

A Coarse-to-Fine Framework for Dual-Arm Manipulation of Deformable Linear Objects with Whole-Body Obstacle Avoidance

Mingrui Yu, Kangchen Lv, Changhao Wang, Masayoshi Tomizuka, and Xiang Li

Abstract—Manipulating deformable linear objects (DLOs) to achieve desired shapes in constrained environments with obstacles is a meaningful but challenging task. Global planning is necessary for such a highly-constrained task; however, accurate models of DLOs required by planners are difficult to obtain owing to their deformable nature, and the inevitable modeling errors significantly affect the planning results, probably resulting in task failure if the robot simply executes the planned path in an open-loop manner. In this paper, we propose a coarse-to-fine framework to combine global planning and local control for dual-arm manipulation of DLOs, capable of precisely achieving desired configurations and avoiding potential collisions between the DLO, robot, and obstacles. Specifically, the global planner refers to a simple yet effective DLO energy model and computes a coarse path to find a feasible solution efficiently; then the local controller follows that path as guidance and further shapes it with closed-loop feedback to compensate for the planning errors and improve the task accuracy. Both simulations and real-world experiments demonstrate that our framework can robustly achieve desired DLO configurations in constrained environments with imprecise DLO models, which may not be reliably achieved by only planning or control.

I. INTRODUCTION

Robotic manipulation methods for rigid objects have been extensively studied and reliably applied to manufacturing and services [1]. In addition to rigid objects, deformable linear objects (DLOs) are also common in human life, such as cables, wires, ropes, rods, etc [2]; however, there are new challenges coming from the deformable properties of DLOs when applying classical manipulation methods to DLO manipulation [3], [4].

We focus on a general problem of DLO manipulation: dual-arm manipulating a DLO from a start configuration to a desired (goal) configuration, which contains both moving and shaping of the DLO. In our previous work [5], [6], we have proposed an adaptive controller which can stably and efficiently achieve DLO shape control in unobstructed environments. In this work, we further deal with a more complex and practical scenario: manipulating DLOs in constrained environments with obstacles (see Fig. 1), which requires not only accurate final control results but also proper collision-free moving paths of both DLOs and robot bodies.

A series of works have studied the DLO shaping from the perspective of control, which use the real-time error between

M. Yu, K. Lv, and X. Li are with the Department of Automation, Tsinghua University, Beijing, China. C. Wang and M. Tomizuka are with the Department of Mechanical Engineering, University of California, Berkeley, CA, USA. This work was supported in part by the Science and Technology Innovation 2030-Key Project under Grant 2021ZD0201404, in part by the National Natural Science Foundation of China under Grant U21A20517 and 52075290, and in part by the Institute for Guo Qiang, Tsinghua University. Corresponding author: Xiang Li (xiangli@tsinghua.edu.cn)

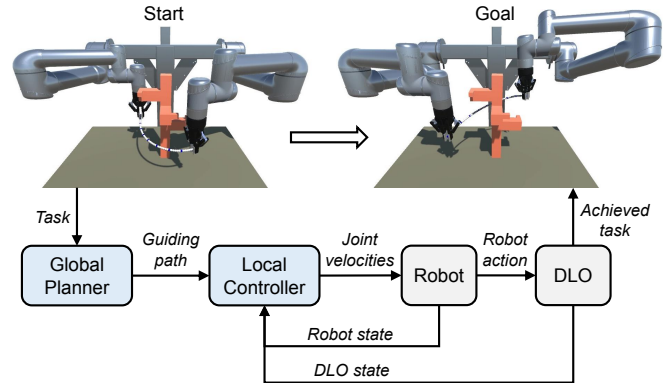


Fig. 1. Overview of the proposed framework for dual-arm manipulation of DLOs in constrained environments with obstacles, in which the global planning and local control complement each other to robustly achieve complex tasks even using imprecise DLO models. The collision avoidance of both the DLO and robot body is considered in the planning and execution.

the current shape and desired shape as the feedback for closed-loop servoing [7]–[11]. However, these control-only methods usually consider a simplified scenario: no obstacles exist, and the robot end-effectors can move freely without considering the arm bodies; moreover, local controllers cannot prevent DLOs from falling into local optimal shapes when desired deformations are large.

To achieve DLO manipulation in constrained environments, global path planning is necessary. Existing studies have proposed different DLO models and incorporated them into classical high-dimensional path planning methods [12]–[14], such as rapidly-exploring random trees (RRTs) [15] and probabilistic roadmap methods (PRMs) [16]. Paths of DLO configurations from the start to the goal are planned, and corresponding end-effector trajectories are extracted for open-loop executions [17]–[19]. However, the planning-only methods are more affected by the inevitable DLO modeling errors than control, since no real-time feedback compensates for the errors. The planned end-effector trajectory may not move the DLO exactly along the expected configuration path, which may cause failures of the actual execution. As a result, most of these approaches are restricted to simulations with sufficiently accurate models.

To overcome the respective shortcomings of the existing planning-only and control-only methods, we propose a new framework that combines both global planning and local control, with the structure shown in Fig. 1. The planner efficiently computes a coarse and global collision-free path to the desired configuration, based on a simple yet effective DLO model; Then, the controller tracks the planned path as

guidance while using closed-loop feedback to compensate for the modeling errors in the planning phase and locally avoid obstacles. In such a framework, the global planning and local control complement each other: the former achieves the finding of feasible solutions and the latter improves the accuracy of the task, which hence enables the robot to robustly and precisely manipulating DLOs to desired configurations in constrained environments while avoiding potential collisions. Specifically, for the planner, we use an RRT planning framework with a DLO energy model for projecting a random DLO configuration to a stable one; for the controller, we formulate it as a model predictive control (MPC) problem with artificial potentials, in which we apply our previously proposed Jacobian model learning method [5]. The full configurations of the DLO and robot are considered in the global planning and local control. We carry out both simulations and real-world experiments to demonstrate that our method can achieve various DLO manipulation tasks in constrained environments which may not be reliably achieved by only planning or control. The video is available on the project website¹.

II. RELATED WORK

A. DLO Path Planning

To use sampling-based planning methods for high-dimensional DLO path planning, an approach to generate valid DLO configurations is required. Bretl and McCarthy [14] formulated the DLO static equilibrium as an optimal control solution and derived that the configuration space of an one-end-fixed Kirchhoff elastic rod is a six-dimensional manifold, but it requires to solve complicated differential equations to determine whether a sample is stable. Roussel et al. [17] used this model for global DLO motion planning in complex tasks at the cost of very huge computation time. To reduce the online time cost, [18] pre-computed a roadmap of stable DLO configurations. However, the roadmap takes a lot of time to compute and memory space to store, and the planning performance highly depends on the roadmap size. Furthermore, this ideal model cannot consider gravity, and [20] showed that it may only be appropriate for DLOs of very high stiffness, such as nitinol rods.

An alternative way is to compute stable DLO equilibrium configurations by locally minimizing the deformation energy using general optimization approaches [12], [13]. We hold the view that such methods are more general and suitable for DLO path planning since we can use it to efficiently project a random invalid configuration to a valid one and incorporate it into standard single-query planning algorithms. In this work, we employ a discrete DLO energy model and an RRT planning framework, achieving efficient and practical DLO path planning.

Recently, learning non-physical DLO models for planning from data using neural networks has also been studied. In [19], three networks are trained for planning and execution, including a state-action forward prediction model, a classifier

for predicting the validity of the model, and a recovery policy. The problems are that such methods need to pre-collect a lot of data of the manipulated DLO for training, and their effectiveness on different untrained DLOs and scenarios cannot be guaranteed.

B. DLO Shape Control

Controlling DLOs to desired shapes using local controllers has been widely researched [7]–[11], and related methods have been detailedly reviewed in [5]. However, most of them cannot handle large deformations, and the manipulation processes are usually not smooth.

In our previous work [5], we proposed an efficient and adaptive approach to learning a state-action DLO model and achieved stable and smooth large deformation control, which makes it possible to track planned paths. In this work, we additionally apply the artificial potential method for local obstacle avoidance, since the actual execution path may not be exactly the planned collision-free path.

C. Combining Planning and Control for DLO Manipulation

McConachie et al. [21] proposed another framework for interleaving planning and control for manipulating deformable objects. Their controller first attempts to perform tasks directly; if their deadlock predictor predicts the controller will get stuck, their planner will be invoked to move the object to a new region. Note that their aim and combination approach are different from ours: first, their planner uses a more simplified model called “virtual elastic band” focusing on the overstretch of deformable objects caused by grippers or obstacles; second, their planned path of grippers is open-loop executed without local adjustment using real-time feedback; as a result, the DLO may be over-compressed or hooked by obstacles. In contrast, our method considers the full DLO configurations both in planning and control, and employs feedback control during the whole manipulation process.

III. PRELIMINARIES

In this work, a dual-arm robot grasps the two ends of the DLO and quasi-statically manipulates the DLO to the desired configuration, during which both the DLO and the arms should not collide with obstacles. The configuration of the DLO is represented by the positions of multiple features uniformly distributed along the DLO.

Some frequently-used notations are listed as follows. The position of the k^{th} feature is represented as $\mathbf{x}_k \in \mathbb{R}^3$. The configuration vector of the DLO is represented as $\mathbf{x} := [\mathbf{x}_1; \dots; \mathbf{x}_m] \in \mathbb{R}^{3m}$, where m is the number of the features. The joint position of the dual-arm robot is represented as $\mathbf{q} \in \mathbb{R}^n$, where n is the degrees of freedom (DoFs) of the robot. The configuration vector of the robot end-effectors is represented as $\mathbf{r} \in \mathbb{R}^{12}$.

The control input \mathbf{u} to the system is the joint velocity of the dual-arm robot. That is

$$\dot{\mathbf{q}} = \mathbf{u} \quad (1)$$

¹Project website: https://mingrui-yu.github.io/DLO_planning

IV. GLOBAL PLANNING

The planner is first invoked to efficiently find a global collision-free path in constrained environments using a simple yet effective DLO energy model. This coarse path will be sent to the controller as global guidance, to achieve the solving of the complex manipulation task.

A. Projection to Stable DLO Configuration Manifold

As defined in Section III, the raw configuration space of the DLO is $3m$ dimensional; however, only a subspace contains stable equilibrium configuration, which can be called a manifold. Thus, the planning is constrained on this manifold, i.e., the DLO configuration of each node should be guaranteed stable. Finding a stable configuration by randomly sampling in the raw space with rejection strategies is unlikely because the dimension of the stable space is much less than that of the raw space. Alternatively, it is more appropriate to use projection methods to move a random sample in the raw space onto its neighbor constrained manifold [22].

Such projections can be achieved by utilizing a DLO energy model. Denote the potential energy of an elastic DLO as E , which is assumed to be fully determined by the DLO configuration \mathbf{x} . A stable equilibrium of the DLO with two end poses fixed is where the DLO's internal configuration locally minimizes the potential energy E [12], [13]. That is

$$\frac{\partial E}{\partial \mathbf{x}_k} = \mathbf{0}, \quad \forall k = 3, \dots, m-2 \quad (2)$$

where the first two and last two feature points are fixed to represent the fixed end poses in discrete DLO models [23].

A random sample in the raw space \mathbf{x}^0 can be projected onto the stable configuration manifold by formulating it as a local minimization problem of the energy with \mathbf{x}^0 as the initial value:

$$\begin{aligned} \mathbf{x}^{\text{stable}} &= \arg \min_{\mathbf{x}} E(\mathbf{x}) \\ \text{s.t. } \mathbf{x}_k &= \mathbf{x}_k^0, \quad \forall k = 1, 2, m-1, m \end{aligned} \quad (3)$$

More constraints may be incorporated for specific tasks. We denote the process of projecting a raw configuration to a neighbor stable configuration as $\mathbf{x}^{\text{stable}} = \text{ProjectStableConfig}(\mathbf{x}^0)$.

In this work, we use a simple mass-spring model [3] as the DLO energy model (illustrated in Fig. 2(a)), where the energy E for a configuration \mathbf{x} is specified as

$$\begin{aligned} E &= \sum_{k=1}^{m-1} \frac{1}{2} \lambda_1 \left(\|\mathbf{x}_{k+1} - \mathbf{x}_k\|_2 - \frac{L}{m-1} \right)^2 \\ &+ \sum_{k=1}^{m-2} \frac{1}{2} \lambda_2 \left(\|\mathbf{x}_{k+2} - \mathbf{x}_k\|_2 - \frac{2L}{m-1} \right)^2 \end{aligned} \quad (4)$$

where λ_1 and λ_2 are the stiffness of Type 1 and Type 2 springs, respectively, and L is the DLO length. In our simulations and experiments, we use the same stiffness values in the planning for all different DLOs. We will show that planning using such a simple model with imprecise parameters can provide acceptable guiding paths for the controller to achieve manipulation tasks. We visualize two examples of the ProjectStableConfig in Fig. 2(b) and 2(c).

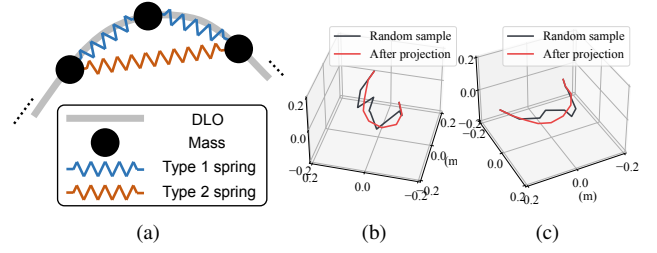


Fig. 2. (a) Illustration of the mass-spring model which models a DLO as a series of masses connected by springs. The i^{th} mass is connected to $(i \pm 1)^{\text{th}}$ masses by Type 1 springs and $(i \pm 2)^{\text{th}}$ masses by Type 2 springs. (b)(c) Illustrations of projecting a randomly sampled raw DLO configuration to a stable configuration by locally minimizing the potential energy of the mass-spring model as (3).

B. Planning Algorithm

Our planning algorithm uses the same framework as the Constrained Bi-directional Rapidly-Exploring Random Tree (CBiRRT) method proposed in [22]. The constraint projection guarantees the validity of nodes, and the bi-directional RRT significantly improves the planning efficiency. Two trees, respectively beginning from the start configuration and goal configuration, grow toward each other, until they are connected. Each node \mathcal{N} in the trees contains both the DLO configuration \mathbf{x} and the robot configuration \mathbf{q} .

During each iteration, one tree is selected as \mathcal{T}^A and the other one as \mathcal{T}^B . First, a random node $\mathcal{N}_{\text{rand}}$ is sampled. The node in \mathcal{T}^A nearest to $\mathcal{N}_{\text{rand}}$ is selected as $\mathcal{N}_{\text{near}}^A$. Then, \mathcal{T}^A grows from $\mathcal{N}_{\text{near}}^A$ towards $\mathcal{N}_{\text{rand}}$ one or more steps and terminates at $\mathcal{N}_{\text{reached}}^A$ using a ConstrainedExtend function. Next, the node in \mathcal{T}^B nearest to $\mathcal{N}_{\text{reached}}^A$ is selected as $\mathcal{N}_{\text{near}}^B$. Then, \mathcal{T}^B grows from $\mathcal{N}_{\text{near}}^B$ towards $\mathcal{N}_{\text{reached}}^A$ and terminates at $\mathcal{N}_{\text{reached}}^B$. If $\mathcal{N}_{\text{reached}}^B$ can be connected to $\mathcal{N}_{\text{reached}}^A$, then a path is found; otherwise, the two tree are swapped and next iteration starts.

Specifically, when sampling an $\mathcal{N}_{\text{rand}}$, we first sample a stable DLO configuration using the ProjectStableConfig function. A heuristic raw-space sampling method is preferred to sample initial configurations more like a DLO. If the DLO is in collision, the node is discarded and a new one is generated. In addition, the distance between two nodes is specified as the Euclidean distance between their DLO configuration vectors.

The function $\mathcal{N}_{\text{reached}} = \text{ConstrainedExtend}(\mathcal{N}_{\text{from}}, \mathcal{N}_{\text{to}})$ iteratively moves from $\mathcal{N}_{\text{from}}$ to \mathcal{N}_{to} and returns reached $\mathcal{N}_{\text{reached}}$. In each step, a new DLO configuration is calculated by interpolating from the last reached configuration \mathbf{x}_{last} to \mathbf{x}_{to} with a small step size, where we use linear interpolation for the centroid positions and spherical interpolation for the relative deformations irrelevant to translations to keep the overall DLO shapes and avoid over-compression [13]. The ProjectStableConfig is then applied to project the interpolated configuration to a stable \mathbf{x}_{new} . The corresponding robot configuration \mathbf{q}_{new} is obtained as the inverse kinematics (IK) solution for the end poses of \mathbf{x}_{new} , and the solution closest to \mathbf{q}_{last} is chosen if multiple solutions exist. If the IK succeeds and no collision occurs, a new node \mathcal{N}_{new} is generated. Such iteration repeats until the node generation fails or \mathcal{N}_{new} is

farther from \mathcal{N}_{to} than the last reached node \mathcal{N}_{last} .

For the start node, both the DLO and robot configurations are known; however, for the goal node, only the goal DLO configuration is given. To use the bi-directional RRT, we set the goal robot configuration as the IK solution for the estimated DLO end poses. If multiple solutions exist, we choose the one closest to a pre-default robot configuration. Note that in actual executions, the robot with this goal configuration may not bring the DLO to the desired shape.

When the planning succeeds, the path from the start node to the goal node is extracted from the two trees and then smoothed. The final path is denoted as $\mathcal{P} = \{\mathcal{N}_{start}, \mathcal{N}_1, \dots, \mathcal{N}_{goal}\}$ and sent to the local controller as a guiding global path.

V. MANIPULATION WITH LOCAL CONTROL

Because of the DLO modeling errors, the planned path may not be exactly achievable. We apply an MPC to track the planned path as guidance while using real-time feedback to adjust the robot motion to ensure that the actual path of both the DLO and robot is still collision-free and the final desired DLO configuration is precisely reached.

A. Manipulation Process

Each node of the planned guiding path is regarded as an intermediate desired configuration for the controller. During the actual execution, the DLO and the manipulators are controlled to move along the path, i.e., to move to the next intermediate configuration iteratively until reaching the final desired configuration.

Each intermediate desired configuration (except for the final desired configuration) doesn't need to be exactly reached, since it may be unachievable owing to modeling errors. If the distance between the current and desired DLO configuration is less than a threshold, the control objective will be switched to the next planned configuration. Besides, if the controller gets stuck (i.e., the control input is close to zero), the global planner will be invoked again to re-plan a path from the current configuration to the final desired configuration.

B. Controlling to an Intermediate Configuration

While moving to the next planned node \mathcal{N}_i , the desired DLO configuration \mathbf{x}_d is set as $\mathcal{N}_i.\mathbf{x}$ and the desired robot configuration \mathbf{q}_d is set as $\mathcal{N}_i.\mathbf{q}$.

1) *Artificial Potentials*: The attractive potential for driving the DLO to \mathbf{x}_d and the attractive potential for driving the arms to \mathbf{q}_d is respectively designed as

$$U_a^{dlo}(\mathbf{x}) = \frac{1}{2}(\mathbf{x} - \mathbf{x}_d)^\top (\mathbf{x} - \mathbf{x}_d) \quad (5)$$

$$U_a^{arm}(\mathbf{q}) = \frac{1}{2}(\mathbf{q} - \mathbf{q}_d)^\top (\mathbf{q} - \mathbf{q}_d) \quad (6)$$

To achieve local obstacle avoidance of the DLO and robot, we design repulsive potentials. The repulsive potential for prevent a point \mathbf{p} from colliding with obstacles \mathcal{CO} in the cartesian space is designed as

$$U_r(\mathbf{p}) = \begin{cases} \frac{1}{2} \left(\frac{1}{c(\mathbf{p})} - \frac{1}{\epsilon_c} \right)^2, & c(\mathbf{p}) < \epsilon_c \\ 0, & \text{otherwise} \end{cases} \quad (7)$$

where $c(\mathbf{p}) = \min_{\mathbf{p}' \in \mathcal{CO}} \|\mathbf{p} - \mathbf{p}'\|$ is the distance between the point and the closest obstacle, and ϵ_c is the maximum distance that the repulsive potential affects. The repulsive potential for the DLO is then designed as

$$U_r^{dlo}(\mathbf{x}) = \sum_{k=1}^m U_r(\mathbf{x}_k) \quad (8)$$

Denote the Cartesian-space positions of y control points on the arm bodies as $\boldsymbol{\xi}(\mathbf{q}) = [\boldsymbol{\xi}_1(\mathbf{q}); \dots; \boldsymbol{\xi}_y(\mathbf{q})]$, where $\boldsymbol{\xi}_i(\cdot)$ represents the forward kinematics for the i^{th} point. In this work we choose one control point from each arm link [24]. The repulsive potential for the arms is then designed as

$$U_r^{arm}(\mathbf{q}) = \sum_{i=1}^y U_r(\boldsymbol{\xi}_i(\mathbf{q})) \quad (9)$$

2) *Control Law*: The controller is formulated as an one-step MPC. The control input is specified as the locally optimal solution of the following optimization problem:

$$\begin{aligned} \min_{\mathbf{u}} \quad & \mathcal{J} = \lambda_a^{dlo} U_a^{dlo}(\mathbf{x}(t + \delta t)) + \lambda_r^{dlo} U_r^{dlo}(\mathbf{x}(t + \delta t)) \\ & + \lambda_a^{arm} U_a^{arm}(\mathbf{q}(t + \delta t)) + \lambda_r^{arm} U_r^{arm}(\mathbf{q}(t + \delta t)) \\ & + \frac{1}{2} \mathbf{u}^\top \mathbf{K}_u \mathbf{u} \\ \text{s.t.} \quad & \mathbf{C}_{dof} \mathbf{J}^{arm} \mathbf{u} = \mathbf{0} \\ & \mathbf{u}^\top \mathbf{K}_u \mathbf{u} \leq u_{max}^2 \end{aligned} \quad (10)$$

where $\lambda_a^{dlo}, \lambda_r^{dlo}, \lambda_a^{arm}, \lambda_r^{arm}$ are weighting coefficients for the potentials, and δt is the step interval. The \mathbf{K}_u is a weighting matrix for the input joint velocities, and related terms are for penalizing and constraining the input magnitude. The first equality constraint is to constrain the allowed DoFs of the robot end-effectors for specific tasks (e.g., 2-D tasks), where \mathbf{J}^{arm} is the robot Jacobian matrix.

To estimate the DLO configuration at the next time step given the control input \mathbf{u} , we use a DLO Jacobian matrix which relates the velocity of DLO features to the velocity of robot end-effectors as $\dot{\mathbf{x}} = \mathbf{J}^{dlo} \dot{\mathbf{r}}$. We employ our previously proposed offline-online data-driven method [5] to obtain an estimation $\hat{\mathbf{J}}^{dlo}$. In addition, $\dot{\mathbf{r}} = \mathbf{J}^{arm} \dot{\mathbf{q}}$. Then, the configuration at the next step is approximated as

$$\mathbf{q}(t + \delta t) \approx \mathbf{q}(t) + \mathbf{u} \delta t \quad (11)$$

$$\mathbf{x}(t + \delta t) \approx \mathbf{x}(t) + \hat{\mathbf{J}}^{dlo} \mathbf{J}^{arm} \mathbf{u} \delta t \quad (12)$$

To avoid the huge computation of incorporating the distance calculation into each optimization iteration, we linearize the repulsive potential terms in the cost function as

$$U_r^{arm}(\mathbf{q}(t + \delta t)) \approx U_r^{arm}(\mathbf{q}(t)) + \nabla_{\mathbf{q}} U_r^{arm}(\mathbf{q}(t)) \mathbf{u} \delta t \quad (13)$$

$$U_r^{dlo}(\mathbf{x}(t + \delta t)) \approx U_r^{dlo}(\mathbf{x}(t)) + \nabla_{\mathbf{x}} U_r^{dlo}(\mathbf{x}(t)) \hat{\mathbf{J}}^{dlo} \mathbf{J}^{arm} \mathbf{u} \delta t \quad (14)$$

Then, only the signed distances at the current time t are required, and the MPC can be efficiently solved using gradient-based optimization solvers.

When controlling to the final desired configuration, there is no need to consider the planned arm configuration, so λ_a^{arm} is set to zero.

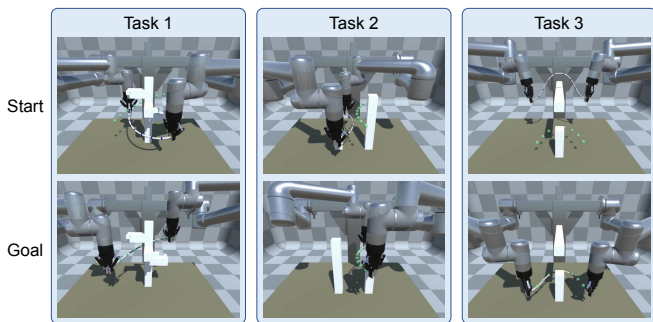


Fig. 3. Three simulated tasks for testing the overall performance of the proposed method. The top row is the start configurations and the bottom row is the goal configurations. The pictures are taken from actual manipulations.

TABLE I

PERFORMANCE OF OUR METHOD IN THE THREE SIMULATED TASKS.

Task	Planning success rate	Planning time (s) ^a	Manipulation success rate	Final task error (cm) ^b
1	8/10	3.29 ± 1.45	8/8	0.55 ± 0.22
2	8/10	4.14 ± 4.31	8/8	0.55 ± 0.11
3	10/10	2.66 ± 1.25	10/10	0.60 ± 0.18

^{a,b} The mean value ± standard deviation over the successful cases.

VI. RESULTS

In both the simulations and experiments, the DLO is grasped by a dual-UR5 robot and its configuration is represented by 10 features ($m = 10$). The control frequency is 10 Hz. The stiffness of the mass-spring model is set as $\lambda_1 = 1.0$ and $\lambda_2 = 1.0$. Other hyper-parameters are set as $\epsilon_c = 0.1$, $\lambda_a^{\text{dlo}} = 10.0$, $\lambda_a^{\text{arm}} = 1.0$, and $\lambda_r^{\text{dlo}} = \lambda_r^{\text{arm}} = 10^{-3}$. The planning algorithm is written in C++ and runs on an Intel i7-10700 CPU (2.9GHz), where the optimization is based on the Ceres Solver [25].

We define the following criteria for evaluation: 1) planning success: the planner finds a feasible path within 1000 RRT iterations; 2) planning time: the time cost for finding a feasible path; 3) final task error: the Euclidean distance between the final desired DLO configuration and reached configuration in a manipulation; 4) collision time: the time (second) of collisions between the DLO, robot, and obstacles during a manipulation; 5) manipulation success: the final task error is less than 5cm. Note that the task error contains errors of all features on the DLO without averaging.

A. Simulations

The simulation is built in Unity [26] with Obi [27] for simulating DLOs and Unity Robotics Hub [28] for integration with ROS. The simulator is utilized as a black box.

1) *Overall Performance*: We first test the overall performance of our method in three different tasks with different scenarios, DLO properties, and start/goal configuration, as shown in Fig 3. The lengths of the DLOs are 0.5/0.4/0.6m for Task 1/2/3, respectively. For each task, the planner tries 10 queries and all found paths are executed by the controller. The performance is summarized in Table I. The planning success rate over all queries is 87%, and most of the successful queries take about 2~5 seconds to find a feasible

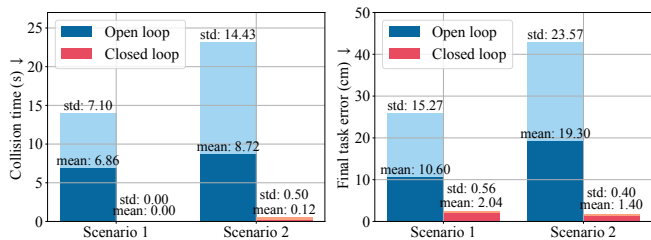


Fig. 4. Quantitative results of the comparison between the proposed closed-loop framework and the open-loop manner. The bars with darker colors refer to the mean values over all 25 manipulations, and those with lighter colors refer to the standard deviations (std).

path. All 26 paths are successfully executed by the controller, and the final tasks are achieved with an average final task error of 0.57cm. Re-planning is invoked 2 times during all 26 manipulations. One ProjectStableConfig takes 0.7ms on average. It is demonstrated that our planning and control framework can efficiently and robustly achieve various DLO manipulation tasks.

2) *Closed-Loop v.s. Open-Loop*: We verify the significance of the control module by comparing the actual manipulation results of directly executing the planned robot path in an open-loop manner and using the proposed controller to track the planned guiding path in a closed-loop manner.

We both do quantitative tests to show the statistical comparison results and visualize two specific cases which clearly illustrate the differences. The quantitative tests use the same two scenarios as the two cases in Fig. 5, and the DLO lengths are 0.4/0.6m, respectively. The DLOs start from straight lines in front of the obstacles. Five different desired shapes behind the obstacles are tested for each scenario, and five paths are planned for each desired shape, so totally 25 paths are executed in both open-loop and closed-loop manners for each scenario. Fig. 4 shows the results of the collision time during manipulations and the final task errors. It can be seen that 1) the proposed closed-loop manner can robustly accomplish DLO manipulation tasks using coarse planned paths and achieve precise final control results (1.77cm on average), while the open-loop manner cannot guarantee that the DLO moves as expected like the case in Fig. 5(a); 2) the proposed closed-loop manner can effectively avoid collisions (0.06s on average), while collisions often happen in the open-loop manipulations, which may cause failures of tasks like the case in Fig. 5(b).

For better illustration, we visualize the manipulation processes of two specific cases in Fig. 5. As for the first case shown in Fig. 5(a), in the open-loop manipulation, although the robot moves along the planned robot path, the DLO gradually deviates from the planned DLO path and finally terminates at another stable but undesired shape; while in the closed-loop manipulation, the controller receives the real-time DLO configuration feedback and adjusts the robot motion, during which the actual robot path is not exactly as planned. As for the second case shown in Fig 5(b), in the open-loop manipulation, unexpected collisions between the DLO and obstacles happen and the DLO is hooked, resulting in the failure of the task; while in the closed-loop

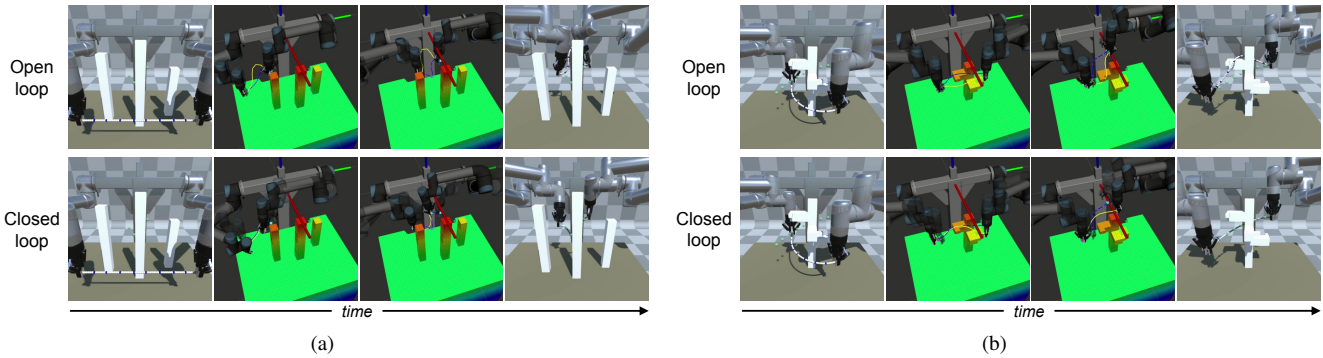


Fig. 5. Two cases for illustrating the comparison between the open-loop manner and proposed closed-loop manner using the same planned paths (more clearly shown in the supplementary video). The first column is the start configurations; the second and third columns are the manipulation processes, where the blue lines and translucent robots refer to the planned intermediate configurations and the yellow lines and non-translucent robots refer to the real-time configurations; and the last column is the final reached configurations, where the translucent green points refer to the final desired DLO configurations.



Fig. 6. Manipulation processes of three real-world tasks using the proposed planning and control framework. The green+black circles refer to the final desired DLO configurations. For Task 3, we also show the comparison between the proposed closed-loop manner and the open-loop manner.

manipulation, thanks to the local obstacle avoidance achieved by the controller using artificial potentials, the DLO avoids the obstacles and finally reaches the desired configuration.

B. Real-World Experiments

We also validate the proposed framework in the real world. To simplify the perception of DLOs, we put markers along the DLOs and conduct 2-D tasks where the DLOs are placed on a table to avoid visual occlusion.

Three representative tasks are achieved using the proposed method, as shown in Fig. 6. In Task 1, the DLO needs to be rotated around 180 degrees; in Task 2, the method in [21] may fail, since the "virtual elastic band" connecting the two end-effectors will not collide with the obstacle along a straight-line robot path, but the actual DLO body will; and in Task 3, the DLO needs to be manipulated from an upper-semicircle shape to a lower-semicircle shape. A 0.3m-long nylon rope, a 0.6m-long HDMI cable, and a 0.45m-long electric wire are used in Task 1, 2, and 3, respectively. For Task 3, we also show the comparison between the proposed closed-loop manner and the open-loop manner using the same planned path, in which the open-loop execution moves the DLO to an unexpected shape and also causes the

collision, while the closed-loop execution succeeds.

VII. CONCLUSION

This paper proposes a coarse-to-fine framework for dual-arm manipulation of DLOs with whole-body obstacle avoidance. A coarse global path is first efficiently planned using an RRT planning framework, in which a DLO energy model is employed to generate stable DLO configurations. Then, an MPC controller tracks the planned path as global guidance, during which it reacts to the real-time states in a closed-loop manner to compensate for the modeling errors and avoid collisions. This formulation considers both the feasibility and accuracy of long-horizon manipulation tasks in constrained environments. The results show that our method can robustly and precisely achieve various DLO manipulation tasks in both simulations and real-world experiments, where the necessity of such a closed-loop framework is also demonstrated by comparative studies.

Future work will be devoted to incorporating more accurate discrete DLO energy models and adding gravity energy terms, further tuning each component of the planning algorithm, employing long-horizon MPC for local control, or applying marker-free perception methods [29].

REFERENCES

- [1] A. Billard and D. Kragic, "Trends and challenges in robot manipulation," *Science*, vol. 364, no. 6446, p. eaat8414, 2019.
- [2] J. Sanchez, J.-A. Corrales, B.-C. Bouzgarrou, and Y. Mezouar, "Robotic manipulation and sensing of deformable objects in domestic and industrial applications: a survey," *The International Journal of Robotics Research*, vol. 37, no. 7, pp. 688–716, 2018.
- [3] H. Yin, A. Varava, and D. Kragic, "Modeling, learning, perception, and control methods for deformable object manipulation," *Science Robotics*, vol. 6, no. 54, 2021.
- [4] J. Zhu, A. Cherubini, C. Dune, D. Navarro-Alarcon, F. Alambeigi, D. Berenson, F. Ficuciello, K. Harada, J. Kober, X. Li, *et al.*, "Challenges and outlook in robotic manipulation of deformable objects," *IEEE Robotics and Automation Magazine*, 2021.
- [5] M. Yu, K. Lv, H. Zhong, S. Song, and X. Li, "Global model learning for large deformation control of elastic deformable linear objects: An efficient and adaptive approach," *IEEE Transactions on Robotics*, vol. 39, no. 1, pp. 417–436, 2023.
- [6] M. Yu, H. Zhong, and X. Li, "Shape control of deformable linear objects with offline and online learning of local linear deformation models," in *2022 IEEE International Conference on Robotics and Automation (ICRA)*, 2022, pp. 1337–1343.
- [7] D. Navarro-Alarcon, H. M. Yip, Z. Wang, Y. Liu, F. Zhong, T. Zhang, and P. Li, "Automatic 3-d manipulation of soft objects by robotic arms with an adaptive deformation model," *IEEE Transactions on Robotics*, vol. 32, no. 2, pp. 429–441, 2016.
- [8] C. Wang, Y. Zhang, X. Zhang, Z. Wu, X. Zhu, S. Jin, T. Tang, and M. Tomizuka, "Offline-online learning of deformation model for cable manipulation with graph neural networks," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 5544–5551, 2022.
- [9] J. Zhu, D. Navarro-Alarcon, R. Passama, and A. Cherubini, "Vision-based manipulation of deformable and rigid objects using subspace projections of 2d contours," *Robotics and Autonomous Systems*, vol. 142, p. 103798, 2021.
- [10] R. Lagneau, A. Krupa, and M. Marchal, "Automatic shape control of deformable wires based on model-free visual servoing," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5252–5259, 2020.
- [11] S. Jin, C. Wang, and M. Tomizuka, "Robust deformation model approximation for robotic cable manipulation," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 6586–6593.
- [12] H. Wakamatsu and S. Hirai, "Static modeling of linear object deformation based on differential geometry," *The International Journal of Robotics Research*, vol. 23, no. 3, pp. 293–311, 2004.
- [13] M. Moll and L. Kavraki, "Path planning for deformable linear objects," *IEEE Transactions on Robotics*, vol. 22, no. 4, pp. 625–636, 2006.
- [14] T. Bretl and Z. McCarthy, "Quasi-static manipulation of a kirchhoff elastic rod based on a geometric analysis of equilibrium configurations," *The International Journal of Robotics Research*, vol. 33, no. 1, pp. 48–68, 2014.
- [15] S. M. LaValle *et al.*, "Rapidly-exploring random trees: A new tool for path planning," 1998.
- [16] N. M. Amato and Y. Wu, "A randomized roadmap method for path and manipulation planning," in *1996 IEEE International Conference on Robotics and Automation (ICRA)*, vol. 1, 1996, pp. 113–120.
- [17] O. Roussel, P. Fernbach, and M. Taïx, "Motion planning for an elastic rod using contacts," *IEEE Transactions on Automation Science and Engineering*, vol. 17, no. 2, pp. 670–683, 2020.
- [18] A. Sintov, S. Macenski, A. Borum, and T. Bretl, "Motion planning for dual-arm manipulation of elastic rods," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6065–6072, 2020.
- [19] P. Mitrano, D. McConachie, and D. Berenson, "Learning where to trust unreliable models in an unstructured world for deformable object manipulation," *Science Robotics*, vol. 6, no. 54, p. eabd8170, 2021.
- [20] I. Mishani and A. Sintov, "Real-time non-visual shape estimation and robotic dual-arm manipulation control of an elastic wire," *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 422–429, 2022.
- [21] D. McConachie, A. Dobson, M. Ruan, and D. Berenson, "Manipulating deformable objects by interleaving prediction, planning, and control," *The International Journal of Robotics Research*, vol. 39, no. 8, pp. 957–982, 2020.
- [22] D. Berenson, S. S. Srinivasa, D. Ferguson, and J. J. Kuffner, "Manipulation planning on constraint manifolds," in *2009 IEEE International Conference on Robotics and Automation (ICRA)*, 2009, pp. 625–632.
- [23] M. Bergou, M. Wardetzky, S. Robinson, B. Audoly, and E. Grinspun, "Discrete elastic rods," in *ACM SIGGRAPH 2008 Papers*, New York, NY, USA, 2008.
- [24] B. Siciliano, L. Sciavicco, L. Villani, and G. Oriolo, *Robotics: Modelling, Planning and Control*. London: Springer, 2009.
- [25] S. Agarwal, K. Mierle, and T. C. S. Team, "Ceres Solver," 2022. [Online]. Available: <https://github.com/ceres-solver/ceres-solver>
- [26] U. Technologies. (2022) Unity real-time development platform. [Online]. Available: <https://unity.com/>
- [27] V. M. Studio. (2022) Obi - Unified particle physics for Unity 3D. [Online]. Available: <http://obi.virtualmethodstudio.com/>
- [28] U. Technologies. (2022) Unity Robotics Hub. [Online]. Available: <https://github.com/Unity-Technologies/Unity-Robotics-Hub>
- [29] K. Lv, M. Yu, Y. Pu, X. Jiang, G. Huang, and X. Li, "Learning to estimate 3-d states of deformable linear objects from single-frame occluded point clouds," *arXiv preprint arXiv:2210.01433*, 2022.