequently, the action space is over-constrained, and the robot can only perform limited motions, inhibiting its ability to maneuver around obstacles in cluttered environments.

To address these limitations, we propose Spill-Free RRT* (*SFRRT**)[1], for spill-free manipulation of open-top liquid-filled containers, particularly in cluttered environments. Our method enables these advantages (Fig. 1): (i) *Increased action space*, so that the robot can operate effectively in cluttered scenarios. We do this by training a model to implicitly learn

liquid dynamics to capture the correlation between container trajectory and spillage risk, (ii) *No additional information needed.* Our method relies solely on the robot and the container shape to operate and does not require constant monitoring of the liquid surface.

To achieve this, SFRRT* generates spill-free trajectories through these key contributions: (i) it first uses the container shape to inform the generation of a spill-free path via RRT*. By considering the container geometry, the algorithm calculates the maximum tilt angle beyond which the container will spill and constrains the path search accordingly (Section III-B). (ii) Subsequently, our method generates a spill-free trajectory from this path by using a time-parameterization algorithm in conjunction with a transformer-based neural network classifier. The neural network, trained on a dataset of container trajectories, predicts whether a given trajectory is spill-free or not, allowing SFRRT* to set the trajectory's jerk to avoid spills (Section III-C), (iii) Lastly, to train the transformer-based neural network, we produce a dataset of 700 container trajectories collected from individuals handling different container scenarios.

II. RELATED WORK

Transporting open-top, liquid-filled containers requires the modeling of liquid motion along with the planning of time-parameterized motion paths. A common approach in liquid modeling relies on using Computational Fluid Dynamics (CFD) simulations to achieve spill-free motions by designing controllers to suppress or limit slosh effects, i.e. liquid vibration effects [13], [14]. Despite the accuracy of CFD simulations, they are computationally demanding and require considerable time to converge, making real-time control applications in robotics unfeasible [15]–[17].

To overcome the computational cost associated with CFD, prior work considers two approaches for modeling the sloshing dynamics inside a container: a spherical pendulum [3], [4], [11] and a two Degrees-of-Freedom (2DoF) mass-spring-damper system [5], [10]. These models are inherently non-linear and computational constraints necessitate linearization for practical implementation. However, linearization introduces several kinematic constraints that restrict robots' range of motions, thereby limiting their ability to navigate through cluttered scenes.

Spherical pendulums. These approaches present controllers that are designed to prevent liquid spillage [4], [11]. They limit the cup's freedom to only one axis of orientation, while also relying on continuous sensor input to monitor the liquid surface for modeling slosh dynamics. In contrast, our approach dispenses with the need for continuous liquid surface observation and allows the container to retain all its degrees of freedom. Additional work assumes that a higher-level motion planner already provides an obstacle-free trajectory [3]. A quadratic program is then employed to optimize the trajectory within a linearized spherical pendulum dynamics model. While the method eliminates the need for continuous liquid surface monitoring, the linearization restricts object yaw to be constant and allows only tilting

the container up to $15°$ to maintain acceptable linearization accuracy. In contrast, our approach enables motion along all three rotation axes and allows the container to be tilted to significantly larger angles, up to the point of spillage, by modeling the quasi-static fluid dynamics.

Mass-spring-dampers. In these approaches, a sloshing-height evaluation method is employed for time-optimal trajectory planning [5]. However, this method is tailored for cylindrical containers subject to 1D and 2D motions. This sloshing estimation technique is extended to accommodate 3D motions but the model assumes that the container remains upright during motion [10]. In contrast, our approach allows the container to be tilted on each axis and is compatible with various container shapes.

Finally, a notable limitation of both mass-spring-dampers and pendulum-based slosh dynamics models is their inability to navigate around obstacles consistently. These models rely on maintaining pendulum-like or mass-spring-damper-like motions for spill-free movement, and obstacles can disrupt these motions, rendering these approaches ineffective.

Other Approaches. Among other approaches, [18] presents an effective method for the rapid, spill-free transport of containers. However, it explicitly states that this technique is not intended for the transportation of liquids, as it does not incorporate constraints to inhibit the accumulation of energy within the cup, which could subsequently lead to sloshing. Additionally, the method proposed in [19] aims to accurately learn slosh dynamics through machine learning, taking into account various liquid properties including thermodynamics. However, its applicability to motion planning with robots has not been evaluated, and its practical implementation necessitates continuous accurate liquid surface monitoring.

Overall, our approach outperforms prior methods by implicitly learning the dynamics of liquid spillage. It does not impose kinematic constraints that would limit the robot's range of motion such as setting a fixed container orientation, thereby enabling its maneuverability in cluttered environments. Moreover, it does not require constant observation of the liquid surface to generate motions and it is not restricted to a particular container shape.

III. METHODS

This section introduces SFRRT*, a motion planner that generates spill- and collision-free trajectories (Fig. 2). SFRRT* accomplishes this through two key contributions: a) the informed sampler, mathematically formalized based on container geometry, to efficiently explore the space of spill-free robot states (Section III-B), b) a transformer-based classifier that models the underlying nonlinear dynamics of liquid motion and classifies whether a planned trajectory is spill-free or not (Section III-C).

A. Mathematical Formulation of the Spillage Scenario

A challenge in spill-free trajectory generation is that the volume of the robot configuration space that contains spill-free trajectories is small and therefore inefficient to explore with random sampling alone [20]. This dictates the need
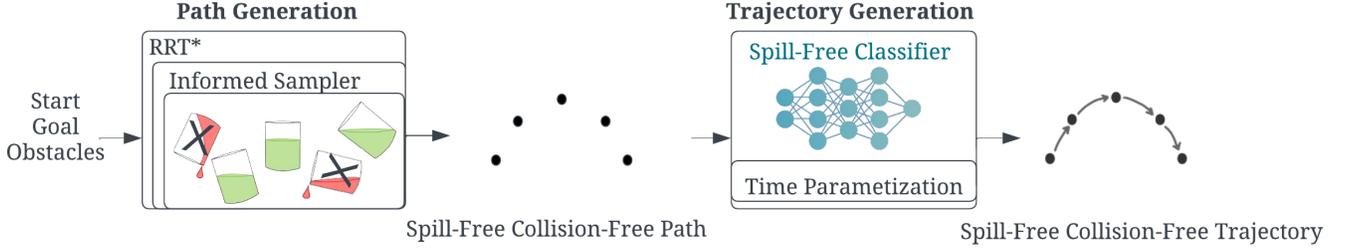
Path Generation

RRT*

Informed Sampler

Start
Goal
Obstacles

Spill-Free Collision-Free Path

Trajectory Generation

Spill-Free Classifier

Time Parametization

Spill-Free Collision-Free Trajectory

Fig. 2: **System diagram**. The inputs are the start, goal, and obstacles, and the output is the trajectory. First, (i) a spill- and collision-free path is generated using RRT* and the informed sampler. The informed sampler utilizes the container properties to avoid sampling states such as an upside-down container. Then, (ii) a spill- and collision-free trajectory is generated using a time parameterization algorithm alongside the spill-free classifier model. The SFC model classifies trajectories as either spill-free or not and is utilized to set the jerk of the trajectory, ensuring the generation of a spill-free trajectory.

for an informed approach to explore dynamically-valid robot states. However, developing a model for combined robot and object dynamics, specifically spillage dynamics, is computationally expensive, task-specific, and impractical for real-world control. To address this, we need to mathematically formalize how container properties inform the construction of the *spill-free space*. For example, for a given container that contains a greater volume of liquid, it cannot be tilted as much as when it contains less liquid. Additionally, in scenarios where containers have the same liquid volume, taller containers can be tilted more than shorter ones.

To sample spill-free configurations from the *spill-free space*, we first calculate the maximum tilt angle that a container can withstand without spilling. Then, we sample states ensuring that their orientation remains within this range. Specifically, under quasi-static conditions, we consider the scenario that involves tilting a container to the point where it cannot retain all its content, i.e., to the angle θ_{max}, such that any $\theta_t \leq \theta_{max}$ would be spill-free and any $\theta_t > \theta_{max}$ would cause spillage (Fig. 3). Quasi-static conditions are assumed because they simplify analytical modeling without requiring expensive computational fluid dynamics calculations.

To formalize the spillage scenario, we model the container as a frustum of a cone, which accurately represents most common container shapes and can provide an upper bound on the maximum tilt angle of containers with other shapes, such as round bottom flasks or curvy wine glasses (Fig. 3a). Specifically, we simplify curved containers to the largest possible frustums of cones that can fit inside the original shape, leading to predicting spillage at a lesser tilt angle than the real container's capabilities. This guarantees a θ less than θ_{max}, ensuring the generation of spill-free states. Lastly, we simplify the frustum of the cone representation from 3D to 2D, resulting in a trapezoid. This simplification also allows us to represent planar symmetric containers such as rectangular prisms and can be directly extended to 3D.

To solve for the maximum tilt angle, we assume that the information we have from the container is the bottom circle radius r_b, the top circle radius r_u, the cup height h_c, and the water height h_w (Fig. 3a). With this information, we can derive other relevant variables for calculating the maximum tilt angle. As shown in Fig. 3, two spillage scenarios could

occur: (i) one where the liquid forms a trapezoid and a triangle (Fig. 3b), and (ii) another where the liquid only forms a triangle (Fig. 3c). This happens because static liquid stays parallel to the horizontal ground and depending on the liquid level, these cases can occur.

To calculate the maximum tilt angle, for both spillage cases, we leverage the key insight that the total amount of liquid remains constant before and after tilting the container. By equating the liquid area in the initial and tilted configurations, we derive an algebraic formula to compute the angle. The resulting equations, Eq. (1) for case I and Eq. (2) for case II, provide the respective maximum tilt angles.

Maximum tilt angle for case I:

$$\theta_{max} = tan^{-1}\left[\frac{h_c(r_u - r_{w'})}{(r_u - r_b)(r_u + r_{w'})}\right] \quad (1)$$

where $r_{w'} = \frac{h_c r_u r_b - h_c r_2 r_b + h_w r_2 r_w + h_w r_2 r_b}{h_c r_2 + h_c r_b}$.

Maximum tilt angle for case II:

$$\theta_{max} = tan^{-1}\left[\frac{h_c}{r_u - r_b + \left(\frac{2h_w(r_w + r_b)}{h_c}\right)}\right] \quad (2)$$

where $r_w = r_u - [(\frac{r_u - r_b}{h_c})(h_c - h_w)]$.

The next section explains how these equations are incorporated in the sampling-based path planner to inform the construction of a spill-free path.

B. Informed Geometric Path Planning for Spill-Free Motions

Having formalized the maximum tilt angle θ_{max}, it is now used to guide the exploration of states that avoid spills, such as an upside-down container, or a container that is tilted beyond its capacity (i.e., $> \theta_{max}$). To achieve this, we implemented an informed sampler to sample robot configurations where the container orientation is less than the spill threshold θ_{max}. Orientation refers to the container's Euler angles, and it is sampled such that the angle between the vertical axis and the axis about which the container is symmetric, is less than θ_{max}. This informed sampler is then incorporated into a geometric RRT* path planner, to generate spill-free paths. RRT* is chosen due to its asymptotic optimality properties. By preferring transitions with minimal cost, RRT* generates smooth motion profiles that maintain orientation stability

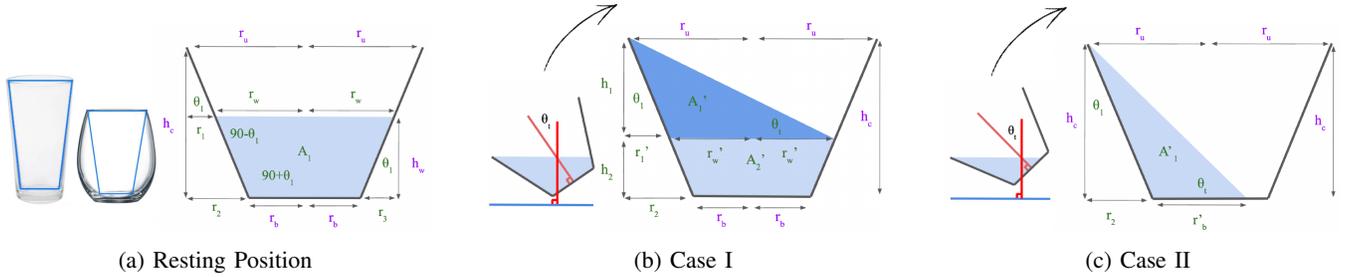(a) Resting Position (b) Case I (c) Case II

Fig. 3: (a) depicts modeling containers as frustum of cones. Two cases occur when tilting the container as illustrated in (b) and (c). Case I is when the liquid forms both a trapezoid and a triangle shape. Case II shows when the liquid forms only a triangular shape. In both cases, θ_t represents the maximum tilt angle being calculated.
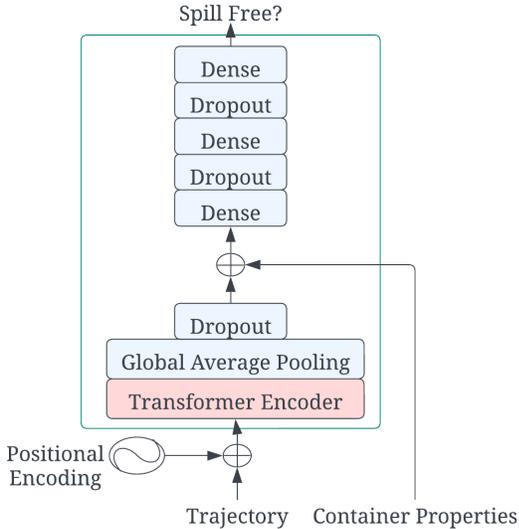

Fig. 4: **Spill-Free Classifier Model Architecture** The input is the trajectory and the container properties. The output is a boolean stating whether the trajectory is spill-free or not.

[21]. The next section describes how the path generated by RRT* is parameterized to ensure a spill-free trajectory.

C. Time Parameterization To Create Spill-Free Trajectories

While RRT* generates spill-free paths, to ensure the generation of a spill-free trajectory, we must carefully parameterize the path such that the velocity, acceleration, and jerk, do not cause spillage. To do this, it is necessary to first model the risk of spillage with respect to the trajectory. To capture this relationship, we employed a transformer-based machine learning model trained to classify whether a container's trajectory is spill-free or not. This model is termed the *Spill-Free Classifier* (SFC), and utilizes the transformer architecture, as depicted in Fig. 4.

The input to SFC is the container trajectory and container properties, where the trajectory is the position and orientation over time and the properties represent the container's top and bottom radius, container height, and liquid fill level. Furthermore, as shown in Fig. 4, the trajectory is first passed through the transformer encoder layer. Then, the container properties are concatenated with it. The SFC model employs the transformer architecture because its input is

sequence data (the trajectory) and transformers have become the leading architecture for handling sequential data due to their efficient parallel training capabilities and long context window [22].

Lastly, to train the SFC model, we utilized a motion capture system to collect container trajectories (comprised of position and orientation over time) performed by individuals across different container types and liquid volumes (Fig. 6). We collected a total of 700 trajectories, including 655 human-performed trajectories and 45 trajectories from the Panda robot. These trajectories encompassed a variety of paths from holding cups upside down and following zig-zag patterns to transporting cups in straight or typical paths between two points.

Algorithm 1: Spill Free Time Parametization (SFTP)

1: **Input:** Waypoints $\mathcal{W}$, Max Jerk J_{max}, Max Acceleration A_{max}, Max Velocity V_{max}, Jerk Decrease Factor $rate$, Max Iterations N,
2: **Output:** Trajectory $\mathcal{T}$
3: $P_s, P_f \leftarrow W[0], W[length(W) - 1]$
4: $V_s, V_f, A_s, A_f \leftarrow 0$
5: **for** $i = 1...N$ **do**
6:     $\mathcal{T} \leftarrow$ Ruckig($\mathcal{W}, J, V, A, P$)
7:     $\mathcal{T}_{fk} \leftarrow ForwardKinematics(P)$
8:     $spillFree \leftarrow$ SpillFreeClassifier($\mathcal{T}_{fk}$)
9:     **if** $spillFree$ **then**
10:         **return** $\mathcal{T}$
11:     **else**
12:         $J_{max} \leftarrow \frac{J_{max}}{rate}$
13:     **end if**
14: **end for**

Having modeled the relationship between the container trajectory and spillage risk, we leverage SFC to generate spill-free trajectories. We introduce the *Spill-Free Time Parameterization* (SFTP), which guarantees the creation of spill-free trajectories by utilizing a time-parameterization algorithm alongside the SFC model. The key idea is to find the maximum jerk that does not cause spillage. To achieve this, SFTP works by iteratively (line 5, Algorithm 1) reducing the maximum jerk of the trajectory (line 12, Algorithm 1), querying SFC after each iteration (line 8, Algorithm 1) until SFC predicts the resulting trajectory is spill-free. Jerk constraints are especially critical in preventing sudden changes in acceleration that could cause spillage. SFTP only modifies the jerk parameter since it directly affects acceleration and
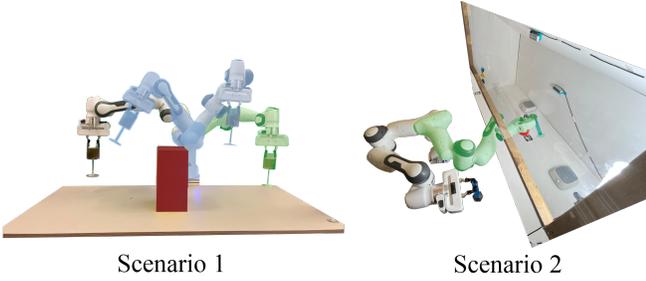
Scenario 1   Scenario 2

Fig. 5: The two scenarios from Table I: The original colored robot for start position, green for goal. Blue indicates intermediate robot states.
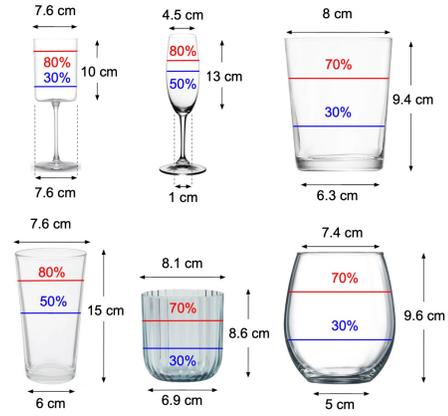

Fig. 6: The entire set of containers used to train SFC. Notably, the top three containers, identified as the wine glass, flute glass, and basic glass, are used in our experiments.

velocity. Lastly, for time parameterization the Ruckig algorithm [23] is used since it can enforce constraints on higher-order derivatives including jerk, acceleration, and velocity.

In summary, to generate a spill-free collision-free trajectory (Fig. 2), first, the container shape and fill level are leveraged to inform the generation of a path using RRT*. The spill-free classifier (SFC) model is then used alongside the Ruckig time-parameterization algorithm to ensure the generation of a spill-free trajectory. We refer to this combined process of generating the path and the trajectory, the *spill-free RRT** (SFRRT*) algorithm.

IV. EVALUATION

We evaluate our spill-free collision-free trajectory planner, SFRRT*, on the 7DoF Franka Emika Panda robotic arm. We design two scenarios and utilize three different containers with different fill levels (Fig. 6, top row) to quantify the performance of SFRRT*. These scenarios are designed to encapsulate real-world use cases of our motion planner, specifically cases that exhibit a diverse range of robot motions: Scenario 1 (Fig. 5, left) involves the robot transporting a drink while avoiding the obstacles on the table; and, Scenario 2 (Fig. 5, right) involves the robot in a chemistry lab transporting a container while avoiding collisions with the fume hood. Additionally, the containers (wine, flute, and basic glasses) are picked based on their diverse geometries, while the fill levels, ranging from 30% to 80% are chosen to cover tasks with varying risks of spillage. The objective across our evaluation tasks is to move these containers from the start to the goal states without spilling the contents of the container being transported. Each experiment, as enumerated in Table I, is repeated five times. We measure the success rate of SFRRT* (Section IV-B) and compare its obstacle avoidance performance (Section IV-C) against the SpillNot method from [3]. We also evaluate the accuracy of the SFC model on containers that were not part of the training set, which are also used in the video demo (Fig. 9) to display a real-world chemistry lab procedure (Section IV-D). Lastly, we present an ablation study that compares modified versions of the SFRRT* (Section IV-E).

A. Implementation Details

The performance of SFRRT* depends on the spill constraint enforced during path planning. This constraint re-

stricts the container's orientation by imposing maximum allowable tilt angles to prevent spillage. As described previously, these angles are calculated using Eq. (1) and Eq. (2). Table I, last column, presents the maximum tilt angles for the containers used in the experiments. To train the spill-free classifier (SFC) model, we deployed it on TensorFlow and optimized it using a binary cross-entropy loss function for 400 epochs with a batch size of 32. SFC achieved an accuracy of 94%. Detailed training parameters are available on the open-source code.

Experiment	Scenario	Container	Fill Level	θ_{max}
1	1	Wine glass	80%	39°
2	1	Wine glass	30%	65°
3	1	Flute glass	80%	49°
4	1	Flute glass	50%	72°
5	2	Basic glass	70%	35°
6	2	Basic glass	30%	62°

TABLE I: **Experiments designed for assessing SFRRT*:** Each is repeated 5 times using a Panda robotic arm.

Finally, several key parameters are set in SFTP (Algorithm 1). The bounds on velocity, acceleration, jerk and are set according to the robotic arm's kinematic constraints [24]. Moreover, the jerk decrease factor $rate$ is set to 2, meaning when SFTP fails to find a spill-free trajectory, J_{max} will be halved and SFTP will reevaluate for spillage under the updated constraint. This $rate$ selection balances execution time and planning time: a smaller rate would extend the time to find a spill-free trajectory, while a larger rate could lead to unnecessarily slow trajectories due to large jerk constraints.

B. Success Rate

To measure the success rate of the experiments presented in Table I, each experiment was repeated five times and the result is presented in Table II. Success is defined as the absence of spillage and environmental collisions during the executed motion. Overall, we achieved a 100% success rate for all experiments except experiment 5. For experiment 5, spill-free trajectories were accomplished in four out of five trials. Experiment 5 faced challenges since it used a container with a maximum tilt angle of 35°, smaller than
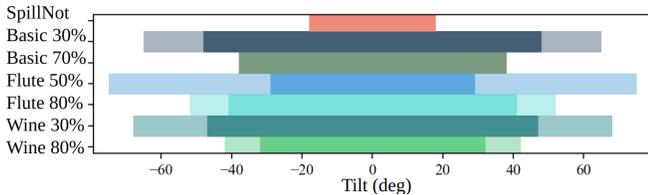
Fig. 7: Maximum tilt angles executed by the robot across all 30 trials from Table I. The Y Axis labels correspond to the experiments, for example, "Basic 30%" represents experiment 6. The shaded region represents the allowable maximum tilt angle (Table I, last column). The solid-colored line corresponds to the executed maximum tilt angle. Spill-Not's [3] maximum tilt angle is included for comparison.

others, limiting the trajectory planning options and increasing spillage risk. Additionally, the SFC model, with its 94% accuracy, had difficulty adapting to this constrained scenario, contributing to the lower success rate.

Experiment	1	2	3	4	5	6
Success	5/5	5/5	5/5	5/5	4/5	5/5

TABLE II: Success rates for different experiments.

C. Avoiding Obstacles

A key capability of SFRRT* is its ability to navigate around obstacles in the environment while ensuring the container does not spill. To analyze this ability, we focus on the end effector orientation allowable ranges. In SFRRT*, the container orientation is constrained by the maximum tilt angle of the container, i.e. it can generate spill-free trajectories, so long as the container is not tilted past the maximum tilt angle. By enabling the robot to tilt the container to higher angles, it can more flexibly maneuver in cluttered spaces, such as a fume hood, which may require tilting chemistry vessels to fit them inside.

To illustrate this expanded ability, Fig. 7 shows the maximum possible and maximum utilized tilt angles of all 30 trials of Table I. In every experiment, SFRRT* utilizes a larger tilt range than the SpillNot [3]. Notably, SpillNot's tilt angles are not contingent on the fill level and are always limited to a maximum of 15° to maintain low linearization error ([3], Fig 4). We also chose [3] for comparison because, unlike other methods, it does not rely on sensor-based slosh modeling, and allows translational motion in all axes.

As seen in Fig. 7, using SFRRT* the robot achieves a maximum tilt angle of 45° when transporting the basic container with a 30% fill level (Table I, experiment 6), surpassing the SpillNot method's 15°. The range of orientation for this specific container is 3x larger than [3]. Crucially, the maximum tilt angle for the flute glass with a 70% fill level is 72°, a 4.8x increase in orientation range.

Furthermore, it is important to enable orientation changes along all axes to avoid obstacles. For example, maintaining a fixed pitch angle throughout the trajectory prevents avoidance of obstacles in certain cluttered environments. SFRRT* can generate trajectories where all three axes of orientation are
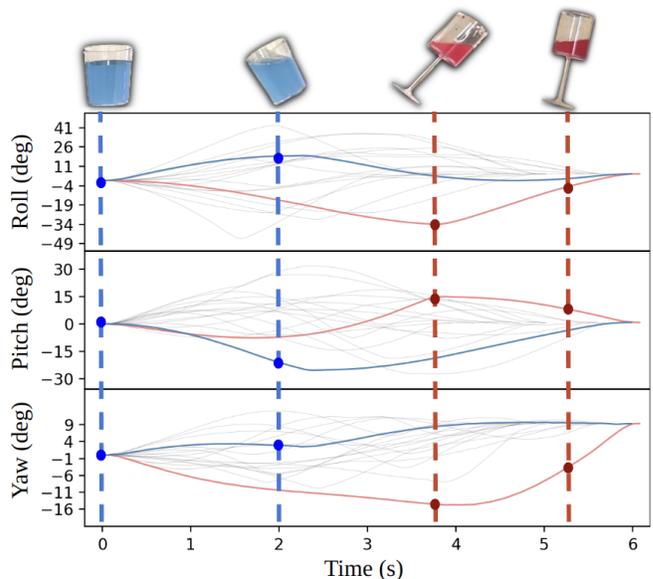

Fig. 8: Container orientation over time across all 30 trials from Table I. The red and blue lines are two trajectories generated for experiment 1 and experiment 6 respectively. Snapshots of the container at specific Euler angles are displayed at the top of the plot. As seen, there are angle changes in each axis of orientation, highlighting the ability of SFRRT* to avoid obstacles.

Fig. 9: The robot transporting a graduated cylinder into the fume hood while avoiding spills and collisions

utilized. Fig. 8 displays the container orientation over all the 30 trajectories generated for Table I. As shown, there are changes in Euler angles across all axes, the range for roll, pitch, and yaw angles are $[-49, 41]$, $[-30, 30]$, and $[-16, 9]$ degrees, respectively. In contrast, the SpillNot method imposes a fixed yaw angle throughout the trajectory [3].

D. SFC: Generalization Capability and Velocity Adaptation

The Spill-Free Classifier (SFC) model is trained using 75% of the 700 trajectories and achieved an accuracy of 94% for the validation set. To evaluate the generalization capability of the SFC, its accuracy was tested on novel containers not present in the training dataset. Specifically, a round bottom flask, beaker, and graduated cylinder were used to gather multiple trajectories using a motion capture system (Table III). For each container type, 12 trajectories were gathered with varying fill levels.

As shown in Table III, SFC demonstrates good generalization capability. It achieves strong performance for the beaker and graduated cylinder, with 11 out of 12 data points accurately predicted. However, its accuracy is lower for the round bottom flask, with only 8 out of 12 data points predicted correctly. This is because the SFC model input for container properties only includes details that capture curvy-

shaped containers by estimating them to a frustum of a cone. By representing these shapes as frustums of cones, the SFC model ensures a conservative yet safe portrayal of container properties. While this cautious approach may occasionally lead to conservative false negatives, it guarantees the generation of spill-free trajectories.

We also explore a different model architecture for SFC. The current architecture (Fig. 4) passes only the trajectory to the transformer layer before concatenating with container properties. We modified this by concatenating properties at each trajectory step before passing it to the transformer. We call this modified architecture SFC_m. SFC_m achieved only 88% accuracy, lower than the 94% of SFC, and demonstrated lower accuracy on novel containers than SFC (Table III).

	Beaker	Graduated Cylinder	Round Bottom Flask
SFC	11/12	11/12	8/12
SFC_m	8/12	9/12	8/12

TABLE III: Model accuracy using unseen containers. SFC demonstrates robust performance in predicting outcomes with previously unseen containers, outperforming SFC_m.

Lastly, SFC is able to take advantage of the container shape to generate trajectories with different velocities. As seen in Fig. 10, considering the same containers, with different fill levels we see that it consistently generates higher velocities for containers with lower fill levels, which are less likely to spill. Moreover, when comparing different containers in the same scenario (as in Table I), those with higher tilt angles consistently achieve higher velocities. For example, in the comparison between the wine glass and the flute glass, the wine glass consistently achieves lower mean velocities. This is attributed to the geometry of the containers. As shown in Table I, the wine glass has a lower maximum tilt angle compared to the flute glass, indicating smaller room to maneuver and a higher susceptibility to spillage. Conversely, the flute glass, with its higher maximum tilt angle consistently achieves higher mean velocities.

E. Ablation Study

In this section, two key elements of SFRRT* are modified to evaluate their impact on performance: the informed sampler and the spill-free classifier (SFC) model.

Task	**SFRRT***	SFRRT*$_u$	SFRRT*$_r$
1	5/5	0/5	1/5
2	5/5	1/5	1/5

TABLE IV: **Success Rate Comparison:** SFRRT*$_u$ uses a uniform sampler, while SFRRT*$_r$ skips the SFC spillage checks and uses random sampling for jerk constraints in time parameterization. Overall, SFRRT* outperforms all variants.
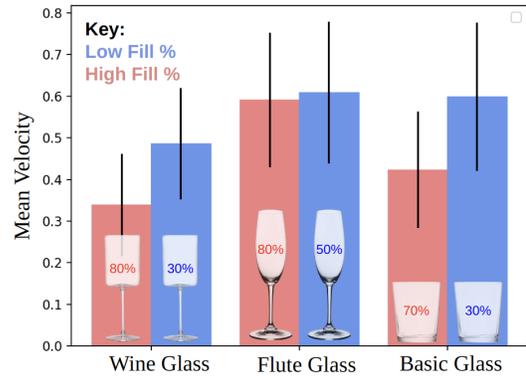

Fig. 10: **Executed Mean Velocity for Containers from Table I Using SFRRT*** As depicted, SFRRT* performs faster trajectories for containers with lower liquid fill levels.

Informed Sampler Variation: SFRRT$_u$.* To avoid sampling states such as an upside-down container, or a container tilted beyond its capacity, the informed sampler samples states within the maximum allowable tilt angles for the container (Table I, last column). To measure its impact, we replace it with a uniform sampler, creating SFRRT*$_u$. Then, experiments 1 and 2 are executed 5 times on the robot, and the success rate is shown in Table IV. The success rate drops significantly since the path generation lacks state sampling constraints, as a result, the time parameterization algorithm struggles to create a spill-free trajectory using this path.

Spill-Free Classifier Variation: SFRRT$_r$.* In SFRRT*, the SFC model assists with time parameterization by selecting the jerk constraints to ensure a spill-free trajectory. To evaluate the impact of SFC, we introduce SFRRT*$_r$ by replacing the SFC model with a random jerk selection. More specifically, a random jerk is selected out of all the jerks that were ever chosen by SFC. The path is then parameterized using that jerk. Table IV shows the success rate for experiments 1 and 2. The observed drop in success rate is due to the random selection of jerk constraints without considering the path itself, in contrast to using the SFC model which takes the path into account when choosing the jerk constraints.

V. Conclusion and Future Work

This work presents SFRRT*, a sampling-based motion planner for the spill-free transport of open-top, liquid-filled containers, particularly in cluttered scenes. SFRRT* distinguishes itself by utilizing implicitly learned spillage dynamics via a transformer-based machine-learning approach. This outperforms existing approaches by increasing the success rate in cluttered scenes. Specifically, SFRRT* enables motions across all axes of end effector orientations using the learned model instead of simplified liquid models. It also permits the container to tilt to its maximum tilt angle, thus enabling flexible maneuvering to avoid obstacles. It also does not need visual monitoring of the liquid, which could get obstructed in cluttered scenes.

However, like many machine learning models, our approach faces a common limitation— the need for a sub-

stantial amount of training data. Moreover, containers are approximated as frustums of cones to estimate the maximum tilt angles. This could lead to conservative estimates for containers with more complex curvy shapes. An alternative would be training a model on image inputs to predict tilt angles more precisely. Lastly, circular upside-down motions (loop-de-loop) can avoid spilling if executed with adequate speed. Still, SFRRT* cannot support generating them because it constrains container orientations to be less than the maximum tilt angle. This limitation stems from the fact that SFC is trained on data collected from humans, who typically exhibit caution when handling liquids. Consequently, the model lacks training data to classify loop-the-loop motions as spill-free.

For future work, we are interested in improving the Spill-Free Classifier model to estimate spillage probability rather than solely classifying it, thus allowing for a more precise time-parameterized trajectory. Furthermore, the path planning phase currently only constrains the orientation range within the maximum tilt angle. However, if the sampled orientation in the resulting path approaches the maximum tilt angle, the generated trajectory will be very slow. To address this, we plan to analyze how the sampled orientation affects velocity, acceleration, and jerk during path planning. We aim to generate a time-optimal parameterized path by considering these kinematic factors in the path-planning phase. Moreover, our approach is versatile and can be expanded to a variety of applications. We plan to extend our method to include spillage dynamics for granular materials such as sand and sugar, to transport a tablespoon of these materials without spillage. Furthermore, our method is adaptable to different path planners and not solely dependent on RRT*. Moving forward, we plan to evaluate our approach using various planning algorithms and compare success rates and planning times. Finally, to improve our SFC accuracy and SFRRT* success rate further, we aim to optimize our data collection process by exploring automation through Computational Fluid Dynamics (CFD) frameworks.

REFERENCES

[1] M. Kyrarini *et al.*, "A survey of robots in healthcare," *Technologies*, vol. 9, no. 1, 2021, ISSN: 2227-7080.

[2] M. T. Mason, "Toward robotic manipulation," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, no. 1, pp. 1–28, 2018.

[3] R. I. C. Muchacho, R. Laha, L. F. Figueredo, and S. Haddadin, "A solution to slosh-free robot trajectory optimization," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 223–230.

[4] J. Reinhold, M. Amersdorfer, and T. Meurer, "A dynamic optimization approach for sloshing free transport of liquid filled containers using an industrial robot," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 2336–2341.

[5] R. Di Leva *et al.*, "Time-optimal trajectory planning for anti-sloshing 2-dimensional motions of an industrial robot," in *2021 20th International Conference on Advanced Robotics (ICAR)*, 2021, pp. 32–37.

[6] Y.-H. Lan, B. Wu, Y.-X. Shi, and Y.-P. Luo, "Iterative learning based consensus control for distributed parameter multi-agent systems with time-delay," *Neurocomputing*, vol. 357, pp. 77–85, 2019, ISSN: 0925-2312.

[7] Z. Ge and X. Ge, "Controllability of singular distributed parameter systems in the sense of mild solution," *Journal of Systems Science and Complexity*, vol. 33, no. 5, pp. 1485–1496, Oct. 2020, ISSN: 1559-7067.

[8] J. Han, "A study on the coffee spilling phenomena in the low impulse regime," *Achievements in the Life Sciences*, vol. 10, no. 1, pp. 87–101, 2016, ISSN: 2078-1520.

[9] L. Moriello, L. Biagiotti, C. Melchiorri, and A. Paoli, "Manipulating liquids with robots: A sloshing-free solution," *Control Engineering Practice*, vol. 78, pp. 129–141, 2018, ISSN: 0967-0661.

[10] R. Di Leva, M. Carricato, H. Gattringer, and A. Mueller, "Sloshing dynamics estimation for liquid-filled containers performing 3-dimensional motions: Modeling and experimental validation," *Multibody System Dynamics*, vol. 56, Aug. 2022.

[11] R. Maderna, A. Casalino, A. M. Zanchettin, and P. Rocco, "Robotic handling of liquids with spilling avoidance: A constraint-based control approach," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 7414–7420.

[12] N. Montazeri, A. C. Oliveira, B. H. Himelbloom, M. B. Leigh, and C. A. Crapo, "Chemical characterization of commercial liquid smoke products," *Food science & nutrition*, vol. 1, no. 1, pp. 102–115, 2013.

[13] J. H. Ferziger, *Computational Methods for Fluid Dynamics*, 2nd. Springer, 2002.

[14] J. D. Anderson, *Fundamentals of Aerodynamics*, 7th. McGraw-Hill Education, 2015.

[15] M. Müller, D. Charypar, and M. Gross, "Particle-based fluid simulation for interactive applications," vol. 2003, Jul. 2003, pp. 154–159, ISBN: 1581136595.

[16] M. K. M.H. Djavareshkian, "Simulation of sloshing with the volume of fluid method," 4, vol. 2, 2006, pp. 299–308.

[17] Y. Kuriyama, K. Yano, and M. Hamaguchi, "Trajectory planning for meal assist robot considering spilling avoidance," in *2008 IEEE International Conference on Control Applications*, 2008, pp. 1220–1225.

[18] J. Ichnowski, Y. Avigal, Y. Liu, and K. Goldberg, "Gomp-fit: Grasp-optimized motion planning for fast inertial transport," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 5255–5261.

[19] B. Moya, D. Gonzalez, I. Alfaro, F. Chinesta, and E. Cueto, "Learning slosh dynamics by means of data," *Computational Mechanics*, vol. 64, no. 2, pp. 511–523, Aug. 2019, ISSN: 1432-0924.

[20] Z. Kingston, M. Moll, and L. E. Kavraki, "Sampling-based methods for motion planning with constraints," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, no. 1, pp. 159–185, 2018.

[21] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *CoRR*, vol. abs/1105.1186, 2011.

[22] T. Wolf *et al.*, "Huggingface's transformers: State-of-the-art natural language processing," *ArXiv*, vol. abs/1910.03771, 2019.

[23] L. Berscheid and T. Kröger, "Jerk-limited real-time trajectory generation with arbitrary target states," *CoRR*, vol. abs/2105.04830, 2021.

[24] F. Emika. "Control parameters - franka emika documentation." (2024), [Online]. Available: https : / / frankaemika . github . io / docs / control _ parameters.html.