

Pseudo-rigid body networks: learning interpretable deformable object dynamics from partial observations

Shamil Mamedov^{1,*}, A. René Geist^{2,*}, Jan Swevers¹, Sebastian Trimpe²

Abstract—Accurately predicting deformable linear object (DLO) dynamics is challenging, especially when the task requires a model that is both human-interpretable and computationally efficient. In this work, we draw inspiration from the pseudo-rigid body method (PRB) and model a DLO as a serial chain of rigid bodies whose internal state is unrolled through time by a dynamics network. This dynamics network is trained jointly with a physics-informed encoder that maps observed motion variables to the DLO’s hidden state. To encourage the state to acquire a physically meaningful representation, we leverage the forward kinematics of the PRB model as a decoder. We demonstrate in robot experiments that the proposed DLO dynamics model provides physically interpretable predictions from partial observations while being on par with black-box models regarding prediction accuracy. The project code is available at: tinyurl.com/prb-networks

I. INTRODUCTION

Deformable linear objects (DLOs), such as cables, ropes, and threads, are prevalent in various promising applications within the field of robotics [1, 2]. Dynamics models empower robots to interact with DLOs, showcasing remarkable dexterity at tasks such as intricate knotting [3, 4], precise manipulation of ropes [5], and surgical suturing [6, 7]. In addition, the parts of many robots can be modeled as DLOs. This naturally applies to numerous soft robots but is also relevant for controlling industrial robots [8].

Learning robot dynamics is challenging because the amount of real data is often limited, and the data’s coverage of the robot’s state space depends on the control strategy used for data collection. Therefore, purely data-driven models, such as neural networks (NNs), often fall short of expectations when learning robot dynamics [1, 2]. Furthermore, these models lack interpretability. A promising approach to enhancing both sample efficiency and interpretability is to incorporate first-principles knowledge into the model’s architecture [9, 10]. Physics-informed models are generally evaluated on rigid robots whose states are directly observed [11, 12, 13]. Unlike rigid robots, a DLO is a continuous accumulation of mass, making it unclear what the state representation should be and how to estimate it in the first place. Fortunately, we can draw inspiration from the literature on the pseudo-rigid body (PRB) method and the articulated body algorithm to pursue the simple idea:

* Equal contribution.

This work was largely done during an exchange stay of S. Mamedov at RWTH Aachen University, which was supported by the FWO-Vlaanderen through SBO project ELYSA for cobot applications (S001821N).

¹MECO Research Team, KU Leuven and Flanders Make@KU Leuven.

²Institute for Data Science in Mechanical Engineering, RWTH Aachen University.




Fig. 1: A DLO – an aluminium rod – is actuated by a robot arm. Given an initial observation of $\{e\}$ ’s translational state y_k and $\{b\}$ ’s position, velocity and acceleration vector x_k , a PRB-Net predicts y_{k+1} while also estimating the DLO’s hidden state h_k .

Model a DLO as a chain of rigid bodies whose dynamics are learned from data.

This approach combines the strengths of the PRB method and machine learning. First, it provides a physically interpretable state. Second, it infers the interactions between the bodies from data instead of solely relying on the expressiveness of analytical physics models.

A. Problem Statement and Contributions

Figure 1 gives an overview of the problem setup. The inertial coordinate frame reads $\{B\}$ while $\{b\}$ and $\{e\}$ denote frames fixed to the DLO’s start and end, respectively. The DLO’s hidden state is $h \in \mathbb{R}^{n_h}$. A dataset \mathcal{D} contains trajectories consisting of the position and velocity of $\{e\}$ relative to $\{B\}$, denoted by $y = [p_e, \dot{p}_e] \in \mathbb{R}^{n_y}$ and position, velocity, and acceleration of $\{b\}$ relative to $\{B\}$, denoted as $x = [q_b, \dot{q}_b, \ddot{q}_b] \in \mathbb{R}^{n_x}$. In short, we tackle the problem: *given measurements of position, velocity, and acceleration of a DLO’s end, predict position and velocity of its other end over the next N time steps* or, using our notation:

given y_k and x_k, \dots, x_{k+N-1} predict y_{k+1}, \dots, y_{k+N} ,

where k is the discrete time index (or step). To this end, we propose the PRB network (PRB-Net) which consists of four essential components: 1) a physics-informed encoder that maps y_k and x_k to h_k , 2) a dynamics network that predicts h_{k+1} from h_k and x_k , 3) a decoder that maps h_k and x_k to y_{k+1} , and 4) a loss function that regularizes h during training. The contributions of this work are:

- 1) We introduce PRB-Net, a model designed to predict the dynamics of DLOs from partial observations.
- 2) We demonstrate that using the forward kinematics of a DLO’s PRB discretization as a decoder effectively enforces a physically consistent hidden state.
- 3) Through experiments on two different DLOs, we show that PRB-Net’s predictions are on par with black-box models in terms of accuracy and computation times, while being physically interpretable.

II. RELATED WORK

The main contribution of our work revolves around learning hidden state dynamics of a DLO from partial observations using PRBs. Therefore, special emphasis is placed on the type of observations that are used by the different DLO dynamics models.

a) Analytical physics-based models: Numerous works explore physics-based dynamics modeling of DLOs [14, 15, 16]. An in-depth discussion of the mechanics underlying DLO is given in [17]. The dynamics of DLOs can be conceptualized as an aggregation of infinitesimally small mass particles interacting through forces. To render the complicated particle interactions tractable for computational analysis, spatial discretization methods aggregate the continuous distribution of particles into finite elements. The PRB method assumes that the particles inside an element form a rigid body such that their distances to each other remain fixed [18]. The PRB approximation results in a simplified representation of the DLO’s dynamics, but its accuracy depends on the chosen force function that determines element interactions and its parameters. In particular, the oftenly used linear spring-damper force models are inaccurate for coarse spatial discretizations or modeling DLOs with heterogeneous material properties.

Cosserat rods represents an alternative method for modeling DLOs, particularly favored by the soft robotics community. Although these models are often quite accurate and geometrically exact, their large prediction times due to numerical integration renders them impractical for model-based control [19]. To mitigate the computational demands of Cosserat rod models, finite difference [17, 20] and piecewise constant curvature [21] approximations have been proposed. Stella et al. employed the latter to model the dynamics of a 2D underwater tentacle [22] and the kinematics of a 3D soft-continuum robot [23].

This paper adopts the state-representation and kinematics from the analytical methods, specifically from PRB method. Instead of depending on analytical dynamics, which require numerical integration, struggle to capture certain real-world phenomena, and are computationally unsuitable for control, we learn its discrete-time approximation from data.

b) Data-driven models: Table I presents a summary of related data-driven DLO dynamics models. Our research is inspired by [27], which employed a long short-term memory (LSTM) network [28] to capture the input-output dynamics of a robot-arm-held foam cylinder. In this study, the authors utilize the LSTM to predict the DLO’s end positions based on the pitch and yaw angles of its other end. The LSTM’s hidden state lacks a physically meaningful representation hindering the formulation of constraints on the DLO’s shape, particularly during manipulation in an environment with obstacles. Additionally, prediction errors of a black-box model can quickly accumulate causing the trajectory prediction to become unstable. PRB-Net alleviates these limitations by resorting to kinematics.

In related works, [24, 25] employed camera-based techniques to estimate the DLO’s center-line as a sequence of

	System	Dynamics model	State type	Observation
[22]	Submerged tentacle (2D)	constant curvature	Splines	Full state
[24]	Submerged tentacle (2D)	echo state network	Points	Full state
[25]	DLO in magnetic field	LSTMs	Points	Full state
[5]	Rope moved by robot arm	bi-LSTM	Mass particles	Full state
[26]	DLO attached to robot arm	bi-LSTM + Graph neural network	Mass particles	Full state
[27]	DLO attached to robot arm	LSTM	Black box	DLO ends
Ours	DLO attached to robot arm	flexible (e.g., <i>RNN, ResNet</i>)	Rigid bodies	DLO ends

TABLE I: Overview of related literature on data-driven modeling of DLO dynamics. While numerous works focus on the dynamics identification and observe the DLO’s full state through cameras, [27] and our work learn the dynamics by solely observing frames attached to the DLO’s start and end. Moreover, while previous literature focuses on a particular model architecture to learn dynamics, we focus on leveraging kinematics such that *any* dynamics model learns a physically meaningful representation.

points. In [25], the dynamics of each point are predicted by an LSTM that shares with the other LSTMs its output history. Similarly, [24] introduced echo state networks, where the output of an RNN serves as the input for a second RNN, in addition to control inputs and bending sensor measurements. In [5], the authors developed a quasi-static model of a DLO using a bi-directional LSTM (bi-LSTM) whose state is estimated from camera images using a convolutional neural network. Unlike previous works that used LSTMs to propagate states in time, the bi-LSTM propagates information along both directions of the DLO’s spatial dimension. Inspired by [5], [26] combined bi-LSTMs with graph neural networks [29]. In particular, [26] discretizes the DLO into a sequence of cylindrical elements similar to Cosserat theory [6] while the spatial and temporal element interactions are modeled through the NNs. However, these seminal works require cameras to obtain an estimation of several points on the DLO and operate at low velocities where quasi-static models are sufficient for predicting states. In contrast, the proposed PRB-Net allows for extending these methods to operate on partial observations and in highly dynamic scenarios. In particular, the network in PRB-Net could be replaced by bi-LSTMs and graph neural networks.

III. BACKGROUND: PRB METHOD FOR MODELING DLOS

The PRB method approximates a DLO as a chain of $n_{el} + 1$ rigid bodies connected via elastic joints. PRB-Net explicitly uses forward kinematics of a DLO’s PRB discretization and uses the same input-output structure as found in the dynamics of rigid-body chains. The following subsections briefly describe the kinematics and dynamics of a DLO using the PRB approximation.




Fig. 2: Generalized coordinate representations of a body chain with $n_{\text{el}} = 2$. The chain's joint states are described through rotation parameters Ψ_i that provide two rotational DOF to each element.

A. Kinematics

As illustrated in Fig. 2, the pose of the i -th body in the chain is described via a body-fixed coordinate frame $\{i\}$. The pose of $\{e\}$ with respect to $\{B\}$ is described by the homogeneous transformation matrix

$${}^B T_e = \begin{bmatrix} {}^B R_e & {}^B p_{e,B} \\ 0 & 1 \end{bmatrix} \in \text{SE}(3)$$

with the position vector ${}^B p_{e,B} \in \mathbb{R}^3$, rotation matrix ${}^B R_e$, and $\text{SE}(3)$ denoting the special Euclidean group [30]. Note that we omit the superscripts and subscripts if a vector or matrix is either expressed in the inertia frame $\{B\}$ or it points from $\{B\}$'s origin such that $p_e := {}^B p_{e,B}$ and $R_e := {}^B R_e$. In a chain of $n_{\text{el}} + 1$ bodies ${}^b p_{e,b}$ is obtained as

$$\begin{bmatrix} {}^b p_{e,b} \\ 1 \end{bmatrix} = {}^b T_1(q_1) \cdot {}^1 T_2(q_2) \cdot \dots \cdot {}^{n_{\text{el}}-1} T_{n_{\text{el}}}(q_{n_{\text{el}}}) \begin{bmatrix} {}^{n_{\text{el}}} p_{e,n_{\text{el}}} \\ 1 \end{bmatrix},$$

where $q_{\text{prb}} = [q_1, \dots, q_{n_{\text{el}}}]$ denotes the generalized coordinates of the DLO's rigid body chain approximation and ${}^{i-1} T_i(q_i)$ transforms a (homogeneous) vector from $\{i\}$ to $\{i-1\}$ [30, p. 93]. In addition, $q_b = [p_b, \Psi_b]$ defines the six degrees of freedoms (DOF) of the floating-base $\{b\}$.

For simplicity's sake, we employ two DOF revolute joints as illustrated in Fig. 2 that are parameterized by Euler angles. Note that any other joint parametrization could be used such as free joints that provide six DOFs. In particular, three dimensional rotational joints should be used if the DLO is subject to torsion. In this case, recent works on learning with 3D rotations advocate against using Euler angles for large changes in rotation [31, 32].

Using the forward kinematics of the PRB approximation, the DLO's endpoint position is obtained as follows

$$p_e = \text{fk}(q_b, q_{\text{prb}}) := R_b(\Psi_b) {}^b p_{e,b}(q_{\text{prb}}) + p_b. \quad (1)$$

The velocity the DLO's end is obtained from differential kinematics which describes the relation between \dot{q}_b , \dot{q}_{prb} , and \dot{p}_e via the joint Jacobian matrix $J_{\text{lin}}(q_b, q_{\text{prb}})$ as

$$\dot{p}_e = J_{\text{lin}}(q_b, q_{\text{prb}}) \begin{bmatrix} \dot{q}_b \\ \dot{q}_{\text{prb}} \end{bmatrix} = \begin{bmatrix} \frac{\partial \text{fk}}{\partial q_b} & \frac{\partial \text{fk}}{\partial q_{\text{prb}}} \end{bmatrix} \begin{bmatrix} \dot{q}_b \\ \dot{q}_{\text{prb}} \end{bmatrix}. \quad (2)$$

Equation (1) and (2) together form the system's *forward kinematics* (FK) that maps elements in h_k and x_k to y_k .

B. Dynamics

The dynamics of a pseudo-rigid body chain can be categorized into forward, inverse, or hybrid dynamics. Forward dynamics uses joint torques as inputs to determine joint accelerations, while inverse dynamics calculates joint torques

from known joint accelerations. The dynamics of a DLO being manipulated by a robot arm is governed by *hybrid dynamics* being the combination of forward and inverse dynamics [33]. For the first body attached to $\{b\}$, the acceleration \ddot{q}_b is given and τ_b is unobserved. For the other bodies, \ddot{q}_{prb} is a function of $\tau_{\text{prb}}(q_{\text{prb}}, \dot{q}_{\text{prb}})$. In turn, the hybrid dynamics of a body chain takes the form

$$\ddot{q}_{\text{prb}} = f_{\text{HybDyn}}(h, x, \theta_P), \quad (3)$$

where θ_P denotes physical parameters *e.g.*, element length, mass, and inertia. A prediction of the next state is obtained by rewriting (3) in continuous state-space form $\dot{h} = f(h, x, \theta_P)$ and resorting to numerical integration, reading

$$h_{k+1} = \text{ODESolve}(f, h_k, x_k, \Delta t), \quad (4)$$

where $\Delta t = t_{k+1} - t_k$ is the integration step size and $\text{ODESolve}(\cdot)$ refers to a numerical integration method.

Two main weaknesses of PRB and other analytical models are their limited expressiveness and slow prediction times. Expressiveness of a PRB is limited by the choice of interaction force model. Commonly used interaction models such as linear spring-dampers require fine spatial discretization to yield tolerable approximations, which increases prediction times [18]. Computation is further slowed down as PRB dynamics are often stiff, which necessitates the use of implicit integration. PRB-Net improves prediction speed by modeling DLO dynamics in *discrete-time* with neural networks.

IV. METHOD

PRB-Net consists of three essential components as shown in the Fig. 3: a physics-informed encoder that maps initial observation y_k and input x_k into a hidden state h_k ; a discrete-time dynamics network that propagates hidden states h_k in time using inputs x_k, \dots, x_{k+N} ; and an analytical decoder that maps hidden states and inputs into observations y_k, \dots, y_{k+N} . This decoder equals the forward kinematics of a DLO's PRB approximation.

A. Decoder

The decoder in the PRB-Net is inspired by the kinematics description of a serial chain of rigid bodies. The decoder equals the FK as given by (1) and (2). By using such a decoder, we ensure the interpretability of a neural network that models the DLO's dynamics, while maintaining a physically meaningful relationship between the hidden states.

To derive an FK decoder, we specify the type of joint (*e.g.*, YZ Euler angles), provide the DLO's overall length, select the number of elements n_{el} , and use [18] to obtain the initial element lengths.

The uniform discretization method [18] assumes constant element lengths L_{el} . As illustrated in Fig. 4 (Left), fixing the first body to $\{b\}$ causes an error ϵ_{el} in the PRB-Net's prediction of the DLO end's position p_e . An upper bound for ϵ_{el} is obtained using the law of cosines, writing

$$\epsilon_{\text{el}} < \sqrt{L^2 + L_1^2 - 2LL_1 \cos(\zeta)} - L_2, \quad (5)$$

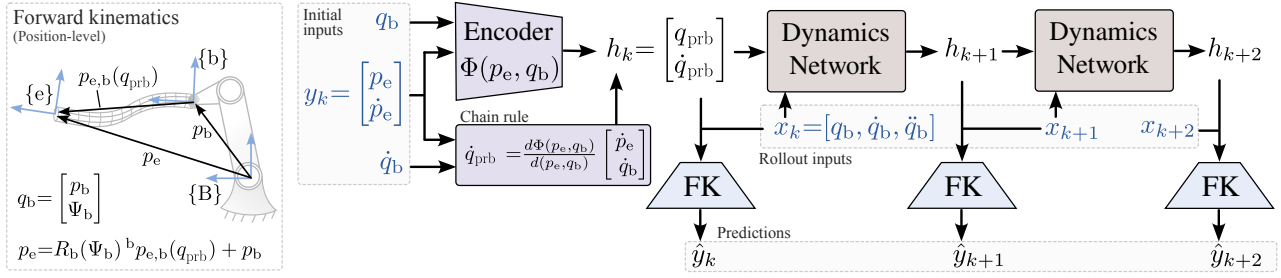


Fig. 3: Architecture of PRB-Net with rollout length $N = 2$. By resorting to the forward kinematics (FK) of a chain of PRBs and taking inspiration from hybrid dynamics to determine the inputs, the model learns a physically-plausible hidden state from partial observations.




Fig. 4: Errors arising in a rigid body chain approximation of a DLO with $n_{el} = 2$ and fixed element length. Left: If the chain's first body is fixed to $\{b\}$ then errors ϵ_{el} arise in the prediction of the DLO's end position p_e . Right: Upper bound on ϵ_{el} for $L = 1920$ mm caused by fixing the first body to $\{b\}$ and setting its length to $L_{el}/2$ plotted over different element counts n_{el} and deflection angles ζ .

with the DLO's length $L = L_1 + L_2$, the length of the first body L_1 , the length of all other elements L_2 , and the deflection angle ζ as depicted in Fig. 4.

We reduce ϵ_{el} by learning the body lengths $\theta_{el} \in \mathbb{R}_{>0}^{n_{el}}$ along the PRB-Net's network parameters. During the training of a PRB-Net, we regularize the sum of estimated body lengths as detailed in Section IV-D.

B. Encoder

The encoder predicts the inverse kinematics of the serial chain. Internally, the encoder performs two key transformations when the data is fed to a multilayer perceptron (MLP) $q_{prb} = \Phi(p_e, q_b)$. The transformations include: *i*) compute the vector $p_{e,b} = p_e - p_b$ to leverage that q_{prb} is independent to translations of the DLO in space; *ii*) transform the Euler angles Φ_b to $\sin(\cdot)$ and $\cos(\cdot)$ to avoid discontinuities in the angle parametrization. The velocities \dot{q}_{prb} are obtained via partial differentiation as

$$\dot{q}_{prb} = \frac{\partial \Phi}{\partial p_e} \dot{p}_e + \frac{\partial \Phi}{\partial q_b} \dot{q}_b, \quad (6)$$

which is straightforward to implement with automatic differentiation libraries such as JAX [34]. Note, (6) ensures that the velocity \dot{q}_{prb} is the time derivative of q_{prb} , which is not the case if an MLP directly predicts \dot{q}_{prb} . If $\dim(q_{prb}) > \dim([p_e, q_b])$, multiple chain configurations q_{prb} yield the same pose of $\{e\}$. In this case, learning inverse kinematics from observations is an ill-posed problem for deterministic networks [35].

C. Dynamics Network

The dynamics network unrolls the hidden state $h = [q_{prb}, \dot{q}_{prb}]$ – being the generalized coordinates of the




Fig. 5: Aluminium rod DLO shape predicted by a PRB-Net with $n_{el} = 7$. Top: $\alpha_q = 0$. Bottom: $\alpha_q = 1$.

pseudo-rigid bodies – over time. In turn, the dynamics network takes up a similar role as the analytical hybrid dynamics (3) and (4). Assuming that PRB is a suitable model for DLOs, these equations indicate that the only input needed for unrolling h in time is $x = [q_b, \dot{q}_b, \ddot{q}_b]$. Therefore, the dynamics network F is a neural network that receives as inputs x_k and h_k to predict h_{k+1} : $h_{k+1} = F(h_k, x_k)$. The network F can essentially have *any* architecture; in this paper, we consider recurrent and residual NNs.

While in our experiments, we attain precise estimates of \ddot{q}_b , future versions of PRB-Net could draw inspiration from [36] and replace \ddot{q}_b with past velocity observations. In addition, if one observes the robot's end-effector force τ_b , then the dynamics network could learn *forward dynamics* by replacing x with the input $x' = [q_b, \dot{q}_b, \tau_b]$.

D. Loss Regularization

While we hypothesized that simultaneously training forward kinematics with dynamics may regularize q_{prb} , we observed that acquiring a physically plausible state representation in PRB-Net requires loss regularization. Therefore, we introduce a regularization term in the loss used to train PRB-Net, writing

$$\mathcal{L} = \frac{1}{n_b N} \sum_{j=1}^{n_b} \sum_{k=1}^N w_k \|y_{k,j} - \hat{y}_{k,j}\|_{W_y}^2 + \mathcal{L}_{reg} \quad (7)$$

with the regularization loss

$$\mathcal{L}_{reg} = \underbrace{\alpha_q \|q_{prb,k,j}\|_2^2}_{\mathcal{L}_P} + \underbrace{\alpha_{\dot{q}} \|\dot{q}_{prb,k,j}\|_2^2}_{\mathcal{L}_K} + \mathcal{L}_{FK}, \quad (8)$$

the batch size n_b , the model's predictions $\hat{y}_{k,j}$, the weighting matrix $W_y \in \mathbb{R}^{6 \times 6}$, the state weight w_k , regularization parameters α_q and $\alpha_{\dot{q}}$, and additional kinematics regularization losses \mathcal{L}_{FK} . While the first two terms in (8) can be simply seen as ℓ^2 regularization of the hidden state h ,




Fig. 6: Influence of the element count n_{el} on the PRB-Net’s prediction error of p_e (blue) and \hat{p}_e (red) for a 1 s prediction horizon.

we can look at these terms through the lens of the PRB method. In this regard, \mathcal{L}_P in (8) represents a potential energy of linear spring forces acting between the elements while \mathcal{L}_K represents the chain’s kinetic energy. While the kinetic energy of the rigid body chain depends on the inertia matrix $M(q_{prb})$, it is lower and upper bounded by positive real constants $\kappa_1 \leq M(q_{prb}) \leq \kappa_2$ [37]. We can reconstruct the DLO’s shape from the PRB-Net’s hidden state h using FK (1). As illustrated in Fig. 5, the chain’s shape depends on the choice of α_q . Finally, \mathcal{L}_{FK} reads

$$\mathcal{L}_{FK} = \alpha_L \left(\sum \theta_{el} - L \right) + \alpha_{el} \|\theta_{el} - \theta'_{el}\|_2^2 + \alpha_{cb} \|\theta_{cb} - \theta'_{cb}\|_2^2, \quad (9)$$

with θ'_{el} denoting the initial element lengths obtained as in [18], θ_{cb} denoting the position of the marker attached to the end of a DLO ${}^b p_{e,b}$, and θ'_{cb} denoting the same quantity obtained from prior calibration routines. The first two terms in (9) enable the PRB-Net to reduce ϵ_{el} by estimating the element lengths θ_{el} as shown in Fig. 4. To account for calibration errors in the calibration of ${}^b p_{e,b}$, we introduce θ_{cb} to the kinematics decoder model and also optimize its entries during training.

V. EXPERIMENTAL EVALUATION

To evaluate the performance of different PRB-Net, we learn to predict the motion of two physical DLOs with different properties that are rigidly connected to a Franka Panda robot arm. The two DLO’s are: i) an aluminum rod as depicted in Fig. 1, ii) a polyethylene foam cylinder as depicted in Fig. 10. Kinematics parameters of DLOs (length, inner and outer diameters) are shown in Table II. The project code and data are available at: tinyurl.com/prb-networks.

A. Data

We collected twelve different trajectories for the aluminum rod and eight trajectories for the foam cylinder. During the first half of each trajectory (about 15 seconds) the robot arm moved the DLO; during the second half, the robot arm rested while the DLO kept moving. For the aluminum rod, nine trajectories were used for training and validation, while for the foam cylinder, seven trajectories were collected. The trajectories were divided into rollouts of length $N = 250$ using a rolling window approach, corresponding to one second of motion. Then, each rollout was split into 85% for training and 15% for validation. The remaining trajectories of the DLO’s were used for testing and divided into rollouts using a sliding window.




Fig. 7: RMSE of PRBN- and black-box models for a rollout length of $N = 250$ corresponding to a 1 s prediction horizon.

DLO	L , m	d_{in} , m	d_{out} , m	ρ , kg/m ³	E , Pa	η , Ns/m
Aluminum rod	1.92	0.004	0.006	2710	$5.15 \cdot 10^{10}$	$1.2 \cdot 10^8$
Foam cylinder	1.90	0	0.06	105	$1.8 \cdot 10^6$	$3.5 \cdot 10^5$

TABLE II: Physical parameters of a PRB model approximating the DLO dynamics with Young’s modulus E and normal damping coefficient η . d_{in} and d_{out} are the inner and outer diameters, respectively. The PRB model with these parameters is used as a baseline for the experiments in Section V.

B. Model Settings

As the PRB-Net’s dynamics network, we used RNN with Gated Recurrent Units [38] and ResNets [39] – and denote the resulting models as PRBN-RNN and PRBN-ResNet, respectively. As “black-box” baseline, we combine the same models with fully connected MLPs acting as encoder and decoder denoting the resulting models as RNN and ResNet. Note that the MLP encoder predicts the full state h without enforcing $\dot{q}_{prb} = dq_{prb}/dt$. As a first-principles baseline, we employ the classical PRB method with the manually tuned parameters ρ , E and η provided in Table II. The PRB-Net models are trained using the loss function (7) while the baseline NNs use an ℓ^2 loss. All models were implemented in JAX [34] using the Equinox library [40]. Further details and hyperparameters are provided in the attached code repository.

C. Experimental Results

a) *Element Count*: Fig. 6 depicts the RMSE of the PRBN-RNN model’s predictions on the test data. Initially, increasing the number of elements substantially reduces the error. This observation aligns with the common knowledge that discretization-based methods, like PRB method, show improved model accuracy for increasing element counts. However, beyond seven elements, the marginal improvement in prediction accuracy does not justify the increased computational cost associated with model evaluation time, which is crucial for enabling efficient model-based control. Therefore, for the remainder of this paper, we adopt $n_{el} = 7$.

b) *Prediction Accuracy*: Fig. 7 shows the prediction accuracy of the considered models evaluated on test trajectories of the same length ($N = 250$ corresponding to a prediction horizon of 1 second) as the training rollouts. Data-driven models significantly outperform the first-principles model (PRB). While physical interpretability often comes at the cost of performance (see PRB results), this is not the case for the PRB-Nets that provide essentially the same performance as black-box models, while being physically interpretable. The models more accurately predict the aluminum rod’s motion




Fig. 8: RMSE of p_e for different prediction horizons.

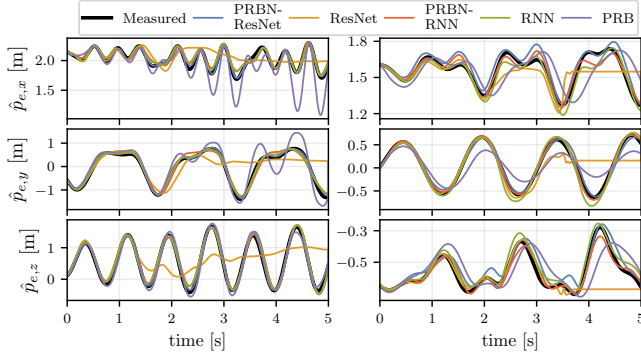


Fig. 9: Models rollouts with a prediction horizon five times longer than during training. Left: Aluminum rod. Right: Foam cylinder.

than the foam cylinder, possibly due to variations in training data or material properties.

All the models yield stable predictions during rollouts that extend significantly beyond N , as depicted in Fig. 8 and 9. However, the ResNet exhibits a noticeable drop in prediction accuracy as the prediction length exceeds N . In contrast, the PRBN-ResNet does not show this issue. Figure 8 shows the model’s RMSE of the prediction of p_e for both DLOs. Our findings indicate that the PRB-Net and models incorporating black-box encoders and decoders show similar performance. Among models with a black-box decoder, the RNN has a smaller RMSE than the ResNet for all rollouts except for the one of length N . Additionally, the PRB-Net exhibits slightly higher RMSEs compared to the RNN at $20N$. These observations align with literature on position encodings [41] that suggest that resorting to a low dimensional latent space whose dimensionality is determined by physics reduces the model’s expressiveness in favor of physical interpretability.

c) Training and Model Evaluation Times: In terms of training time, PRB-Net is slower (about 1.5 times) to train than standard neural networks as the nonlinear forward kinematics and the physics-informed encoder require gradient computations during the forward pass as shown in Fig. 3. Since the trained models may be used for model-based control, we analyzed the prediction times for one second-ahead prediction of the best data-driven models and of analytical PRB models. The PRBs analytical continuous-time dynamics were integrated using a constant step-size implicit 3rd order Runge-Kutta integrator.

Table III shows that the PRBN-RNN and RNN models have similar evaluation times, with the latter being slightly faster. However, even with just a two-element discretization, the evaluation time for the PRB method is at least two orders of magnitude slower. While integration times could be




Fig. 10: Left: Shape of a foam cylinder predicted by a PRB-Net with $n_{el} = 7$ and the element length being learned alongside other parameters. Right: Foam cylinder actuated by a robot arm.

	RNN	PRBN-RNN	PRB-2	PRB-5	PRB-7	PRB-10
Pred. time [s]	$4.3 \cdot 10^{-3}$	$5.7 \cdot 10^{-3}$	1.5	2.3	3.1	4.3

TABLE III: Model prediction times for a one second-ahead prediction. For the analytical PRB method several spatial discretizations were considered *e.g.*, $n_{el} = 2$ for “PRB-2”.

reduced by selecting a dedicated integrator, they will likely not match the discrete-time models like PRB-Net which requires only a single function evaluation.

d) Physical interpretability: PRB-Nets possess a physically meaningful state representation compared to black-box models such as RNNs and ResNets. In this regard, PRB-Nets hidden state represents elastic joint positions and velocities between rigid bodies. Therefore, one can approximately reconstruct the DLO’s shape via forward kinematics (1) from a partial observation y and an input x , as shown in Fig. 10.

VI. CONCLUSION

We presented the first data-driven DLO dynamics model capable of accurately learning the motion of a DLO from partial observations while also acquiring a physically interpretable hidden state representation. This model was obtained by first deriving the forward kinematics of a rigid body chain approximating the DLO. At the first step of a rollout, the DLO’s hidden state is reconstructed from partial observations of the DLO’s start and end frames using an MLP encoder. In the subsequent time steps, the hidden state is unrolled through time using a neural network that effectively learns the system’s dynamics. All networks are trained jointly with the kinematics parameters while we regularize the hidden state to ensure that it acquires a physically plausible hidden state representation.

Our model successfully predicted the position and velocity of a 1.92 m long aluminum rod and a 1.90 m long polyethylene foam cylinder. Despite the endpoints of the aluminium rod moving up to 1.5 m from their equilibrium position, the trained PRB-Net achieved a sub-decimeter RMSE in a twenty-second simulation.

ACKNOWLEDGMENT

The authors thank Sebastian Giedyk for helping with the setup, Henrik Hose for his inputs on ROS, and Wilm Decré for his valuable feedback on the manuscript.

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 International Journal of Robotics Research*, vol. 37, no. 7, pp. 688–716, 2018.
- [2] P. Jiménez, "Survey on model-based manipulation planning of deformable objects," *Robotics and computer-integrated manufacturing*, vol. 28, no. 2, pp. 154–163, 2012.
- [3] W. Wang, D. Berenson, and D. Balkcom, "An online method for tight-tolerance insertion tasks for string and rope," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 2488–2495.
- [4] E. Yoshida, K. Ayusawa, I. G. Ramirez-Alpizar, K. Harada, C. Duriez, and A. Kheddar, "Simulation-based optimal motion planning for deformable object," in *2015 IEEE international workshop on advanced robotics and its social impacts (ARSO)*. IEEE, 2015, pp. 1–6.
- [5] 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.
- [6] D. K. Pai, "Strands: Interactive simulation of thin solids using cosserat models," in *Computer graphics forum*, vol. 21. Wiley Online Library, 2002, pp. 347–352.
- [7] M. Saha and P. Isto, "Manipulation planning for deformable linear objects," *IEEE Transactions on Robotics*, vol. 23, no. 6, 2007.
- [8] A. Verl, A. Valente, S. Melkote, C. Brecher, E. Ozturk, and L. T. Tunc, "Robots in machining," *CIRP Annals*, vol. 68, no. 2, pp. 799–822, 2019.
- [9] M. Lutter, C. Ritter, and J. Peters, "Deep lagrangian networks: Using physics as model prior for deep learning," in *International Conference on Learning Representations*, 2019.
- [10] M. Cranmer, S. Greydanus, S. Hoyer, P. Battaglia, D. Spergel, and S. Ho, "Lagrangian neural networks," in *ICLR 2020 Workshop on Integration of Deep Neural Models and Differential Equations*, 2019.
- [11] L. Rath, A. R. Geist, and S. Trimpe, "Using physics knowledge for learning rigid-body forward dynamics with gaussian process force priors," in *Conference on Robot Learning*. PMLR, 2022, pp. 101–111.
- [12] D. Nguyen-Tuong and J. Peters, "Using model knowledge for learning inverse dynamics," in *2010 IEEE International Conference on Robotics and Automation*, 2010, pp. 2677–2682.
- [13] M. Lutter and J. Peters, "Combining physics and deep learning to learn continuous-time dynamics models," *The International Journal of Robotics Research*, vol. 42, no. 3, pp. 83–107, 2023.
- [14] 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*, 2007, pp. 63–72.
- [15] S. Drücker and R. Seifried, "Application of stable inversion to flexible manipulators modeled by the absolute nodal coordinate formulation," *GAMM-Mitteilungen*, vol. 46, no. 1, 2023.
- [16] J. Kim and N. S. Pollard, "Fast simulation of skeleton-driven deformable body characters," *ACM Transactions on Graphics (TOG)*, vol. 30, no. 5, pp. 1–19, 2011.
- [17] M. Bergou, M. Wardetzky, S. Robinson, B. Audoly, and E. Grinspun, "Discrete elastic rods," in *ACM SIGGRAPH 2008*, 2008, pp. 1–12.
- [18] E. Wittbrodt, I. Adamiec-Wójcik, and S. Wojciech, *Dynamics of flexible multibody systems: rigid finite element method*. Springer Science & Business Media, 2007.
- [19] F. Boyer, A. Gotelli, P. Tempel, V. Lebastard, F. Renda, and S. Briot, "Implicit time-integration simulation of robots with rigid bodies and cosserat rods based on a newton–euler recursive algorithm," *IEEE Transactions on Robotics*, vol. 40, pp. 677–696, 2024.
- [20] H. Lang, J. Linn, and M. Arnold, "Multi-body dynamics simulation of geometrically exact cosserat rods," *Multibody System Dynamics*, vol. 25, no. 3, pp. 285–312, 2011.
- [21] R. J. Webster III and B. A. Jones, "Design and kinematic modeling of constant curvature continuum robots: A review," *The International Journal of Robotics Research*, vol. 29, no. 13, pp. 1661–1683, 2010.
- [22] F. Stella, N. Obayashi, C. D. Santina, and J. Hughes, "An experimental validation of the polynomial curvature model: Identification and optimal control of a soft underwater tentacle," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11 410–11 417, 2022.
- [23] F. Stella, Q. Guan, C. Della Santina, and J. Hughes, "Piecewise affine curvature model: a reduced-order model for soft robot-environment interaction beyond pcc," in *2023 IEEE International Conference on Soft Robotics (RoboSoft)*, 2023, pp. 1–7.
- [24] K. Tanaka, Y. Minami, Y. Tokudome, K. Inoue, Y. Kuniyoshi, and K. Nakajima, "Continuum-body-pose estimation from partial sensor information using recurrent neural networks," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11 244–11 251, 2022.
- [25] A. Tariverdi, V. K. Venkiteswaran, M. Richter, O. J. Elle, J. Tørresen, K. Mathiassen, S. Misra, and Ø. G. Martinsen, "A recurrent neural-network-based real-time dynamic model for soft continuum manipulators," *Frontiers in Robotics and AI*, vol. 8, p. 631303, 2021.
- [26] Y. Yang, J. A. Stork, and T. Stoyanov, "Learning to propagate interaction effects for modeling deformable linear objects dynamics," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 1950–1957.
- [27] J. A. Preiss, D. Millard, T. Yao, and G. S. Sukhatme, "Tracking fast trajectories with a deformable object using a learned model," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 1351–1357.
- [28] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [29] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini, "The graph neural network model," *IEEE transactions on neural networks*, vol. 20, no. 1, pp. 61–80, 2008.
- [30] K. M. Lynch and F. C. Park, *Modern robotics*. Cambridge University Press, 2017.
- [31] D. Q. Huynh, "Metrics for 3d rotations: Comparison and analysis," *Journal of Mathematical Imaging and Vision*, vol. 35, pp. 155–164, 2009.
- [32] R. Bréquier, "Deep regression on manifolds: a 3d rotation case study," in *2021 International Conference on 3D Vision (3DV)*. IEEE, 2021, pp. 166–174.
- [33] J. Kim, "Lie group formulation of articulated rigid body dynamics," Technical report, Carnegie Mellon University, Tech. Rep., 2012.
- [34] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, "JAX: composable transformations of Python+NumPy programs," 2018. [Online]. Available: <http://github.com/google/jax>
- [35] M. I. Jordan, "Constrained supervised learning," *Journal of Mathematical Psychology*, vol. 36, no. 3, pp. 396–425, 1992.
- [36] A. Sanchez-Gonzalez, J. Godwin, T. Pfaff, R. Ying, J. Leskovec, and P. Battaglia, "Learning to simulate complex physics with graph networks," in *International conference on machine learning*. PMLR, 2020, pp. 8459–8468.
- [37] B. Siciliano, L. Sciavicco, L. Villani, and G. Oriolo, "Modelling, planning and control," *Advanced Textbooks in Control and Signal Processing*. Springer, 2009.
- [38] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using rnn encoder-decoder for statistical machine translation," *arXiv preprint arXiv:1406.1078*, 2014.
- [39] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [40] P. Kidger and C. Garcia, "Equinox: neural networks in JAX via callable PyTrees and filtered transformations," *Differentiable Programming workshop at Neural Information Processing Systems 2021*, 2021.
- [41] P. Duffer, M. Schmitt, and H. Schütze, "Position Information in Transformers: An Overview," *Computational Linguistics*, vol. 48, no. 3, pp. 733–763, 09 2022. [Online]. Available: <https://doi.org/10.1162/coli.a.00445>