DLO-Splatting: Tracking Deformable Linear Objects Using 3D Gaussian Splatting

Holly Dinkel^{*,1}, Marcel Büsching^{*,2}, Alberta Longhini², Brian Coltin³, Trey Smith³ Danica Kragic², Mårten Björkman², Timothy Bretl¹

Abstract—This work presents DLO-Splatting, an algorithm for estimating the 3D shape of Deformable Linear Objects (DLOs) from multi-view RGB images and gripper state information through prediction-update filtering. The DLO-Splatting algorithm uses a position-based dynamics model with shape smoothness and rigidity dampening corrections to predict the object shape. Optimization with a 3D Gaussian Splatting-based rendering loss iteratively renders and refines the prediction to align it with the visual observations in the update step. Initial experiments demonstrate promising results in a knot tying scenario, which is challenging for existing vision-only methods.

I. INTRODUCTION

This work presents DLO-Splatting, an algorithm for tracking the shapes of Deformable Linear Objects (DLOs) such as rope for manipulation shape planning and control tasks such as knot tying [1-5]. These tasks are common in applications including robotic surgery, industrial automation, and human habitat maintenance [6-10]. Inspired by recent work in cloth shape estimation [11, 12], the DLO-Splatting algorithm approaches the tracking problem using dynamics to predict the state and Gaussian Splatting-inspired rendering to update the state to track through visually tricky phenomena, such as accurate topology tracking through dense knotting. DLO-Splatting replaces the Graph Neural Network dynamics model used in the prediction step of these previous works with a position-based dynamics model to bypass the reliance on model training and learning. This work makes the following contributions:

- DLO-Splatting synthesizes observations from multiple perspectives to track visually complex topologies which cannot be disambiguated by vision-based tracking alone.
- DLO-Splatting uses position-based dynamics with shape smoothness and rigidity dampening corrections to predict the state of the rope one time step into the future without relying on any training data or model learning.
- DLO-Splatting uses a 3D Gaussian Splatting-based rendering loss to iteratively refine the predicted state based on new sensor information.

*Equal Contribution

¹Holly Dinkel and Timothy Bretl are with the Department of Aerospace Engineering and the Coordinated Science Laboratory at the University of Illinois Urbana-Champaign, Urbana, IL 61801 USA. e-mail: {hdinkel2, tbretl}@illinois.edu

²Marcel Büsching, Alberta Longhini, Danica Kragic, and Mårten Björkman are with the Division of Robotics, Perception and Learning at KTH Royal Institute of Technology, Stockholm, 114 28 Sweden, e-mail: {busching, albertal, dani, celle}@kth.se.

³Brian Coltin and Trey Smith are with the NASA Ames Research Center, Moffett Field, CA, 94035 USA. e-mail: {brian.coltin, trey.smith}@nasa.gov.

Existing works demonstrate DLO tracking on real data [13–21]. These algorithms emphasize vision without physics to track the shape of moving DLOs under various types of occlusion, or otherwise use fiducial markers and physics simulation to ground estimation. Other work extends real-time single-DLO tracking to multi-DLO tracking, estimating the state of multiple DLOs as they are braided while maintaining correct topologies for each object [22]. One of the most challenging aspects of perceiving thin objects, including DLOs, is segmenting them from each other or from the background.

II. RELATED WORK

Existing work on learning unknown object dynamics focuses on revealing properties such as inertial parameters or friction through interaction [23, 24] while analytical methods using mass-spring or position-based dynamics integrate the shape over time, assuming a small integration time step [3, 25– 27]. Predicting deformable object shapes enables the design of physics-informed state estimators [12, 20, 28–30].

Vision-based 2D tracking is a well-explored problem with many proposed solutions [31–35]. Image-only 2D trackers struggle with occlusions and require additional multi-view or precise depth information to project the results to 3D. Rendering-based losses for tracking have been explored by extensions of NeRF to state estimation [36], dynamic scenes [37– 41], and dynamic extensions of 3D Gaussian Splatting [11, 42– 44]. These methods enable 3D tracking of scene content but often struggle with real-world applicability and long training times, often in the magnitude of hours [45, 46].

III. THE DLO-SPLATTING ALGORITHM

The DLO-Splatting algorithm estimates the 3D state $\hat{\mathbf{X}}^{t+1} \in \mathbb{R}^{N \times 3}$ of a DLO represented by N nodes over time given an initial state \mathbf{X}^0 , the actions applied to the DLO \mathbf{a}^t , and visual observations consisting of RGB images $\mathcal{I}_k^t \in \mathbb{R}^{H \times W \times 3}$, where H and W are the camera resolution, from K cameras with camera perspective projection matrices $\mathbf{P}_k \in \mathbb{R}^{3 \times 4}$. The DLO-Splatting algorithm estimates the 3D state of a DLO using a prediction-update framework akin to Bayesian filtering. First, the DLO state is predicted using position-based dynamics (Section III-A). This state is iteratively updated using 3D Gaussian Splatting-based rendering (Section III-B).

A. Prediction with Position-Based Dynamics

Position-Based Dynamics (PBD) derived from physics firstprinciples are used to predict the next state of the rope given the current estimated state and a gripper action. The position of the grasped node, \mathbf{x}_{a}^{t} , is computed as the closest node in the set of nodes describing the shape of the rope, $\mathbf{X}^t \in \mathbb{R}^{N \times 3}$, to the position of the center of the gripper, \mathbf{p}_q^t , as

$$\mathbf{x}_{g}^{t} = \operatorname*{argmin}_{i} \|\mathbf{x}_{i}^{t} - \mathbf{p}_{g}^{t}\|.$$
(1)

The predicted position of the rope given gripper action $\mathbf{a}^t = \mathbf{p}_g^t - \mathbf{p}_g^{t-1}$ is integrated through Verlet velocity integration of each node in \mathbf{X}^t for time step Δt according to

$$\mathbf{X}^{t+1} = \mathbf{X}^t + \left(\mathbf{X}^t - \mathbf{X}^{t-1}\right)\Delta t + \frac{1}{2}\mathbf{F}^t\Delta t^2, \qquad (2)$$

where each $\mathbf{f}_i^t \in \mathbf{F}^t$ for $\mathbf{F}^t \in \mathbb{R}^{N \times 3}$ are the summed external forces acting on node \mathbf{x}_i^t including gravity $\mathbf{f}_{i,g}^t$, the normal force $\mathbf{f}_{i,N}^t$, and friction $\mathbf{f}_{i,f}^t$. Given the constant mass of each node m_i as the total mass of the object m divided by the number of nodes N, gravitational constant g, and friction coefficient μ_f , the external forces acting on node i are

where binary c_i^t indicates whether node *i* is in contact with the manipulation plane, i.e., $\mathbf{x}_{i,z}^t \leq 0$, and $\mathbf{v}_{i,xy}^t$ is the velocity of node \mathbf{x}_i^t if $c_i^t = 1$. Furthermore, $\mathbf{F}^t = \mathbf{F}_g^t + \mathbf{F}_N^t + \mathbf{F}_f^t$.

The rope length is constrained to maintain segment lengths by applying constraint forces, $\mathbf{n}_i^t \in \mathbf{N}^t$, as

$$\mathbf{X}^t = \mathbf{X}^t + \mathbf{\Lambda}^t \mathbf{N}^t. \tag{4}$$

The parameter $\lambda_i^t \in \mathbf{\Lambda}^t$ is used to satisfy the constraint on segment length, *L*, given by

$$\lambda_i^t = \frac{L - \|\mathbf{x}_{i+1}^t - \mathbf{x}_i^t\|}{2},\tag{5}$$

and the constraint direction is

$$n_{i}^{t} = \frac{\mathbf{x}_{i+1}^{t} - \mathbf{x}_{i}^{t}}{\|\mathbf{x}_{i+1}^{t} - \mathbf{x}_{i}^{t}\|}.$$
(6)

When the grasped node moves to the new gripper position, the position correction $\mathbf{x}_g^{t+1} = \mathbf{p}_g^t + \mathbf{a}^t \Delta t$ must also be propagated through the rope. After applying Verlet velocity integration, the positions of nodes along the rope are adjusted to keep the segment lengths between adjacent nodes constant. Given lengths \mathbf{l}^{t+1} ,

$$l_i^{t+1} = \|\mathbf{x}_{i+1}^{t+1} - \mathbf{x}_i^{t+1}\|,$$
(7)

and desired length L, length corrections, $\Delta \mathbf{l}^{t+1}$ are computed for each nonzero length in \mathbf{l}^{t+1} as

$$\Delta l_i^{t+1} = L - l_i^{t+1} \times \frac{\mathbf{x}_{i+1}^{t+1} - \mathbf{x}_i^{t+1}}{l_i^{t+1}}$$
(8)

and applied to the positions of the nodes on the segment as

$$\mathbf{x}_{i}^{t+1} \leftarrow \mathbf{x}_{i}^{t+1} - \frac{1}{2}\Delta l_{i}^{t+1} \\
 \mathbf{x}_{i+1}^{t+1} \leftarrow \mathbf{x}_{i+1}^{t+1} + \frac{1}{2}\Delta l_{i}^{t+1}.$$
(9)

The predicted nodes are finally $\mathbf{X}_{PBD}^t = [\mathbf{x}_1^t, \cdots, \mathbf{x}_N^t]^{\mathsf{T}}$.



Fig. 1. **The DLO-Splatting Algorithm.** The DLO-Splatting algorithm estimates the 3D state of a DLO using a prediction-update framework akin to Bayesian filtering. The DLO state is predicted using position-based dynamics and is iteratively updated using 3D Gaussian Splatting-based rendering.

B. Update with 3D Gaussian Splatting-Based Rendering

In 3D Gaussian Splatting, the scene is represented as a set of M 3D Gaussian distributions, where each distribution $G_j(\mu_j, \Sigma, \mathbf{c}_j, o_j)$ is parameterized by its position μ_j , covariance Σ , color \mathbf{c}_j , and opacity o_j [42]. The 3D Gaussians are distributed on the centerline of the DLO, such that each segment contains d equally spaced Gaussians. This representation establishes a direct coupling between the 3D Gaussians and the DLO's nodes, allowing the Gaussians positions to deform along with the rope. To capture the shape of the rope, we model the size of all Gaussian with covariance matrix $\Sigma = \sigma \mathbf{1}^{3\times3}$ where σ corresponds to the diameter of the rope, resulting in equal-sized spherical Gaussians. Stochastic gradient descent optimizes the Gaussian parameters using the \mathcal{L}_{obs} rendering loss, defined by

$$\mathcal{L}_{obs} = \|\mathcal{I}_k^t - \tilde{\mathcal{I}}_k^t\|_2, \tag{10}$$

between ground truth image \mathcal{I}_k^t and rendered image $\tilde{\mathcal{I}}_k^t$. Rendering $\tilde{\mathcal{I}}_k^t$ projects the 3D Gaussians onto the image plane of perspective k using \mathbf{P}_k as

$$\hat{\mathcal{I}}_{k}^{t} = h(\mathbf{X}_{\text{PBD}}^{t}, \mathbf{P}_{k}) = [\mathbf{c}_{j}(\mathbf{u}, G_{j})]_{\forall \mathbf{u} \in \mathcal{I}_{k}^{t}}.$$
 (11)

The Gaussians are ordered by their distance to the image plane $(G_0 \text{ is nearest})$ which allows for their aggregation using α blending. The value of blending factor α_j is estimated for each Gaussian by multiplying its opacity by its Gaussian function,

$$\alpha_{j}(\mathbf{u}) = o_{j} \exp\left(-\frac{1}{2}\left(\mathbf{u} - \boldsymbol{\mu}_{j}^{\prime}\right)^{\mathsf{T}} (\boldsymbol{\Sigma}^{\prime})^{-1} \left(\mathbf{u} - \boldsymbol{\mu}_{j}^{\prime}\right)\right), \quad (12)$$

where μ'_j and Σ'_j are the 2D projections of the Gaussian parameters onto the image plane. Rendering evaluates for each



Fig. 2. **Qualitative Results.** The DLO-Splatting algorithm is compared to TrackDLO on qualitative tracking during a cross move commonly performed in knot tying. During this move, one tip of the DLO is moved through a loop to create an additional crossing in the topology. After eight seconds, both DLO-Splatting and TrackDLO failed to estimate the correct topology of the DLO, however DLO-Splatting succeeds in tracking the grasped tip.

 $\mathbf{u} \in \mathcal{I}_k^t$ the rendering function given by

$$\mathbf{c}(\mathbf{u}) = \sum_{j \in M} \mathbf{c}_j \alpha_j(\mathbf{u}) \prod_{l=0}^{j-1} (1 - \alpha_l(\mathbf{u})).$$
(13)

For the state update, DLO-Splatting adapts this static Gaussian Splatting formulation to allow for fast optimization of dynamic rope states. The key difference is DLO-Splatting only models the appearance of the rope with Gaussians, while the geometry of the rope is modeled with the node chain \mathbf{X}^t as shown in Figure (1). The rendering loss in Eq. 10 can be used to iteratively update the node positions as

$$\hat{X}^{t+1} = X_{\text{GS}}^{t+1} = X_{\text{PBD}}^{t+1} + \Delta X^{t+1}, \quad (14)$$

where
$$\Delta \mathbf{X}^{t+1} = \underset{\Delta \mathbf{X}}{\operatorname{argmin}} \mathcal{L}_{\operatorname{obs}} \left(\mathbf{X}_{\operatorname{PBD}}^{t+1} + \Delta \mathbf{X} \right)$$
 (15)

is optimized using Stochastic Gradient Descent (SGD). The original Gaussian Splatting implementation uses directiondependent spherical harmonics as color representation, enabling the modeling of non-Lambertian effects such as reflections. Since DLOs do not exhibit these visual effects, DLO-Splatting only models their RGB colors instead. The initial RGB color values are the averages of the initial segmentation of the rope. Representing the object thickness with spherical Gaussians also allows for ignoring Gaussian rotations during deformation, simplifying optimization. This work uses *gsplat* as the rasterizer for Gaussian Splatting rendering for improved speed and memory efficiency [47].

IV. DEMONSTRATION

The performance of the DLO-Splatting algorithm is compared to TrackDLO for tracking the state of a rope during a cross move—a common final step in knot tying [48, 49]. The demonstration uses three cameras: an Intel RealSense D435, an Intel RealSense D405, and a Luxonis Oak-D Pro, each recording at 30 Hz within a calibrated workspace [50]. An ABB IRB120 robot equipped with an OnRobot 2FG7 gripper manipulates the DLO. Joint states from the robot are recorded at 10 Hz and used to compute forward kinematics in the PyBullet physics simulator to obtain the pose of the grasp center point. The data for the move is saved in a Robot Operating System (ROS) bag file [51]. Both methods are initialized using TrackDLO [13]. As shown in Figure 2, DLO-Splatting more accurately tracks the rope tip as it is lifted outof-plane by the robot. However, both methods ultimately fail to recover the rope's final shape and topology.

V. CONCLUSION

This work introduced DLO-Splatting, an in-progress method for estimating the state of a DLO relying on physics and rendering. The DLO-Splatting algorithm is inspired by recent advancements in cloth and DLO state estimation, aiming to handle especially challenging scenarios where vision-only methods often fail. Current limitations include difficulty modeling and constraining self-intersections, a slow update rate of 1 Hz, and reduced performance during occlusions-particularly from the gripper. To address these challenges in the prediction step, future work could use a higher-fidelity simulation to model deformation [52] or sample physics prediction at a much higher rate to improve resolution of physics constraints. To reduce errors caused by visual artifacts and occlusions, future work may incorporate a Gaussian Splatting-based representation of the entire scene into the rendering process [29]. Additionally, the current implementation uses object masking during the update step, and future work could use more recent DLO instance segmentation methods to track multiple objects simultaneously [53–59]. While DLO-Splatting is still early in development, it demonstrates a promising direction for state estimation in tasks involving complex DLO geometries and dynamics-such as knots-where parts of the object are selfcolliding and heavily occluded, making purely vision-based tracking unreliable.

ACKNOWLEDGMENTS

The NASA Space Technology Graduate Research Opportunity award 80NSSC21K1292 supported Holly Dinkel and the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation supported Marcel Büsching. Holly, Marcel, and Alberta thank the 2024 KTH Royal Institute of Technology Robotics, Perception, and Learning Summer School organizing team, including Anna Gautier, Fereidoon Zangeneh, Grace Hung, Jens Lundell, Maciej Wozniak, Miguel Vasco, Olov Andersson, and Patric Jensfelt, for facilitating this collaboration.

References

- [1] M. Yan, G. Li, Y. Zhu, and J. Bohg, "Learning Topological Motion Primitives for Knot Planning," in *IEEE/RSJ Int. Conf. Intell. Robot. Sys. (IROS)*, 2020. 1
- [2] R. Lagneau, A. Krupa, and M. Marchal, "Automatic Shape Control of Deformable Wires Based on Model-Free Visual Servoing," *IEEE Robot. Autom. Lett.*, vol. 5, no. 4, pp. 5252– 5259, 2020.
- [3] H. Yin, A. Varava, and D. Kragic, "Modeling, Learning, Perception, and Control Methods for Deformable Object Manipulation," in *Sci. Robot.*, vol. 6, May 2021, pp. 1–16. 1
- [4] M. Yu, H. Zhong, and X. Li, "Shape Control of Deformable Linear Objects with Offline and Online Learning of Local Linear Deformation Models," in *IEEE Int. Conf. Robot. Autom.* (*ICRA*), 2022, pp. 1337–1343.
- [5] S. Jin, W. Lian, C. Wang, M. Tomizuka, and S. Schaal, "Robotic Cable Routing with Spatial Representation," *IEEE Robot. Autom. Lett.*, vol. 7, no. 2, pp. 5687–5694, 2022.
- [6] B. Lu, H. K. Chu, and L. Cheng, "Dynamic Trajectory Planning for Robotic Knot Tying," in *IEEE Int. Conf. Real-Time Comput. Robot. (RCAR)*, 2016, pp. 180–185.
- [7] V. Viswanath, J. Grannen, P. Sundaresan, B. Thananjeyan, A. Balakrishna, E. Novoseller, J. Ichnowski, M. Laskey, J. E. Gonzalez, and K. Goldberg, "Disentangling Dense Multi-Cable Knots," in *IEEE/RSJ Int. Conf. Intell. Robot. Sys. (IROS)*, 2021, pp. 3731–3738.
- [8] A. Keipour, M. Bandari, and S. Schaal, "Efficient Spatial Representation and Routing of Deformable One-Dimensional Objects for Manipulation," *IEEE/RSJ Int. Conf. Intell. Robot. Sys. (IROS)*, pp. 211–216, 2022.
- [9] J. Guo, J. Zhang, D. Wu, Y. Gai, and K. Chen, "An Algorithm Based on Bidirectional Searching and Geometric Constrained Sampling for Automatic Manipulation Planning in Aircraft Cable Assembly," J. Manuf. Syst., vol. 57, pp. 158–168, 2020.
- [10] P. Chang and T. Padır, "Model-Based Manipulation of Linear Flexible Objects: Task Automation in Simulation and Real World," *Machines*, vol. 8, 2020. 1
- [11] B. P. Duisterhof, Z. Mandi, Y. Yao, J.-W. Liu, J. Seidenschwarz, M. Z. Shou, R. Deva, S. Song, S. Birchfield, B. Wen, and J. Ichnowski, "DeformGS: Scene Flow in Highly Deformable Scenes for Deformable Object Manipulation," *IEEE Int. Work-shop Algorithmic Found. Robot. (WAFR)*, 2024. 1
- [12] A. Longhini, M. Büsching, B. P. Duisterhof, J. Lundell, J. Ichnowski, M. Björkman, and D. Kragic, "Cloth-Splatting: 3D Cloth State Estimation from RGB Supervision," in *Conf. Robot. Learn. (CoRL)*, 2024. 1
- [13] J. Xiang, H. Dinkel, H. Zhao, N. Gao, B. Coltin, T. Smith, and T. Bretl, "TrackDLO: Tracking Deformable Linear Objects Under Occlusion With Motion Coherence," *IEEE Robot. Autom. Lett.*, vol. 8, no. 10, pp. 6179–6186, 2023. 1, 3
- [14] C. Chi and D. Berenson, "Occlusion-Robust Deformable Object Tracking Without Physics Simulation," in IEEE/RSJ Int. Conf. Intell. Robot. Sys. (IROS), 2019, pp. 6443–6450.
- [15] Y. Wang, D. M^cConachie, and D. Berenson, "Tracking Partially-Occluded Deformable Objects While Enforcing Geometric Constraints," in *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2021, pp. 14 199–14 205.
- [16] T. Tang, C. Wang, and M. Tomizuka, "A Framework for Manipulating Deformable Linear Objects by Coherent Point Drift," *IEEE Robot. Autom. Lett.*, vol. 3, no. 4, pp. 3426–3433, 2018.
- [17] T. Tang and M. Tomizuka, "Track Deformable Objects from Point Clouds with Structure Preserved Registration," Int. J. Robot. Res., vol. 41, no. 6, pp. 599–614, 2022.
- [18] Y. Yang, J. A. Stork, and T. Stoyanov, "Learning to Propagate Interaction Effects for Modeling Deformable Linear Objects

Dynamics," IEEE/RSJ Int. Conf. Intell. Robot. Sys. (IROS), pp. 4056–4062, 2021.

- [19] W. Zhang, K. Schmeckpeper, P. Chaudhari, and K. Daniilidis, "Deformable Linear Object Prediction Using Locally Linear Latent Dynamics," in *IEEE Int. Conf. Robot. Autom. (ICRA)*, June 2021, pp. 13 503–13 509.
- [20] J. Schulman, A. Lee, J. Ho, and P. Abbeel, "Tracking Deformable Objects with Point Clouds," in *IEEE Int. Conf. Robot. Autom. (ICRA)*, 2013, pp. 1130–1137. 1
- [21] S. Ge, G. Fan, and M. Ding, "Non-Rigid Point Set Registration with Global-Local Topology Preservation," *IEEE/CVF Int. Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, pp. 245– 251, 2014. 1
- [22] J. Xiang and H. Dinkel, "Simultaneous Shape Tracking of Multiple Deformable Linear Objects with Global-Local Topology Preservation," in *IEEE Int. Conf. Robot. Autom. (ICRA) Work*shop on Representing and Manipulating Deformable Objects, May 2023. 1
- [23] Z. Xu, J. Wu, A. Zeng, J. B. Tenenbaum, and S. Song, "Dense-PhysNet: Learning Dense Physical Object Representations via Multi-Step Dynamic Interactions," in *Robot. Sci. Syst. (RSS)*, 2019. 1
- [24] B. Bianchini, M. Halm, and M. Posa, "Simultaneous Learning of Contact and Continuous Dynamics," in *Conf. Robot. Learn.* (*CoRL*), 2023. 1
- [25] B. Lloyd, G. Székely, and M. Harders, "Identification of Spring Parameters for Deformable Object Simulation," *IEEE Trans. Vis. Comput. Graphics*, vol. 13, no. 5, pp. 1081–1094, 2007.
- [26] M. Macklin, M. Müller, N. Chentanez, and T.-Y. Kim, "Unified Particle Physics for Real-Time Applications," ACM Trans. Graph. (TOG), vol. 33, no. 4, Jul. 2014.
- [27] J. Bender, M. Müller, and M. Macklin, "Position-Based Simulation Methods in Computer Graphics." in *Eurographics (Tutorials)*, 2015, pp. 1–32. 1
- [28] T. Tang, Y. Fan, H.-C. Lin, and M. Tomizuka, "State Estimation for Deformable Objects by Point Registration and Dynamic Simulation," in *IEEE/RSJ Int. Conf. Intell. Robot. Sys. (IROS)*, 2017, pp. 2427–2433. 1
- [29] J. Abou-Chakra, K. Rana, F. Dayoub, and N. Suenderhauf, "Physically Embodied Gaussian Splatting: A Visually Learnt and Physically Grounded 3D Representation for Robotics," in *Conf. Robot. Learn. (CoRL)*, 2024. 3
- [30] M. Zhang, K. Zhang, and Y. Li, "Dynamic 3D Gaussian Tracking for Graph-Based Neural Dynamics Modeling," in *Conf. Robot. Learn. (CoRL)*, 2024. 1
- [31] N. Karaev, I. Rocco, B. Graham, N. Neverova, A. Vedaldi, and C. Rupprecht, "CoTracker: It is Better to Track Together," *Eur. Conf. Comput. Vis. (ECCV)*, 2024. 1
- [32] X. Shi, Z. Huang, W. Bian, D. Li, M. Zhang, K. C. Cheung, S. See, H. Qin, J. Dai, and H. Li, "VideoFlow: Exploiting Temporal Cues for Multi-frame Optical Flow Estimation," in *IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, 2023, pp. 12435– 12446.
- [33] C. Doersch, A. Gupta, L. Markeeva, A. Recasens, L. Smaira, Y. Aytar, J. Carreira, A. Zisserman, and Y. Yang, "TAP-Vid: A Benchmark for Tracking Any Point in a Video," in Adv. Neur. Inf. Proc. (NeurIPS), 2023.
- [34] Q. Wang, Y.-Y. Chang, R. Cai, Z. Li, B. Hariharan, A. Holynski, and N. Snavely, "Tracking Everything Everywhere All at Once," in *IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, 2023, pp. 19738– 19749.
- [35] Z. Teed and J. Deng, "RAFT: Recurrent All-Pairs Field Transforms for Optical Flow," in *Eur. Conf. Comput. Vis. (ECCV)*, 2020, pp. 402–419. 1
- [36] A. Caporali, K. Galassi, and G. Palli, "Deformable Linear Objects 3D Shape Estimation and Tracking From Multiple 2D

Views," IEEE Robot. Autom. Lett., vol. 8, no. 6, pp. 3852–3859, 2023. 1

- [37] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis," *Commun. ACM*, vol. 65, no. 1, p. 99–106, 2022. 1
- [38] A. Pumarola, E. Corona, G. Pons-Moll, and F. Moreno-Noguer, "D-NeRF: Neural Radiance Fields for Dynamic Scenes," in *IEEE/CVF Int. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2021, pp. 10318–10327.
- [39] K. Park, U. Sinha, J. T. Barron, S. Bouaziz, D. B. Goldman, S. M. Seitz, and R. Martin-Brualla, "Nerfies: Deformable Neural Radiance Fields," in *IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, 2021, pp. 5845–5854.
- [40] Z. Li, S. Niklaus, N. Snavely, and O. Wang, "Neural Scene Flow Fields for Space-Time View Synthesis of Dynamic Scenes," in *IEEE/CVF Int. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2021, pp. 6494–6504.
- [41] Z. Li, Q. Wang, F. Cole, R. Tucker, and N. Snavely, "DynIBaR: Neural Dynamic Image-Based Rendering," in *IEEE/CVF Int. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2023, pp. 4273– 4284. 1
- [42] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, "3D Gaussian Splatting for Real-Time Radiance Field Rendering," *ACM Trans. Graph. (TOG)*, vol. 42, no. 4, July 2023. 1, 2
- [43] J. Luiten, G. Kopanas, B. Leibe, and D. Ramanan, "Dynamic 3D Gaussians: Tracking by Persistent Dynamic View Synthesis," in *IEEE Int. Conf. 3D Vis. (3DV)*, 2024, pp. 800–809.
- [44] G. Wu, T. Yi, J. Fang, L. Xie, X. Zhang, W. Wei, W. Liu, Q. Tian, and X. Wang, "4D Gaussian Splatting for Real-Time Dynamic Scene Rendering," in *IEEE/CVF Int. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2024, pp. 20310–20320. 1
- [45] H. Gao, R. Li, S. Tulsiani, B. Russell, and A. Kanazawa, "Monocular Dynamic View Synthesis: A Reality Check," in Adv. Neur. Inf. Proc. (NeurIPS), 2022, pp. 33768–33780.
- [46] M. Büsching, J. Bengtson, D. Nilsson, and M. Björkman, "FlowIBR: Leveraging Pre-Training for Efficient Neural Image-Based Rendering of Dynamic Scenes," in *IEEE/CVF Int. Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, June 2024, pp. 8016–8026. 1
- [47] V. Ye, R. Li, J. Kerr, M. Turkulainen, B. Yi, Z. Pan, O. Seiskari, J. Ye, J. Hu, M. Tancik, and A. Kanazawa, "gsplat: An Open-Source Library for Gaussian Splatting," arXiv preprint arXiv:2409.06765, 2024. 3
- [48] H. Dinkel, R. Navaratna, J. Xiang, B. Coltin, T. Smith, and T. Bretl, "KnotDLO: Toward Interpretable Knot Tying," in *IEEE Int. Conf. on Robot. and Autom. (ICRA) 3D Visual Representations for Manipulation Workshop*, May 2024. 3
- [49] W. Peng, J. Lv, Y. Zeng, H. Chen, S. Zhao, J. Sun, C. Lu, and L. Shao, "TieBot: Learning to Knot a Tie from Visual Demonstration through a Real-to-Sim-to-Real Approach," in *Conf. Robot. Learn. (CoRL)*, 2024. 3
- [50] K. Koide and E. Menegatti, "General Hand-Eye Calibration Based on Reprojection Error Minimization," *IEEE Robot. Autom. Lett.*, vol. 4, no. 2, pp. 1021–1028, 2019. 3
- [51] Stanford Artificial Intelligence Laboratory, "Robotic Operating System: Noetic Ninjemys," https://www.ros.org, 2018. 3
- [52] M. Macklin, "Warp: A High-performance Python Framework for GPU Simulation and Graphics," March 2022, nVIDIA GPU Technology Conference (GTC). 3
- [53] S. Zhaole, H. Zhou, L. Nanbo, L. Chen, J. Zhu, and R. B. Fisher, "A Robust Deformable Linear Object Perception Pipeline in 3D: From Segmentation to Reconstruction," *IEEE Robot. Autom. Lett.*, vol. 9, no. 1, pp. 843–850, 2024. 3
- [54] 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," in *IEEE Int. Conf. Robot. Autom. (ICRA)*

Workshop on Representing and Manipulating Deformable Objects, May 2022.

- [55] R. Zanella, A. Caporali, K. Tadaka, D. De Gregorio, and G. Palli, "Auto-Generated Wires Dataset for Semantic Segmentation with Domain Independence," in *IEEE Int. Conf. Comput. Cont. Robot. (ICCCR).* IEEE, Jan. 2021, pp. 292–298.
- [56] A. Caporali, R. Zanella, D. De Gregorio, and G. Palli, "Ariadne+: Deep Learning-Based Augmented Framework for the Instance Segmentation of Wires," in *IEEE Trans. Ind. Inf.*, February 2022, pp. 1–11.
- [57] A. Caporali, K. Galassi, R. Zanella, and G. Palli, "FASTDLO: Fast Deformable Linear Objects Instance Segmentation," *IEEE Robot. Autom. Lett.*, vol. 7, no. 4, pp. 9075–9082, 2022.
- [58] A. Caporali, K. Galassi, B. L. Žagar, R. Zanella, G. Palli, and A. C. Knoll, "RT-DLO: Real-Time Deformable Linear Objects Instance Segmentation," *IEEE Trans. Ind. Informat.*, vol. 19, no. 11, pp. 11333–11342, 2023.
- [59] V. Viswanath, K. Shivakumar, M. Parulekar, J. Ajmera, J. Kerr, J. Ichnowski, R. Cheng, T. Kollar, and K. Goldberg, "HAND-LOOM: Learned Tracing of One-Dimensional Objects for Inspection and Manipulation," in *Conf. Robot. Learn. (CoRL)*, 2023, pp. 341–357. 3