

Learning to Estimate 3-D States of Deformable Linear Objects from Single-Frame Occluded Point Clouds

Kangchen Lv, Mingrui Yu, Yifan Pu, Xin Jiang, Gao Huang, and Xiang Li

Abstract—Accurately and robustly estimating the state of deformable linear objects (DLOs), such as ropes and wires, is crucial for DLO manipulation and other applications. However, it remains a challenging open issue due to the high dimensionality of the state space, frequent occlusions, and noises. This paper focuses on learning to robustly estimate the states of DLOs from single-frame point clouds in the presence of occlusions using a data-driven method. We propose a novel two-branch network architecture to exploit global and local information of input point cloud respectively and design a fusion module to effectively leverage the advantages of both methods. Simulation and real-world experimental results demonstrate that our method can generate globally smooth and locally precise DLO state estimation results even with heavily occluded point clouds, which can be directly applied to real-world robotic manipulation of DLOs in 3-D space.

I. INTRODUCTION

Robotic manipulation of deformable linear objects (DLOs), such as ropes and wires, has a wide variety of applications in industrial, service, and health-care sectors [1], [2]. An accurate and robust state estimator for DLOs is obviously the prerequisite for subsequent manipulations. Compared to rigid objects, the infinite dimensional DLO state space makes it very challenging to perceive deformations. Besides, occlusions and noises occur frequently in unstructured environments, resulting in higher requirements for robust DLO state estimation.

Commonly used representations to describe DLO states include Fourier-based parameterization [3], implicit latent descriptors learned by neural networks [4], [5], a chain of uniformly distributed nodes [6]–[9], etc. Among these methods, representing a DLO as a chain of 3-D nodes (see Fig. 1) is general in various manipulation tasks and will be adopted in this work.

A complete processing stream to estimate the DLO state can be roughly divided into three procedures: segmentation (i.e., segmenting the DLO from environment), detection (i.e., estimating the DLO state in a single frame), and tracking (i.e., tracking the deformation across several frames). As sensors are RGB or RGB-D cameras in most cases, segmenting the DLO region in image space is the essential first step for consecutive processing. [10]–[13] focus on how to obtain pixel-level DLO masks of high quality using

K. Lv, M. Yu, Y. Pu, G. Huang, and X. Li are with the Department of Automation, Tsinghua University, China. X. Jiang is with the Beijing Academy of Artificial Intelligence, Beijing, China. This work was supported in part by the National Key R&D Program of China under Grant 2020AAA0105200, in part by the Institute for Guo Qiang, Tsinghua University, and in part by the National Natural Science Foundation of China under Grant U21A20517 and 52075290. Corresponding author: Xiang Li (xiangli@tsinghua.edu.cn)

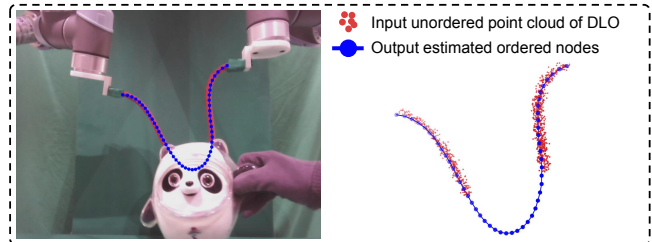


Fig. 1. Illustration of our task: 3-D occlusion-robust DLO state estimation from a single-frame point cloud. Red points are the unordered incomplete point cloud of the occluded rope and blue connected dots represent our estimated ordered node sequence as its current state.

traditional image processing or data-driven methods. As for detection, this step aims at estimating the positions of nodes along the DLO in one frame with the cleaned sensory data as input. For example, [14], [15] use neural networks to encode the DLO into several sequential key-points in the 2-D image space; [16] estimates a skeleton line and 3-D joint positions on it from point cloud to represent the DLO, but not robust against occlusions and different DLO types. As for tracking, various works have also been proposed to track the correspondence of point cloud across video frames in the presence of occlusions and self-intersections [17]–[23]. These works model DLO tracking task as a GMM-based non-rigid point registration problem with some geometric constraints. However, these pure tracking-based methods rely on an accurate initial state which requires manual setting or specific initial conditions. Besides, there are few effective ways to rectify the accumulated drift errors or re-initialize for tracking failure. Therefore, it is necessary to develop an accurate and robust 3-D state estimation method for DLOs from a single frame, which can be independently applied to estimate the DLO state in each frame or combined with tracking methods above to utilize temporal information.

In this paper, we focus on estimating a sequence of ordered and uniformly distributed nodes from single-frame point cloud occlusion-robustly to represent the state of DLO, as shown in Fig. 1. Note that we only use point cloud as our input without any auxiliary physical simulation and robot configurations. The challenges of this task are as follows: 1) there are few distinguishable features in the point cloud of DLOs; 2) occlusions and noises are common in the environment; 3) generalization ability for different DLOs is required. To deal with challenges above, we propose a novel two-branch network architecture to leverage both the global geometry information for guaranteeing smooth and occlusion-robust shape, and local geometry information for

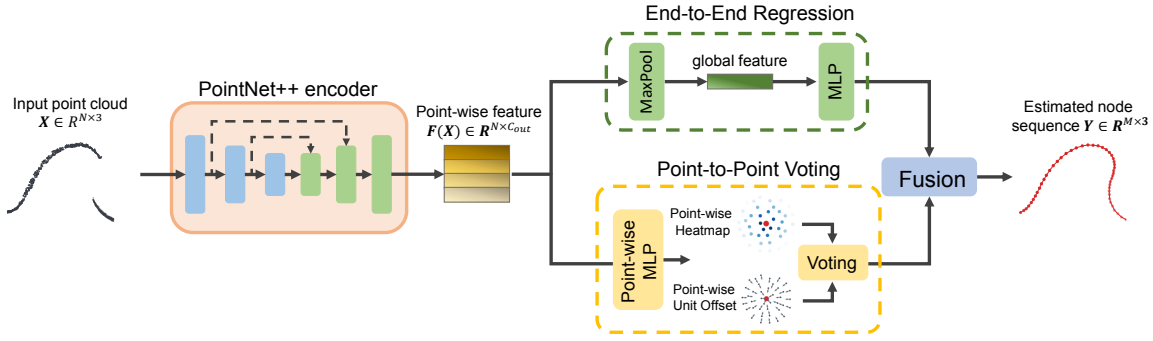


Fig. 2. Overview of the proposed method for occlusion-robustly estimating the 3-D states of DLOs. The input point cloud which might be fragmented due to occlusions is first fed into a PointNet++ encoder and the extracted features are then processed by two parallel branches: *End-to-End Regression* and *Point-to-Point Voting*. The estimation results of these two branches are finally fused with a fusion module to obtain the final output node sequence.

precise estimations. To the best of our knowledge, we are the first to realize accurate and robust 3-D state estimation of DLOs from single-frame point cloud input even with heavy occlusions. Specifically, we first exploit a PointNet++ encoder [24] to extract deep features of the input point cloud and then feed the features into two branches: *End-to-End Regression* and *Point-to-Point Voting*. We encourage these two branches to focus on global and local geometry information respectively and finally fuse their estimations to combine their advantages. The whole framework is trained on synthetic dataset generated in simulation without collecting real-world data. Experiments suggest our method achieves high performance on occlusion-robust state estimation of DLOs and can be directly applied in real-world scenarios.

II. PROBLEM STATEMENT

The goal of our method is to estimate the 3-D states of DLOs from point cloud obtained by an RGB-D camera. In this work, we focus on the state estimation problem and assume that the point cloud of the DLO has already been segmented out of the raw full point cloud by RGB image segmentation. We represent the DLO state as a sequence of M nodes uniformly distributed, where M is a pre-defined number of nodes that can sufficiently describe the DLO state. The problem is to estimate the coordinates of the nodes $\mathbf{Y} = [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_M]^T \in \mathbb{R}^{M \times 3}$ from the input point cloud $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N]^T \in \mathbb{R}^{N \times 3}$ where N is the number of the points in the segmented point cloud. Note that the input point cloud \mathbf{X} is unordered, while the order of estimated nodes in \mathbf{Y} from one end to another end has been represented by the index $1, 2, \dots, M$. In addition, the point cloud of the DLO may be fragmentary and noisy because of the occlusions, imperfect segmentation, and depth images of low quality.

III. METHOD

As shown in Fig. 2, our proposed method contains two branches: an *End-to-End Regression* branch and a *Point-to-Point Voting* branch, which focuses on the global and the local geometry information, respectively. Then, a deformable registration module is designed to leverage the advantages of both branches and fuse the two predictions to output the final estimated node sequence.

A. End-to-End Regression

The most straightforward approach is to train an end-to-end network with the point cloud $\mathbf{X} \in \mathbb{R}^{N \times 3}$ as input and the node sequence $\mathbf{Y} \in \mathbb{R}^{M \times 3}$ as output, which is indicated as *End-to-End Regression*. We exploit a PointNet++ [24] encoder denoted as $F(\cdot)$ to extract deep latent features $F(\mathbf{X}) \in \mathbb{R}^{N \times C_{out}}$ of input point cloud \mathbf{X} , which means that each point in input point cloud has a C_{out} -dimensional feature vector. A max pooling layer is then applied to get the global feature $\text{MaxPool}(F(\mathbf{X})) \in \mathbb{R}^{C_{out}}$ which is irrelevant to the input point order. Finally, a fully-connected layer FC_1 predicts the node sequence \mathbf{Y}_{reg}^{pred} . The whole regression network is defined as

$$\mathbf{Y}_{reg}^{pred} = FC_1(\text{MaxPool}(F(\mathbf{X}))). \quad (1)$$

With the ground-truth node coordinates \mathbf{Y}^{gt} , the training loss function for each sample is

$$L_{reg} = \|\mathbf{Y}_{reg}^{pred} - \mathbf{Y}^{gt}\|^2. \quad (2)$$

It is experimentally found that such an end-to-end network can ensure that the estimated DLO shapes are smooth and look like real DLOs even using heavily occluded point cloud input, which suggests that this network can learn the key global characteristic of DLOs well. However, the predictions are often slightly different from the actual states such that they are not sufficiently accurate for applications (see Fig. 5). This phenomenon is believed to be brought about by the feature max pooling operation, which neglects crucial local information for precise estimation.

B. Point-to-Point Voting

To make up for the shortcomings of the end-to-end regression method, we design a point-to-point voting framework to utilize local geometry information, which is inspired by early works [25], [26]. Instead of using max-pooling layers for direct regression, this method generates point-wise predictions $\mathbf{Y}_{vot}^{pred,1}, \mathbf{Y}_{vot}^{pred,2}, \dots, \mathbf{Y}_{vot}^{pred,N}$ from each input point $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N$ and then uses a point-to-point voting scheme to get the final estimation. Specifically, we can regress an offset vector \mathbf{O}_{ij} which predicts the vector beginning from input point \mathbf{x}_i and ending at node \mathbf{y}_j . During inference, the $\mathbf{y}_j^{pred,i}$ can be calculated by adding

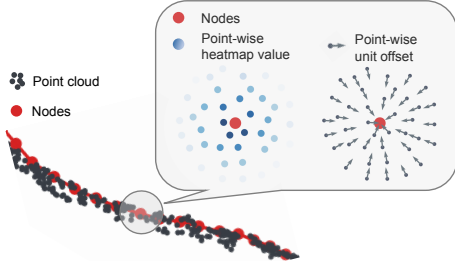


Fig. 3. Demonstration of predicted point-wise heatmap value and unit offset vector. Considering the neighborhood of one node, the points closer to it will have a higher heatmap value (visualized as deeper color), and the unit offset vector represents the normalized direction from the input point to the desired node.

the estimated $\mathbf{O}_{ij}^{\text{pred}}$ to \mathbf{x}_i and then we can apply a voting scheme among the predictions $\mathbf{Y}_{\text{vot}}^{\text{pred},i}$ of all input points to get the final $\mathbf{Y}_{\text{vot}}^{\text{pred}}$.

Similar to [25], we decompose the point-wise offset vector \mathbf{O}_{ij} into a heatmap value H_{ij} for the distance from \mathbf{x}_i to \mathbf{y}_j and a unit offset vector \mathbf{U}_{ij} for the direction, as illustrated in Fig. 3, which makes network training easier. We also restrict the ground truth value of such a point-wise estimation inside the neighborhood of the desired node to exclude noisy estimations of points far away. Given a neighborhood radius r , the ground-truth heatmap value H_{ij}^{gt} is defined as

$$H_{ij}^{\text{gt}} = \begin{cases} 1 - \|\mathbf{x}_i - \mathbf{y}_j\|/r & , \quad \|\mathbf{x}_i - \mathbf{y}_j\| < r, \\ 0 & , \quad \|\mathbf{x}_i - \mathbf{y}_j\| \geq r, \end{cases} \quad (3)$$

and the ground-truth unit offset vector $\mathbf{U}_{ij}^{\text{gt}}$ is defined as

$$\mathbf{U}_{ij}^{\text{gt}} = \begin{cases} (\mathbf{y}_j - \mathbf{x}_i)/\|\mathbf{x}_i - \mathbf{y}_j\| & , \quad \|\mathbf{x}_i - \mathbf{y}_j\| < r, \\ 0 & , \quad \|\mathbf{x}_i - \mathbf{y}_j\| \geq r, \end{cases} \quad (4)$$

We regress the point-wise heatmap $\mathbf{H} \in \mathbb{R}^{N \times M}$ and offset vector $\mathbf{U} \in \mathbb{R}^{N \times M \times 3}$ from the encoded latent feature $\mathbf{F}(\mathbf{X}) \in \mathbb{R}^{N \times C_{\text{out}}}$ using shared point-wise fully-connected layers \mathbf{FC}_2 and \mathbf{FC}_3 , respectively. The prediction of the heatmap and unit offset is denoted as

$$\mathbf{H}^{\text{pred}} = \text{Sigmoid}(\mathbf{FC}_2(\mathbf{F}(\mathbf{X}))), \quad (5)$$

$$\mathbf{U}^{\text{pred}} = \text{Normalized}(\mathbf{FC}_3(\mathbf{F}(\mathbf{X}))), \quad (6)$$

where the Sigmoid activation function and vector normalization are used for regulating the predicted heatmap value $\in [0, 1]$ and the unit offset, respectively. For the point-to-point voting method, the training loss is

$$L_{\text{vot}} = \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M \left[(H_{ij}^{\text{pred}} - H_{ij}^{\text{gt}})^2 + \|\mathbf{U}_{ij}^{\text{pred}} - \mathbf{U}_{ij}^{\text{gt}}\|^2 \right]. \quad (7)$$

The overall training loss for the whole network can be formulated as: $L_{\text{tot}} = L_{\text{reg}} + \alpha L_{\text{vot}}$, where α is a pre-defined coefficient.

As for the inference, the point-wise estimation for node \mathbf{y}_j from input point \mathbf{x}_i is obtained as

$$\mathbf{y}_j^{\text{pred},i} = r(1 - H_{ij}^{\text{pred}})\mathbf{U}_{ij}^{\text{pred}} + \mathbf{x}_i. \quad (8)$$

We assume that the input point closer to the desired node predicts more accurate heatmap and unit offset, so we also use the heatmap value H_{ij}^{pred} as the confidence of the prediction $\mathbf{y}_j^{\text{pred},i}$. To eliminate the influence of unconfident predictions, we choose input points with the highest K heatmap value for the j^{th} node to calculate the final estimation as

$$\mathbf{y}_j^{\text{pred}} = \left(\sum_{i \in \mathcal{K}} H_{ij}^{\text{pred}} \mathbf{y}_j^{\text{pred},i} \right) / \sum_{i \in \mathcal{K}} H_{ij}^{\text{pred}}, \quad (9)$$

where the indexes of the K chosen points form the set \mathcal{K} .

Experiment results show that the point-to-point voting scheme can produce precise state estimations, suggesting that it mainly focuses on local regions and learns the local characteristic for precise estimation well. However, its architecture determines that if there are no enough input points in the local neighborhood of a node because of heavy occlusions, the prediction of the occluded part will be significantly inaccurate (also shown in Fig. 5).

C. Fusion of the Two Branches

To leverage the advantages of both two branches and achieve occlusion-robust state estimations, we further introduce a non-rigid registration-based fusion module. With the known correspondence between the regression and voting results, we aim to estimate a non-rigid transformation from the smooth but imprecise regression results to the accurate unoccluded voting results, as shown in Fig. 4. Note that here we only select the subset of unoccluded nodes from the regression and voting results to exclude unreliable estimations for registration. Then, we transform all the nodes in regression results using the estimated transformation to obtain the final estimations whose unoccluded parts are almost as accurate as the voting results and occluded parts are filled up by the transformed smooth regression results. Details of each step are described as follows:

1) *Select the unoccluded node subset*: Firstly, we have to recognize which node lies in the missing part of the input point cloud. We define the visible possibility p_j of the considered node \mathbf{y}_j as the max value of the heatmap among all input points:

$$p_j = \max_i H_{ij}^{\text{pred}}. \quad (10)$$

The nodes whose p_j is greater than a pre-defined threshold $T \in [0, 1]$ are regarded as unoccluded parts. Denoting the visible possibility of all nodes \mathbf{Y} as \mathbf{P} , the subsets of unoccluded nodes from the regression and voting results for the following non-rigid registration (simplified as *nrr*) are

$$\begin{aligned} \mathbf{Y}_{\text{vot}}^{\text{nrr}} &= \mathbf{Y}_{\text{vot}}^{\text{pred}} [\mathbf{P} > T], \\ \mathbf{Y}_{\text{reg}}^{\text{nrr}} &= \mathbf{Y}_{\text{reg}}^{\text{pred}} [\mathbf{P} > T]. \end{aligned} \quad (11)$$

2) *Estimate the non-rigid transformation*: We utilize a modified *Coherent Point Drift* (CPD) algorithm [27] to estimate the non-rigid transformation with known correspondences. The classical CPD formulates registration as a GMM problem and regards source points ($\mathbf{Y}_{\text{reg}}^{\text{nrr}}$ for us) as the centroids of Gaussian model from which target points ($\mathbf{Y}_{\text{vot}}^{\text{nrr}}$)

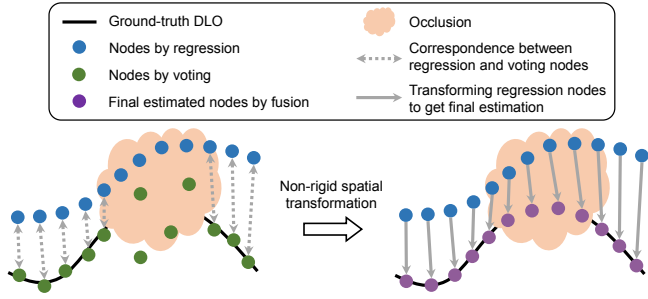


Fig. 4. Illustration of the non-rigid registration-based fusion module. We first identify and exclude the unreliable nodes estimated by voting (green points inside occlusion region). Then, a non-rigid spatial transformation is estimated by applying our modified CPD registration algorithm from the smooth but imprecise regression results (blue points in the left sub-picture) to the precise reliable voting results. Finally, the total sequence estimated by regression is transformed to get the final node sequence (purple points) using this non-rigid transformation.

are sampled. For non-rigid registration, CPD ensures the coherent motion of these centroids by representing the non-linear spatial transformation as

$$\mathcal{T}(\mathbf{Y}_{\text{reg}}^{\text{nrr}}) = \mathbf{Y}_{\text{reg}}^{\text{nrr}} + \mathbf{G}(\mathbf{Y}_{\text{reg}}^{\text{nrr}}) \mathbf{W}, \quad (12)$$

where $\mathbf{G}(\cdot)\mathbf{W}$ represents the displacement function as a Gaussian Radius Basis Function Network. The elements of kernel matrix \mathbf{G} with a parameter β are

$$G_{ij}(\mathbf{Z}) = \exp(-\|\mathbf{z}_i - \mathbf{z}_j\|^2/2\beta^2), \quad (13)$$

and $\mathbf{W} \in \mathbb{R}^{N_z \times D}$ (N_z as the number of points in \mathbf{Z} , and $D = 3$ for this 3-D task) is the weight matrix.

For general-defined non-rigid registration problem, estimating the correspondence between two point sets is very challenging, and CPD uses the EM algorithm to iteratively update the correspondence probability matrix, σ^2 of the Gaussian model, and the weight matrix \mathbf{W} until convergence. In the E-step, CPD fixes the current estimation of σ^2 and \mathbf{W} to estimate more accurate correspondences while σ^2 and \mathbf{W} are later optimized in the M-step. However, the correspondence of our source $\mathbf{Y}_{\text{reg}}^{\text{nrr}}$ and target $\mathbf{Y}_{\text{vot}}^{\text{nrr}}$ has been given naturally by the order of nodes in the sequence. Thus, there is no need to execute the E-step and we can directly update σ^2 and \mathbf{W} iteratively.

Following Eq. (22) and Eq. (23) in [27], we can fix the correspondence probability matrix as an identity matrix and solve \mathbf{W} and σ^2 using

$$(\mathbf{G}(\mathbf{Y}_{\text{reg}}^{\text{nrr}}) + \lambda\sigma^2\mathbf{I})\mathbf{W} = \mathbf{Y}_{\text{vot}}^{\text{nrr}} - \mathbf{Y}_{\text{reg}}^{\text{nrr}}, \quad (14)$$

$$\sigma^2 = \frac{1}{D}(\text{tr}((\mathbf{Y}_{\text{vot}}^{\text{nrr}})^T \mathbf{Y}_{\text{vot}}^{\text{nrr}}) - 2\text{tr}((\mathbf{Y}_{\text{vot}}^{\text{nrr}})^T \mathcal{T}(\mathbf{Y}_{\text{reg}}^{\text{nrr}})) + \text{tr}(\mathcal{T}(\mathbf{Y}_{\text{reg}}^{\text{nrr}})^T \mathcal{T}(\mathbf{Y}_{\text{reg}}^{\text{nrr}}))), \quad (15)$$

where λ is a parameter reflecting the level of smoothness regularization and $D = 3$.

3) *Transform the whole set of regression results:* Finally, we apply the estimated non-rigid transformation to the whole regression node sequence, where the displacement function is given by a Gaussian Radius Basis Function Network and the previously optimized weights \mathbf{W} .

With all the regression nodes $\mathbf{Y}_{\text{reg}}^{\text{pred}}$ and our chosen subset $\mathbf{Y}_{\text{reg}}^{\text{nrr}}$, we reconstruct the kernel matrix $\mathbf{G}(\mathbf{Y}_{\text{reg}}^{\text{pred}}, \mathbf{Y}_{\text{reg}}^{\text{nrr}})$ whose elements are

$$G_{ij}(\mathbf{Z}, \mathbf{Y}) = \exp(-\|\mathbf{z}_i - \mathbf{y}_j\|^2/2\beta^2). \quad (16)$$

The transformed regression node sequence are our final fused state estimations of two branches:

$$\mathbf{Y}_{\text{fus}}^{\text{pred}} = \mathcal{T}(\mathbf{Y}_{\text{reg}}^{\text{pred}}) = \mathbf{Y}_{\text{reg}}^{\text{pred}} + \mathbf{G}(\mathbf{Y}_{\text{reg}}^{\text{pred}}, \mathbf{Y}_{\text{reg}}^{\text{nrr}}) \mathbf{W}. \quad (17)$$

IV. RESULTS

A. Data Collection and Model Training

All the training data are generated in simulations for the convenience of getting the ground-truth node positions. We use the Unity3D [28] as the simulator and the Obi Rope package [29] for simulating DLOs, which is a unified particle physics that models DLOs as chains of oriented particles with stretch, shear, bend, and twist constraints.

During the data collection, RGB-D images captured by a camera and 3-D positions of the DLO particles in the camera frame are recorded per simulation second to generate corresponding point cloud data and ground-truth node positions. The two ends of the DLO are randomly moved using the same strategy as in our previous work [6] which is efficient in covering different shapes. Visualizations of the data examples can be found in the supplemental video on our project website¹. The DLO properties (lengths, radius, and stiffness) and camera poses are also randomized to improve the generalization ability of the model.

We totally generate 25000 frames of data (500 sequences, 50 frames in each sequence) and randomly separated them with a ratio of 0.8 and 0.2 for training and validation, respectively. During the training, the input point cloud is augmented by random jittering with Gaussian noises, random rotation across each axis, and random occlusions. For batch training, we sample $N = 1024$ points from the initial point cloud with the farthest point sampling (FPS) method and predict $M = 50$ nodes from it.

In our implementation, we build up the PointNet++ encoder as the architecture for segmentation tasks in [24], which has 4 point set abstraction levels and 4 feature propagation levels. As for hyper-parameters, we use the Adam optimizer with a initial learning rate 0.01, a weight decay $5e-4$ and a batch size 32 to train the network for 200 epochs. The learning rate decays with a ratio 0.8 every 20 epochs. The training loss weight of two branches is $\alpha = 1.0$. For point-to-point voting branch, we set $r = 0.02$ and $K = 64$. For the registration-based fusion module, we set $T = 0.5$, $\lambda = 0.25$ and $\beta = 0.5$.

B. Simulation Results

We adopt two metrics to evaluate the performance: a) errors of nodes, i.e., the average Euclidean distance between the estimated and ground-truth node positions; and b) uniformity, i.e., the standard deviation of the distances between every adjacent node which illustrates whether the estimated nodes are uniformly distributed.

¹<https://kangchen-lv.github.io/DLO-perception>

TABLE I
QUANTITATIVE RESULTS OF THREE METHODS WITH DIFFERENT LEVELS OF OCCLUSION.

Level of occlusion Method	No occlusion			10% occluded			20% occluded			40% occluded		
	Regression	Voting	Fusion	Regression	Voting	Fusion	Regression	Voting	Fusion	Regression	Voting	Fusion
Errors of all nodes (mm) ↓	25.6	2.7	2.7	25.9	7.3	3.9	26.8	27.5	11.5	32.6	55.6	25.7
Errors of unoccluded nodes (mm) ↓	25.6	2.7	2.7	25.8	3.7	3.8	26.6	9.1	9.8	29.0	17.6	20.0
Errors of occluded nodes (mm) ↓	-	-	-	26.6	39.6	5.2	27.9	101.1	18.2	35.2	112.3	33.4
Uniformity (mm) ↓	1.1	1.1	0.9	1.10	17.5	1.0	1.1	21.3	1.7	1.1	22.6	2.6

TABLE II
ABLATION STUDIES OF DIFFERENT FUSION METHODS.

Level of occlusion	Fusion method	Errors of all nodes (mm) ↓	Uniformity (mm) ↓
No occlusion	Point Concate.	15.9	1.2
	Latent Concate.	18.0	1.2
	CPD	16.5	1.2
	Ours	2.7	0.9
20% occluded	Point Concate.	18.3	1.2
	Latent Concate.	20.7	1.2
	CPD	19.4	1.7
	Ours	11.5	1.7

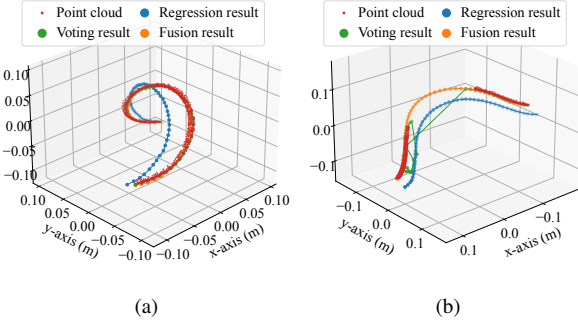


Fig. 5. Visualization of the regression, voting and fusion results of (a) unoccluded point cloud and (b) occluded point cloud.

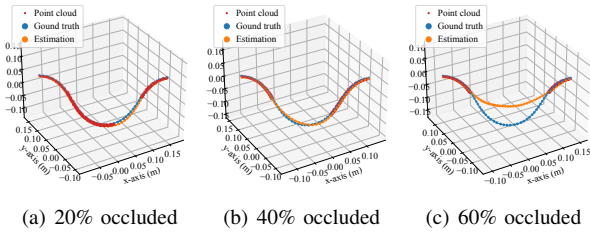


Fig. 6. Estimation results of the same DLO under different occlusion ratio.

1) *Self-comparisons*: We quantitatively evaluate the performance of *End-to-End Regression*, *Point-to-Point Voting* and *Fusion* with different levels of occlusion, in which we further decompose the distance error of all nodes into those of unoccluded and occluded nodes to suggest the specific properties of these methods. As shown in Table I, the errors of regression results are relatively large compared to other methods in unoccluded cases but increase very slowly as the level of occlusion rises. Besides, the regression results keep uniform under all settings, suggesting that it always estimates smooth global DLO shapes even with heavy occlusions. In contrast, although the voting results are precise for complete point cloud input (with an error of 2.7mm), the prediction

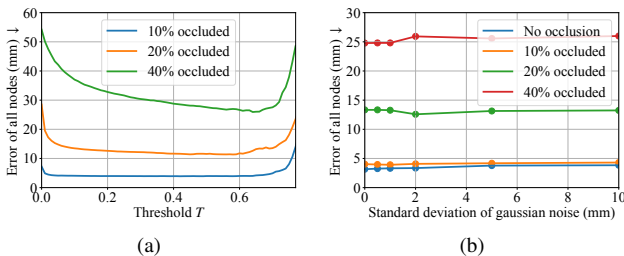


Fig. 7. (a) The relationship between node error and the visible possibility threshold T . (b) The robustness of our method under point cloud noises.

accuracy and uniformity cannot be guaranteed for occluded cases (with an error of 112.3mm for 40% occluded input). Under all settings, our proposed fusion method is the most accurate and robust one, which owes to its ability to combine the advantages of both two branches to predict accurate positions close to the unoccluded voting results and maintain a smooth shape. The uniformity is also well kept even in heavily-occluded scenarios.

Two visualized examples are shown in Fig. 5. Note that the estimated nodes are connected one by one to illustrate their order in all visualizations. It can be seen that the regression results are robust against occlusion but imprecise, while the voting results are accurate outside occlusion but extremely unreliable for occluded parts. However, our fusion method can accurately and robustly estimate the DLO state in both unoccluded and occluded scenarios.

2) *Robustness and sensitivity analysis*: First, we test the relationship between the performance of our method and the occlusion level. As shown in Fig. 6, the occlusion ratio of point cloud is 20%, 40%, and 60%, respectively. When the most of the DLO is occluded (such as Fig. 6(c)), the accurate DLO state is very difficult to estimate since there are multiple possible shapes for the occluded part, but our predicted node sequence still forms a reasonable DLO shape given only the incomplete information of two DLO ends.

In addition, we investigate the sensitivity of our fusion module to the threshold T , which is for determining the visible nodes in Eq. 11, as shown in Fig. 7(a). As T increases, more unreliable voting nodes for occluded parts will be excluded from registration so that the distance error will decrease. However, too many valid nodes will also be excluded and the performance will be hurt severely with a high T value. We also conduct a test of the robustness against Gaussian noises on the input point cloud. As demonstrated in Fig. 7(b), despite that the standard deviation of noise is increasing, our performance still keeps stable.

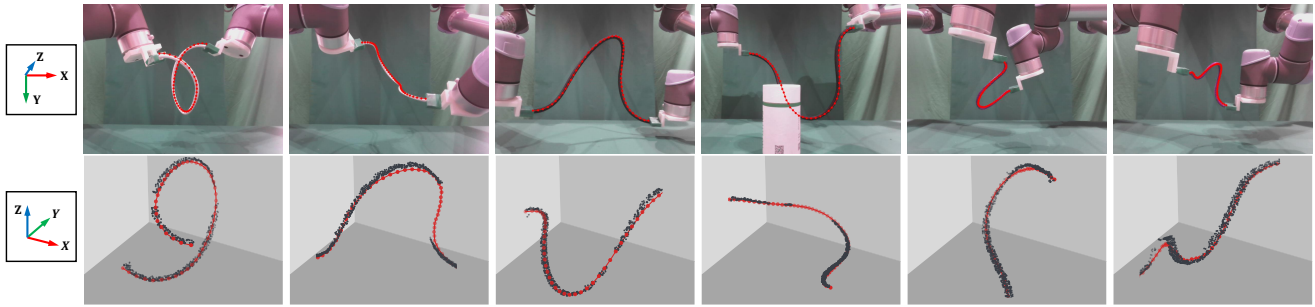


Fig. 8. State estimation results of three different real-world DLOs from occluded and fragmentary point clouds. The top row shows the raw RGB images and the reprojection of the estimated node sequence (in red color); and the bottom row shows the raw point clouds of the DLOs (in black color) and the 3-D positions of the estimated node sequence (in red color). Each column refers to a shape, in which the top image is in the common camera frame and the bottom point cloud is in a top-view frame to better illustrate the Z-axis shape.

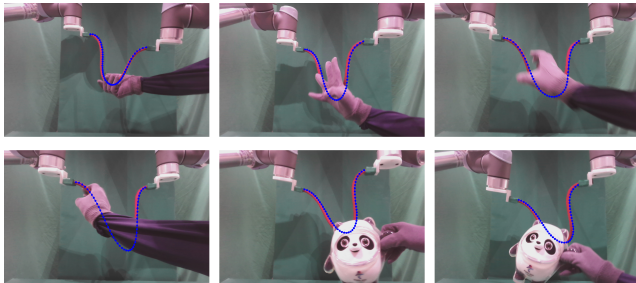


Fig. 9. State estimation results of a moving DLO under partial occlusions by hands or the "Bing Dwen Dwen". The blue points refer to the reprojection results of the estimated 3-D node positions.

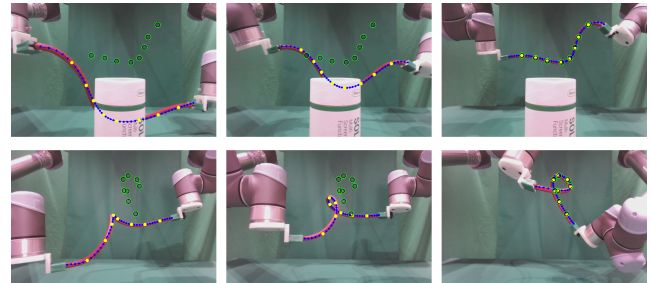


Fig. 10. Applications in downstream DLO shape control task. The blue, yellow, and green points represent the reprojected 3-D node positions, selected control nodes and target positions, respectively.

3) *Ablation on fusion approaches*: Several fusion methods are compared with the proposed modified CPD-based fusion module. We design the baseline fusion methods as follows: a) Point Concatenate: concatenating the voting and regression results and learning a refinement mapping with an MLP; b) Latent Concatenate: concatenating the voting results and the global feature in regression branch to learn the refinement; c) CPD: the classical CPD algorithm [27] with unknown correspondences. Experimental results in Table II suggest that our fusion method outperforms all baselines in both occluded or unoccluded scenarios. Note that the first two baselines perform similar to the regression module so that their uniformity under occlusion is good.

C. Real-World Experiments

We choose three DLOs of different lengths (0.5, 0.7, and 0.3m, respectively) and different materials to examine the generalization ability of our method in real-world applications. The two ends of DLOs are rigidly grasped by dual robot arms and deformed to various complex shapes. The DLO point clouds are first segmented from the background via color thresholding in RGB space and then fed into our two-branch model. Results in Fig. 8 illustrates that our method can be directly applied to estimate the real-world DLO state with small sim-to-real gaps. Even in some cases with self-intersection or occlusion by obstacles, our state estimations are still smooth and precise enough. Experiments on estimating the state of a moving DLO from each frame in a dynamic sequence (see Fig. 9) also suggest the robustness of our method amidst heavy occlusions.

We also integrate our method into the DLO shape control task as the front-end perception module. As shown in Fig. 10, a uniformly-distributed subset of estimated nodes (8 yellow points) is chosen to be controlled to achieve the target positions (corresponding green points) and form the desired DLO shape. The controller is based on our previous work [6] and our method can achieve real-time performance on a GeForce RTX 2060 GPU (average computation time per frame is 34ms). In the presence of occlusions or self-intersections at both initial and middle stage of the manipulation process, our method can steadily output precise DLO states and finally achieve the target positions, which cannot be realized by the existing pure-tracking methods.

V. CONCLUSIONS

In this work, we propose a learning-based method to robustly estimate the 3-D states of DLOs from single-frame point clouds even with heavy occlusions. We use a sequence of ordered nodes as the state representation of DLOs and design a two-branch architecture to estimate it. The two branches are encouraged to utilize the global or local geometry information respectively and their estimations are combined by a fusion module later to get the final output. The simulation and real-world experimental results demonstrate our method can guarantee an accurate and smooth shape of DLOs in both occluded and unoccluded cases with high generalization capability for real-world applications. In the future, we will be devoted to incorporating the temporal information into the framework to get smoother tracking results across a long-term point cloud sequence.

REFERENCES

- [1] 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.
- [2] H. Yin, A. Varava, and D. Kragic, “Modeling, learning, perception, and control methods for deformable object manipulation,” *Science Robotics*, vol. 6, no. 54, p. eabd8803, 2021.
- [3] D. Navarro-Alarcon and Y.-H. Liu, “Fourier-based shape servoing: A new feedback method to actively deform soft objects into desired 2-d image contours,” *IEEE Transactions on Robotics*, vol. 34, no. 1, pp. 272–279, 2017.
- [4] P. Sundaresan, J. Grannen, B. Thananjeyan, A. Balakrishna, M. Laskey, K. Stone, J. E. Gonzalez, and K. Goldberg, “Learning rope manipulation policies using dense object descriptors trained on synthetic depth data,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 9411–9418.
- [5] P. Zhou, J. Zhu, S. Huo, and D. Navarro-Alarcon, “Lasesom: A latent and semantic representation framework for soft object manipulation,” *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5381–5388, 2021.
- [6] 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*, 2022.
- [7] 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)*. IEEE, 2019, pp. 6586–6593.
- [8] 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 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 1337–1343.
- [9] M. Yu, K. Lv, C. Wang, M. Tomizuka, and X. Li, “A coarse-to-fine framework for dual-arm manipulation of deformable linear objects with whole-body obstacle avoidance,” *arXiv preprint arXiv:2209.11145*, 2022.
- [10] H. Dinkel, J. Xiang, H. Zhao, B. Coltin, T. Smith, and T. Bretl, “Wire point cloud instance segmentation from rgbd imagery with mask r-cnn.”
- [11] A. Caporali, K. Galassi, R. Zanella, and G. Palli, “Fastdlo: Fast deformable linear objects instance segmentation,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 9075–9082, 2022.
- [12] D. D. Gregorio, G. Palli, and L. D. Stefano, “Let’s take a walk on superpixels graphs: Deformable linear objects segmentation and model estimation,” in *Asian Conference on Computer Vision*. Springer, 2018, pp. 662–677.
- [13] R. Zanella, A. Caporali, K. Tadaka, D. De Gregorio, and G. Palli, “Auto-generated wires dataset for semantic segmentation with domain-independence,” in *2021 International Conference on Computer, Control and Robotics (ICCCR)*. IEEE, 2021, pp. 292–298.
- [14] S. Huo, A. Duan, C. Li, P. Zhou, W. Ma, H. Wang, and D. Navarro-Alarcon, “Keypoint-based planar bimanual shaping of deformable linear objects under environmental constraints with hierarchical action framework,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 5222–5229, 2022.
- [15] M. Yan, Y. Zhu, N. Jin, and J. Bohg, “Self-supervised learning of state estimation for manipulating deformable linear objects,” *IEEE robotics and automation letters*, vol. 5, no. 2, pp. 2372–2379, 2020.
- [16] M. Wnuk, C. Hinze, A. Lechler, and A. Verl, “Kinematic multibody model generation of deformable linear objects from point clouds,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 9545–9552.
- [17] J. Schulman, A. Lee, J. Ho, and P. Abbeel, “Tracking deformable objects with point clouds,” in *2013 IEEE International Conference on Robotics and Automation*. IEEE, 2013, pp. 1130–1137.
- [18] T. Tang, Y. Fan, H.-C. Lin, and M. Tomizuka, “State estimation for deformable objects by point registration and dynamic simulation,” in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 2427–2433.
- [19] T. Tang and M. Tomizuka, “Track deformable objects from point clouds with structure preserved registration,” *The International Journal of Robotics Research*, p. 0278364919841431, 2018.
- [20] T. Tang, C. Wang, and M. Tomizuka, “A framework for manipulating deformable linear objects by coherent point drift,” *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3426–3433, 2018.
- [21] C. Chi and D. Berenson, “Occlusion-robust deformable object tracking without physics simulation,” in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 6443–6450.
- [22] Y. Wang, D. McConachie, and D. Berenson, “Tracking partially-occluded deformable objects while enforcing geometric constraints,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 14 199–14 205.
- [23] S. Jin, W. Lian, C. Wang, M. Tomizuka, and S. Schaal, “Robotic cable routing with spatial representation,” *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 5687–5694, 2022.
- [24] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, “Pointnet++: Deep hierarchical feature learning on point sets in a metric space,” *Advances in neural information processing systems*, vol. 30, 2017.
- [25] L. Ge, Z. Ren, and J. Yuan, “Point-to-point regression pointnet for 3d hand pose estimation,” in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 475–491.
- [26] C. Wan, T. Probst, L. Van Gool, and A. Yao, “Dense 3d regression for hand pose estimation,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 5147–5156.
- [27] A. Myronenko and X. Song, “Point set registration: Coherent point drift,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 32, no. 12, pp. 2262–2275, 2010.
- [28] U. Technologies. (2022) Unity real-time development platform. [Online]. Available: <https://unity.com/>
- [29] V. M. Studio. (2022) Obi - Unified particle physics for Unity 3D. [Online]. Available: <http://obi.virtualmethodstudio.com/>