

DLOFTBs – Fast Tracking of Deformable Linear Objects with B-splines

Piotr Kicki¹, Amadeusz Szymko¹ and Krzysztof Walas¹

Abstract—While manipulating rigid objects is an extensively explored research topic, deformable linear object (DLO) manipulation seems significantly underdeveloped. A potential reason for this is the inherent difficulty in describing and observing the state of the DLO as its geometry changes during manipulation. This paper proposes an algorithm for fast-tracking the shape of a DLO based on the masked image. Having no prior knowledge about the tracked object, the proposed method finds a reliable representation of the shape of the tracked object within tens of milliseconds. This algorithm’s main idea is to first skeletonize the DLO mask image, walk through the parts of the DLO skeleton, arrange the segments into an ordered path, and finally fit a B-spline into it. Experiments show that our solution outperforms the State-of-the-Art approaches in DLO’s shape reconstruction accuracy and algorithm running time and can handle challenging scenarios such as severe occlusions, self-intersections, and multiple DLOs in a single image.

I. INTRODUCTION

Deformable Linear Objects (DLOs) are a class of objects that are characterized by two main features: deformability, which refers to the fact that the object is not a rigid body and its geometry can change, and linearity, which stands for the fact that the object is elongated and the ratio of its length to its width is substantial [1]. Such objects are ubiquitous both in everyday life and in industry, where one can find ropes, cables, pipes, sutures, etc. While the manipulation of rigid bodies is already solved for a wide range of objects [2], manipulating DLOs is still unsolved even for everyday objects such as cables and hoses. Due to the ubiquity of the DLOs, manipulating them poses a complex and vital challenge, which has been in the scope of researchers for over three decades [3]. The interest in this topic has grown over the last few years, as the automatic wiring harness assembly is crucial for car manufacturers [4], as well as automatic completion of surgical sutures, which could help surgeons [5]. To perform such tasks autonomously, robotic manipulators need to perceive the configuration of the manipulated object, as this is crucial for calculating the adequate control signal. Achieving this requires accurate and fast DLO tracking abilities. However, state-of-the-art DLO tracking algorithms are relatively slow and do not meet the time requirements of the control systems. In addition, they can not properly handle manipulation sequences that contain occlusions [6] and self-intersections [7], or make many assumptions about the model of the tracked object [8], [9].

This work is supported by the European Union’s Horizon 2020 Research and Innovation Programme under grant agreement No 870133, REMODEL.
¹ Institute of Robotics and Machine Intelligence, Poznan University of Technology, Poznan, Poland; e-mail: {name.surname}@put.poznan.pl




Fig. 1. Using the proposed DLO tracking algorithm, one can transform the DLO mask into a low-dimensional B-spline representation in less than 40 ms.

This paper proposes a fast, non-iterative method for estimating a DLO’s shape using a walk through object’s mask and B-spline regression. The proposed algorithm takes as an input the mask of the DLO and returns a sequence of control points of the B-spline curve that approximates the shape of the tracked DLO. Our solution can deterministically identify the shape of a DLO on the HD image within 40 ms while handling non-trivial scenarios, like occlusions, self-intersections, and multiple DLOs in the scene. The general scheme of the proposed approach is presented in Figure 1.

The main contribution of this work is twofold:

- a novel deterministic fast DLO tracking algorithm, which can handle occlusions, self-intersections, and multiple DLOs in the scene while requiring no prior knowledge about the tracked DLO and is faster and more accurate than State-of-the-Art solutions for assumed quality of the output shape,
- dataset of real and artificial 2D and 3D videos and images of several different DLOs, on which we performed a verification of the proposed method and which we share with the community for objective performance evaluation and to encourage the development of real-time DLO tracking¹.

II. RELATED WORK

A. DLO representation

In the literature, there are several ways to represent the geometric shape of the DLO. The most straightforward one is representing it as a sequence of points [10], [11]. However, more complex models are usually necessary for accurate cable modeling and tracking, like a B-spline model with multiple chained random matrices, proposed in [6]. A similar approach, but using Bezier curves and rectangle chains, was proposed in [12], while in [13] NURBS curves were used.

¹https://github.com/PPI-PUT/cable_observer/tree/master

In our research, we use a B-spline representation (similar to the one used in [12], [13]) as it is flexible and enables one to accurately track the shape of a generic DLO while being compact, relatively easy, and cheap to work with. Using this representation, one can build more complex models, which consider the kinematics and dynamics of the DLO [14], [15].

B. DLO tracking

DLO tracking requires transforming the data gathered with sensors into the chosen representation. While there are attempts to use data from tactile sensors [16], the most successful way to perceive the DLO shape is to use vision and depth sensors. One of the most straightforward approaches to DLO shape tracking is to use the fiducial markers located along the DLO, and track them [11] or use them to estimate the shape of a DLO [17]. A similar approach was presented in [18], where colors denote consecutive rope segments. The most common approach is to create a model of the DLO and use images or point clouds as measurements to modify its parameters and track the object deformation iteratively. One of the examples of this approach is the modified expectation-maximization algorithm (EM), proposed in [19], which is used to update the predefined DLO model based on the registered deformations and simulation in the physics engine. Similarly, in [8], the FEM methods were used to track the deformation of the predefined model. Whereas, in [9], a Structure Preserved Registration algorithm with the object represented as a Mixture of Gaussians was used. Authors of [12] performed DLO tracking using the Recursive Bayesian Estimator on the Spatial Distribution Model, built with the Bezier curve and the chain of rectangles. Due to the iterative and often probabilistic character of the model updates, these methods usually have problems tracking rapidly deforming objects and require an appropriate model and accurate initialization. To mitigate the slow initialization problem, authors of [6] used the Euclidean minimum spanning tree and the Breadth-first search method to speed up initialization. However, obtaining the DLO shape estimate still takes hundreds of milliseconds. A much faster EM-based tracking approach, which utilizes a coherent point drift method extended with some geometric-based regularization and physically and geometrically inspired constraints, was presented in [7]. The instance segmentation method for multiple DLOs, which also can serve for tracking, was initially proposed in [20] and extended using Deep Learning solutions in [21] and [22].

The solution presented in this paper utilizes a similar idea to the one presented in [21], [22]. However, using skeletons instead of the super-pixel graphs reduces the computational complexity [21], and the lack of Deep Learning in our solution facilitates better generalization without sacrificing performance [21], [22]. In our work, we do not try to model the DLO, but instead, quickly provide a compact representation of the DLO state, that is consistent and is trackable between frames. Thus, the proposed method can aid the existing model-based methods with fast and accurate structured measurements of the system state.

III. DLO TRACKING

A. Problem Formulation

The problem considered in this paper is to track the DLO on the video sequence. By tracking, we understand transforming consecutive video frames of the DLO's binary mask, obtained from the selected segmentation algorithm, into a 1D curve resembling the object's shape, which representation should be consistent between frames. In this paper, we will not consider the image segmentation problem similarly to [6], [7], [19]. But instead, we will focus on shape tracking only, with the assumption that for homogeneously colored cables, the mask is given by any color-based segmentation algorithm, or in more challenging scenarios, state-of-the-art deep learning method [21] is used.

B. Proposed Method

In this section, we introduce our proposed novel approach to fast tracking of DLO, called DLOFTBs, which by using the walks on the DLO mask's skeleton, enables rapid fitting of the B-spline curve into the masked image of the DLO. The general scheme of the proposed algorithm is presented in Figure 2. To transform the mask image into a B-spline curve, 4 main processing steps are made: (i) morphological open & skeletonization, (ii) walk along the skeleton segments, (iii) filtering and ordering of segments, and finally, (iv) B-spline fitting. In the following subsections, we will describe each of these steps in detail.

1) *Morphological open & skeletonization*: The first operation we perform on the mask of the DLO is a 3×3 (the smallest possible) morphological open. We used it to remove some false positive pixels, which are common because of imperfect segmentation. After that, one of the essential steps in mask processing is performed – skeletonization [23]. This operation takes the mask image as an input and creates its skeleton, i.e., a thin version of the mask, which lies in the geometric centers of the DLO segments, preserves its topology, and reduces its width to a single pixel. This significantly reduces the amount of information about the pixels representing the DLO while preserving the crucial information encoded in central pixels along the DLO. Moreover, using the skeleton, one can easily find the crucial parts of the DLO mask, such as segment endpoints – pixels with only one neighbor or branching – pixels with more than two neighbors. While the segment endpoints will constitute starting points for the walks on segments, the branching points are crucial while performing a walk, as they require the walker to choose one of several possible paths. To avoid this inconvenience, we propose removing branching points and postponing the decision to make connections between segments for further processing. By doing so, we can simplify the segment walk algorithm considerably.

2) *Walk algorithm*: Having the skeleton prepared and segment endpoints determined, we can perform a walk along each segment. To do so, we start from a random segment endpoint and go pixel by pixel till the end of the segment, collecting the subsequent pixel coordinates. Such traversing




Fig. 2. The general scheme of DLOFTBs. The mask image is transformed through several processing stages to obtain a B-spline representation of its shape.

is always possible and unambiguous, as we removed all pixels with more than two neighbors in the previous step. After each walk, we remove two points from the set of endpoints that were the segment's beginning and end. Next, we draw another segment endpoint and perform a walk, which repeats until the end of the segment endpoints. As a result, we obtain a set of paths i.e., ordered lists of pixels representing all segments.

3) *Filtering and ordering of segments*: In this step, we first filter out segments shorter than p pixels, which are likely to represent some artifacts of the mask or resulting from the skeletonization procedure. While this approach may also result in removing a short part of the actual DLO, it will not affect the resultant path significantly, as it will be treated as occlusion and handled by our algorithm at the next stage.

In order to fit a B-spline effectively into a set of segments, we need to order them. As a result of the previous processing step, we have an unordered set of ordered lists of pixels. To order them, we need to find segment endpoint pairs that will most likely connect to each other. While there are many possible criteria and algorithms for deciding about connections, we decided to use a criterion that takes into account both the distance and orientation of the endpoints and is defined by

$$J = mJ_d + (1 - m)J_o, \quad (1)$$

where J_d is a euclidean distance between segment endpoints, J_o is a criterion related to the mutual orientation of the segment endpoints, and $m \in [0; 1]$ is a linear mixing factor. While the definition of J_d is rather straightforward, the exact formula of the J_o is given by

$$J_d = |\pi - \phi_1 - \phi_2|, \quad (2)$$

where ϕ_1, ϕ_2 are approximated orientations of the segment endpoints.

Using the criterion J (1), one can decide about the pairs of the segment endpoints. The most accurate solution would be to check for all possible pairing schemes and find the one with the lowest J . However, it is also the most computationally expensive one, as it requires checking even $(2s - 1)(2s - 3) \dots 1$ pairings, where s is the number of segments. To limit the computational burden, we decided to use a potentially less accurate but much faster approach – a greedy one. Thus, we need to choose at most $s - 1$ connections out of $s(2s - 1)$ pairs of endpoints, considering

the already taken endpoints. Note that we don't want to make highly improbable connections. Therefore, we introduce a threshold J_{th} and consider only the connections for which criterion $J < J_{th}$. This enables us to track a single cable robustly and detect and track multiple DLOs at once. Finally, all detected sequences of segments are separately passed to the B-spline fitting phase, described in the next point. B-spline curves provide a compact representation that can be used to identify DLO instances based on the representation continuity between frames.

4) *B-spline fitting*: To fit the B-spline to the sequence of segments, we need an argument for the B-spline, i.e., the vector t of the relative position of the pixels on the curve we want to define. To do so, we calculate the distance along the segments and euclidean distances between segments and concatenate them into a single vector, the cumulative sum of which serves as the B-spline argument t . Using the Euclidean distance between segments, we introduce an estimate of the distance along the DLO (we do not have access to the true one, as the parts of a DLO are occluded). This procedure prevents sudden changes in the pixel's positions in terms of the B-spline argument t .

Moreover, we need to define the number and positions of knots. In the proposed solution, we defined knots as a sequence of k equidistant, in terms of the element number, elements of the vector t , as it will ensure that the Schoenberg-Whitney conditions [24] are met.

Finally, one can fit two dimensional B-spline using the prepared argument vector t , knots, and x and y coordinates of the ordered pixels. We used cubic splines, as higher continuity is unnecessary for the considered problem.

5) *3D data*: Even though the proposed DLO tracking algorithm is meant to work on images, it can be easily extended to work for the 3D data obtained from the RGBD sensor. In this case, we deal with the mask in the same way as for the 2D case till the moment of the B-spline fitting. Given a sequence of segments in the 2D space, we augment it with the corresponding depth coordinates and then perform the B-spline fitting. Thus, we obtain three dimensional B-spline representing the shape of the cable with respect to the curve length estimate.

IV. EXPERIMENTS

To perform all experiments, we used a single core of the Intel Core i7-9750H CPU and following, heuristically

chosen single set of parameters of our algorithm $m = 0.05$, $p = 10$, and $k = 25$, which are the mixing factor of segments connection criteria, segments length threshold, and the number of knots.

A. Datasets

To evaluate the proposed cable tracking method (DLOFTBs), we conducted several experiments, which show the performance of the proposed algorithm on 4 datasets:

1) *RGB real*: 7 sequences of RGB images (≈ 900 frames in total), collected with the Intel RealSense D435 camera, of a single cable being manipulated by two UR3 manipulators.

2) *RGBD real*: 10 sequences of RGBD images (≈ 2500 frames in total), collected with the Kinect Azure, of the single cable being manipulated by a human.

3) *RGBD artificial*: 5 sequences of the artificially created RGBD images (≈ 1400 frames in total), generated from a reference curve evolving in time. This dataset allows us to compare directly to the reference curve instead to mask.

4) *Ariadne+*: The test set, taken from [21], consists of 62 images of multiple cables. We enriched this dataset with manual annotations of the cable shapes to facilitate direct comparison between curve shapes.

B. Performance criteria

Assessing the quality of the DLO shape tracking is not a trivial task [25], especially when the only ground truth data available is the mask of the DLO (datasets 1 and 2). For dataset 1 and dataset 2, we use two Mean Minimal Distance (MMD) criteria, which build upon the ideas of Modified Hausdorff Distance [26] and are defined by

$$\mathcal{L}_1 = \text{MMD}(\mathcal{M}, \mathcal{C}_d) \quad \text{and} \quad \mathcal{L}_2 = \text{MMD}(\mathcal{C}_d, \mathcal{M}), \quad (3)$$

where

$$\text{MMD}(X, Y) = \frac{1}{|X|} \sum_{x \in X} \min_{y \in Y} d(x, y), \quad (4)$$

where $d(x, y)$ is a Euclidean distance between x and y , \mathcal{M} is a set of pixels belonging to a mask, while \mathcal{C}_d is a set of points on the predicted curve \mathcal{C} .

In turn, for dataset 3 and dataset 4 we have access to the mathematical curve representing the reference shape \mathcal{C}_r . Therefore, we can formulate a much more accurate measure of the performance, which builds upon the Fréchet distance [27], and is defined by

$$\mathcal{L}_3(\mathcal{C}_d, \mathcal{C}_{r_d}) = \frac{F(\mathcal{C}_d, \mathcal{C}_{r_d}) + F(\mathcal{C}_{r_d}, \mathcal{C}_d)}{2}, \quad (5)$$

where \mathcal{C}_{r_d} is a discretized version of the reference path, and where

$$F(X, Y) = \frac{1}{|X|} \sum_{i=0}^{|X|-1} \min_{w \in [0;1]} d(X(i), (1-w)Y(k(i)) + wY(k(i)+1)), \quad (6)$$

where $k(i)$ satisfies $D_Y(k(i)) \leq D_X(i) \leq D_Y(k(i)+1)$ and is monotonically non-decreasing, where $D_X(i)$ is a normalized distance along X curve at i -th discretization point. This function allows for fair alignment of curves despite possible differences in parameterization.

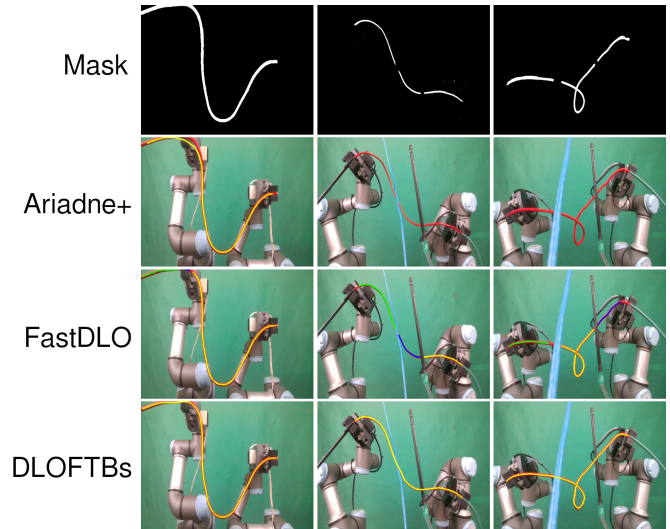


Fig. 3. Comparison of 2D DLO tracking algorithms. Only our approach was able to handle occlusions and self-intersections of the tracked cable.

C. 2D videos of a single cable

In the first stage of the experimental evaluation, we evaluated DLOFTBs on *RGB real* dataset (see Section IV-A.1), with masks generated using hue-based segmentation, and compared it with Ariadne+ [21] and FastDLO [22] learned approaches. We used cables with different widths and lengths and tested all algorithms on challenging setups, shown in Table I, including self-intersection (*scenario 6*) and occlusions (*scenarios 2-6*). Obtained results show that the proposed algorithm achieves the most stable behavior and outperforms all baselines in terms of criterion \mathcal{L}_1 and processing time, and achieves similar results in terms of criterion \mathcal{L}_2 . Huge values of \mathcal{L}_1 for the baselines indicate, that, unlike DLOFTBs, they are unable to cover the whole cable with the predicted spline (extreme case is Ariadne+ which was unable to generate any curve for *scenario 6*). Whereas, small values of \mathcal{L}_2 for almost all cases, ensures that predicted splines do not cover empty areas. Our method achieves a relatively big value of \mathcal{L}_2 only for *scenario 3*, in which large parts of the cable are outside the camera's field of view, therefore, even reasonable and plausible curves generated by the proposed method results in the growth of \mathcal{L}_2 . The behavior of the algorithms for some sample challenging frames is presented in Figure 3. Even though both baselines were provided with a very clean mask of the tracked cable, they were unable to handle occlusions and self-intersections, whereas DLOFTBs handled them perfectly.

D. 2D masks of multiple cables

In this experiment, we evaluated the ability of DLOFTBs to identify multiple cables at once on the masked image and compared directly with Ariadne+ and FastDLO algorithms [21], [22] on the augmented version of the Ariadne+ test set (Section IV-A.4) segmented using DeepLabV3+ network III for all algorithms. The result of this comparison can be found in Table III. We outperformed Ariadne+ and FastDLO in terms of algorithm execution time, and the

TABLE I
PERFORMANCE OF THE PROPOSED DLO TRACKER ON SEVERAL 2D VIDEOS OF DUAL-ARM MANIPULATION.

Algorithm	Scenario	0	1	2	3	4	5	6
DLOFTBs	\mathcal{L}_1 [px]	4.82±0.05	6.15±0.13	2.45±0.09	2.88±0.17	4.97±0.30	2.92±0.08	4.45±0.04
	\mathcal{L}_2 [px]	0.88±2.35	0.39±0.12	1.67±3.13	7.08±14.58	0.56±0.03	0.72±0.07	0.99±0.02
	Time [ms]	38.2±6.4	34.8±4.4	30.0±5.2	33.2±10.0	39.5±4.0	26.0±2.0	35.1±2.0
Ariadne+	\mathcal{L}_1 [px]	8.03±4.38	12.45±16.80	30.01±32.57	29.55±0.64	54.85±71.67	117.60±38.23	–
	\mathcal{L}_2 [px]	0.51±0.8	1.99±4.71	1.22±3.18	47.59±1.55	0.74±0.23	2.63±6.60	–
	Time [ms]	973.9±35.3	956.0±61.3	974.2±30.5	935.5±37.0	962.5±31.0	923.3±30.9	–
FastDLO	\mathcal{L}_1 [px]	9.17±12.8	8.98±10.32	29.3±33.2	29.1±47.5	177.6±16.6	173.7±6.5	171±13
	\mathcal{L}_2 [px]	0.36±0.01	0.37±0.01	0.36±0.01	0.36±0.02	0.37±0.01	0.36±0.01	0.36±0.02
	Time [ms]	86.3±6.3	97.6±10.2	62.2±6.5	64.9±9.5	96.4±9.1	60.8±9.0	88.3±3.6




Fig. 4. Comparison of multiple DLOs tracking methods on Ariadne+ dataset.

accuracy of the DLO shape reconstruction, and scored second in the number of wrongly identified DLOs. The relatively high number of redundant curves fitted by DLOFTBs is a result of extremely noisy masks generated by DeepLabV3+ (see 3rd column of Figure 4). In Figure 4 we present a qualitative analysis of algorithms behavior on 3 challenging images. Ariadne+ has severe problems with handling complex backgrounds and bends at the intersections of cables, while FastDLO cannot solve the intersection in the 2nd image properly and produces a wavy shape for the left cable

in the 1st image. In turn, DLOFTBs generates the most accurate solutions for the first two images, however, if the mask is very noisy (3rd image) it fits curves into linear false positives regions of the mask.

TABLE III
COMPARISON OF DLOFTBs AND ARIADNE+ ALGORITHMS ON MULTIPLE CABLE DETECTION BENCHMARKS.

Algorithm	\mathcal{L}_3	# missing	# redundant	Time [ms]
Ariadne+	45.06	9	16	421.3
FastDLO	51.55	3	75	64.3
DLOFTBs	27.17	3	33	39.2

E. 3D video sequences

To accurately compare the proposed approach with another State-of-the-Art method, we made our algorithm work with 3D data, for which the State-of-the-Art CDCPD2 algorithm [7] was designed.

1) *Real data*: In our experiments, to accurately compare the precision of shape tracking, we adjusted the video frame rate to enable each algorithm to process it at its own pace and reported the times needed to process a single frame. Furthermore, because the performance of the CDCPD2 method is strongly related to the cable length estimate, we tested it for several different lengths for each scenario and reported only the best result, while our algorithm required no parameter tuning.

In Table II we present mean values of the \mathcal{L}_1 and \mathcal{L}_2 errors and mean algorithm running times for 10 scenarios of cable being manipulated by a human, and sample frames for

TABLE II
COMPARISON OF THE DLOFTBs WITH CDCPD2 ALGORITHM ON SEVERAL 3D VIDEOS OF A HUMAN MANIPULATING THE CABLE.

Scenario	0	1	2	3	4	5	6	7	8	9	
CDCPD2 [7]	\mathcal{L}_1	6.26	7.1	6.21	5.73	15.6	5.49	6.03	5.79	19.0	11.19
	\mathcal{L}_2	2.73	2.85	6.55	3.87	16.6	2.83	2.76	2.5	12.3	3.79
	Time [ms]	81	110	125	120	87	85	91	92	92	103
DLOFTBs	\mathcal{L}_1	5.05	5.31	4.41	4.5	4.07	3.79	4.69	4.51	4.79	5.23
	\mathcal{L}_2	1.08	1.04	5.75	2.71	2.08	0.88	1.21	1.03	1.58	0.68
	Time [ms]	26	36	38	33	29	30	19	18	23	26
Sample frame											



Fig. 5. DLOFTBs (green) and CDCPD2 algorithm (red) on real-data 3D cable tracking sequence. With the fast movement of the cable and insufficiently large cable model updates, CDCPD2 could not track the moving cable. In contrast, our approach accurately tracks the cable shape throughout the sequence.

each scenario from the *RGBD real* dataset (see Section IV-A.2). DLOFTB achieves lower errors than CDCPD2 for all considered scenarios and criteria, and its running times are about 3 times shorter. While for many scenarios values of the criteria are rather comparable, there are some cases where the proposed approach outperforms the CDCPD2 by a large margin (scenarios 4, 8, 9). In these cases, the CDCPD2 algorithm lost the track of the cable shape due to the fast movements of the cable (scenarios 8 and 9) or the complexity of the initial shape (scenario 4).

In Figure 5 we present a sample tracking sequence, in which our algorithm is able to keep track of the cable movements and deformations. At the same time, for CDCPD2 the changes are too significant to be possible to follow. Moreover, for the last images in the sequence, CDCPD2 produces a wavy shape, which does not reflect the actual cable shape but does not increase the performance error measures significantly.

2) *Artificial data*: To expose the aforementioned types of errors and accurately measure the quality of tracking, we need to utilize the \mathcal{L}_3 criterion. To do so, we used *RGBD artificial* dataset (see Section IV-A.3), which also includes challenging cases like high cable curvature (*scenarios 0, 1*), self-intersections (*scenarios 1, 2*) and rapid cable moves (*scenarios 3, 4*).

The results of this comparison are presented in Table IV. Also, in this experiment, our proposed approach outperforms the CDCPD2 algorithm. However, the use of the more accurate criterion emphasized the differences between compared methods. DLOFTB achieves mean \mathcal{L}_3 values that are from 6 to 20 times smaller than those achieved by CDCPD2. Minimal mean values of \mathcal{L}_3 show that our approach is, on average, more accurate than the best predictions made by the CDCPD2 in 4 out of 5 scenarios. Moreover, maximal mean values show that throughout 3 out of 5 scenarios, DLOFTBs does not produce any significantly wrong measurements (max mean $\mathcal{L}_3 > 10$), while CDCPD2 does so for all scenarios. In Figure 6 we present a part of the *scenario 2* in which cable was recovering from the self-intersection. Our algorithm was able to accurately track the cable throughout the whole process. In contrast, the CDCPD2 crushed when the cable was occluding itself a moment before the untying and lost track for the rest of the sequence.

V. CONCLUSIONS

This paper proposes a novel approach to DLO tracking on 2D and 3D images and videos called DLOFTBs. Using a




Fig. 6. DLOFTBs (green) and CDCPD2 algorithm (red) on artificially generated data of 3D cable tracking sequence. In the complex case of deforming the cable from the self-intersection into the basic shape, our approach accurately tracks the deformation throughout the whole sequence. In contrast, the CDCPD2 started with accurate tracking but degraded significantly after the last moment when the cable was self-intersecting.

segmented mask of the cable, we can precisely fit a B-spline representation of its shape within tens of milliseconds. The experimental analysis showed that DLOFTB is accurate and can handle tedious cases like occlusions, self-intersections, or even multiple DLOs at one time. Moreover, it outperforms the State-of-the-Art DLO tracking algorithms CDCPD2 [7], Ariadne+ [21], and FastDLO [22] in all considered scenarios both in terms of the quality of tracking, identification of multiple cables, and algorithm running time. Moreover, the proposed solution is able to solve all aforementioned problems with a single set of parameters and does not require any training. Thus it does not depend on the training data, and, unlike the CDCPD2, does not need any prior information about the DLO.

Our method was extensively tested against algorithmic and learned methods. The weakness of the approaches that utilize learning is that they have substantial problems with generalization and are not working with the data outside the training set distribution. Our approach is not suffering from this issue, so it is better suitable for robotics. We claim that there is still some space for non-deep-learning approaches, which are better in generalization and are fully explainable.

TABLE IV

COMPARISON OF DLOFTBs AND CDCPD2 ALGORITHMS ON SEVERAL ARTIFICIALLY GENERATED 3D CABLE MANIPULATION SCENARIOS.

Scenario	0	1	2	3	4
mean \mathcal{L}_3	7.94	10.62	26.8	23.9	30.4
CDCPD2 max mean \mathcal{L}_3	128	75	198	82	100
min mean \mathcal{L}_3	1.96	2.96	2.02	2.82	2.48
mean \mathcal{L}_3	1.59	4.22	1.99	2.42	2.47
DLOFTBs max mean \mathcal{L}_3	2.6	38.1	47.7	5.75	6.91
min mean \mathcal{L}_3	1.23	1.57	1.1	1.48	1.34
Sample frame					

REFERENCES

- [1] 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 Int. J. of Rob. Res.*, vol. 37, no. 7, pp. 688–716, 2018.
- [2] A. Billard and D. Kragic, "Trends and challenges in robot manipulation," *Science*, vol. 364, no. 6446, p. 8414, 2019.
- [3] M. Inaba and H. Inoue, "Rope handling by a robot with visual feedback," *Advanced Rob.*, vol. 2, no. 1, pp. 39–54, 1987.
- [4] P. Kicki, M. Bednarek, P. Lembicz, G. Mierzwiak, A. Szymko, M. Kraft, and K. Walas, "Tell me, what do you see?—interpretable classification of wiring harness branches with deep neural networks," *Sensors*, vol. 21, no. 13, 2021.
- [5] S. Sen, A. Garg, D. V. Gealy, S. McKinley, Y. Jen, and K. Goldberg, "Automating multi-throw multilateral surgical suturing with a mechanical needle guide and sequential convex optimization," in *2016 IEEE Int. Conf. on Rob. and Aut. (ICRA)*, 2016, pp. 4178–4185.
- [6] G. Yao, R. Saltus, and A. Dani, "Shape estimation for elongated deformable object using B-spline chained multiple random matrices model," *Int. J. of Intell. Rob. and Appl.*, vol. 4, no. 4, pp. 429–440, Dec 2020.
- [7] Y. Wang, D. McConachie, and D. Berenson, "Tracking partially-occluded deformable objects while enforcing geometric constraints," in *IEEE Int. Conf. on Rob. and Aut. (ICRA)*, 2021, pp. 14 199–14 205.
- [8] A. Petit, V. Lippiello, and B. Siciliano, "Real-time tracking of 3d elastic objects with an rgb-d sensor," in *2015 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2015, pp. 3914–3921.
- [9] T. Tang and M. Tomizuka, "Track deformable objects from point clouds with structure preserved registration," *Int. J. of Rob. Res.*, 2019.
- [10] H. Han, G. Paul, and T. Matsubara, "Model-based reinforcement learning approach for deformable linear object manipulation," in *2017 13th IEEE Conf. on Aut. Sci. and Eng. (CASE)*, 2017, pp. 750–755.
- [11] Y. Lai, J. Poon, G. Paul, H. Han, and T. Matsubara, "Probabilistic pose estimation of deformable linear objects," in *2018 IEEE 14th Int. Conf. on Aut. Sci. and Eng. (CASE)*, 2018, pp. 471–476.
- [12] A. Zea, F. Faion, and U. D. Hanebeck, "Tracking elongated extended objects using splines," in *2016 19th Int. Conf. on Information Fusion (FUSION)*, 2016, pp. 612–619.
- [13] R. C. Jackson, R. Yuan, D. Chow, W. Newman, and M. Çavuşoğlu, "Real-time visual tracking of dynamic surgical suture threads," *IEEE Trans. on Aut. Sci. and Eng.*, vol. 15, no. 3, pp. 1078–1090, 2018.
- [14] G. Palli, "Model-based manipulation of deformable linear objects by multivariate dynamic splines," in *2020 IEEE Conf. on Industrial Cyberphysical Systems (ICPS)*, vol. 1, 2020, pp. 520–525.
- [15] J. Spillmann and M. Teschner, "Corde: Cosserat rod elements for the dynamic simulation of one-dimensional elastic objects," in *Proceedings of the 2007 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, ser. SCA '07, 2007, p. 63–72.
- [16] J. Sanchez, C. M. Mateo, J. A. Corrales, B. Bouzgarrou, and Y. Mezouar, "Online shape estimation based on tactile sensing and deformation modeling for robot manipulation," in *2018 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2018, pp. 504–511.
- [17] A. Borum, D. Matthews, and T. Bretl, "State estimation and tracking of deforming planar elastic rods," in *2014 IEEE Int. Conf. on Rob. and Aut. (ICRA)*, 2014, pp. 4127–4132.
- [18] W. H. Lui and A. Saxena, "Tangled: Learning to untangle ropes with rgb-d perception," in *2013 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2013, pp. 837–844.
- [19] J. Schulman, A. Lee, J. Ho, and P. Abbeel, "Tracking deformable objects with point clouds," in *2013 IEEE Int. Conf. on Rob. and Aut.*, 2013, pp. 1130–1137.
- [20] D. D. Gregorio, G. Palli, and L. di Stefano, "Let's take a walk on superpixels graphs: Deformable linear objects segmentation and model estimation," in *ACCV*, 2018, pp. 662–677.
- [21] A. Caporali, R. Zanella, D. De Gregorio, and G. Palli, "Ariadne+: Deep learning-based augmented framework for the instance segmentation of wires," *IEEE Transactions on Industrial Informatics*, pp. 1–1, 2022.
- [22] 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.
- [23] T. Lee, R. Kashyap, and C. Chu, "Building skeleton models via 3-d medial surface axis thinning algorithms," *CVGIP: Graphical Models and Image Processing*, vol. 56, no. 6, pp. 462–478, 1994.
- [24] I. J. Schoenberg and A. Whitney, "On polya frequency function. iii. the positivity of translation determinants with an application to the interpolation problem by spline curves," *Trans. of the American Mathematical Society*, vol. 74, no. 2, pp. 246–259, 1953.
- [25] R. Veltkamp, "Shape matching: similarity measures and algorithms," in *Proceedings International Conference on Shape Modeling and Applications*, 2001, pp. 188–197.
- [26] M. Dubuisson and A. Jain, "A modified hausdorff distance for object matching," in *Int. Conf. on Patt. Recog.*, vol. 1, 1994, pp. 566–568.
- [27] M. Fréchet, "Sur quelques points du calcul fonctionnel," *Rendiconti del Circolo Matematico di Palermo (1884-1940)*, vol. 22, pp. 1–72.