Physics-Guided Learning for Tendon State Estimation of Cable-Driven Hyper-Redundant Robots

Jiaxun He, Jun-Guo Lu, *Member, IEEE*, Qing-Hao Zhang, *Member, IEEE*, Zhenpu Zhu, Zhanxuan Peng, and Guoying Gu, *Senior Member, IEEE*

***Abstract*—Accurate estimation of joint-side tendon tensions at every joint is essential for high-performance hybrid control of cable-driven hyper-redundant robots (CDHRRs), yet direct measurement is typically unavailable. Actuator-side tensions are measurable for all tendons, while joint-side measurements are available only at the terminal joint. Under this sensing constraint, analytical transmission models suffer from compounding friction errors, while data-driven learning becomes ill-posed without intermediate joint-side measurements. A physics-guided tension estimator, the Physics-Guided Cable-Hole Transmission Cascade (PG-CHTC), is proposed for instrumentation-limited CDHRRs to estimate joint-side tendon tensions along the chain from actuator-side measurements. A shared-parameter strategy is enforced by learning a positive, element-wise tension-ratio predictor reused at each guide, regularized by an Euler–Eytelwein capstan prior and corrected by a strictly causal temporal residual. This multiplicative cascade improves practical identifiability under terminal-only joint-side measurements, while inference relies solely on actuator-side sensing. The approach is validated on a 12-module CDHRR, where open-loop logged-data evaluation at the terminal joint yields a 70.2% reduction in tension RMSE compared to a tuned physical baseline. Deployment within a closed-loop controller yields settling time reductions of 44.4% to 50.8% across test protocols, improving tracking performance previously constrained by transmission uncertainty.**

***Index Terms*—Tendon/wire mechanism, redundant robots, model learning for control, deep learning methods, flexible robotics.**

Received 5 February 2026; accepted 12 June 2026. Date of publication 2 July 2026; date of current version 10 July 2026. This article was recommended for publication by Associate Editor D. Zanotto and Editor C. Gosselin upon evaluation of the reviewers' comments. This work was supported in part by the National Natural Science Foundation of China under Grant 62303309, Grant 62073217, and Grant62403310, in part by the National Science and Technology Major Project under Grant 2025ZD1606600, and in part by the National Science and Technology Major Project of New Oil and Gas Exploration and Development under Grant 2025ZD1401200 and Grant 2025ZD1401201. *(Corresponding author: Jun-Guo Lu.)*

Jiaxun He, Jun-Guo Lu, and Qing-Hao Zhang are with the School of Automation and Intelligent Sensing, Key Laboratory of System Control and Information Processing, Ministry of Education of China and Shanghai Key Laboratory of Perception and Control in Industrial Network Systems, Shanghai Jiao Tong University, Shanghai 200240, China (e-mail: jiaxun.he@sjtu.edu.cn; jglu@sjtu.edu.cn; zhangqinghao@sjtu.edu.cn).

Zhenpu Zhu, Zhanxuan Peng, and Guoying Gu are with the State Key Laboratory of Mechanical System and Vibration, Robotics Institute, School of Mechanical Engineering, Shanghai Jiao Tong University, Shanghai 200240, China (e-mail: zhuzhenpu@sjtu.edu.cn; steven2606@sjtu.edu.cn; guguoying@sjtu.edu.cn).

This article has supplementary downloadable material available at https://doi.org/10.1109/LRA.2026.3709642, provided by the authors.

Digital Object Identifier 10.1109/LRA.2026.3709642

I. Introduction

CABLE-DRIVEN hyper-redundant robots (CDHRRs) offer high dexterity, slender form factors, and the ability to navigate cluttered workspaces for inspection and manipulation tasks [1], [2]. However, the actuation mechanism, which comprises long, densely routed tendons, creates a transmission system strongly dependent on configuration and path history. This configuration-dependent transmission complicates sensing and control beyond conventional rigid manipulators. For effective closed-loop regulation, the joint-side tensions, which act directly at the joints, are the critical quantities determining joint torques and safety margins against slackness or over-tension. These values serve as essential inputs for modern hybrid tension and configuration control strategies, where joint effort arises from joint-side force differentials and moment arms via the standard Jacobian mapping [3], [4], [5], [6].

In practice, sensing is typically restricted to the actuator-side interface. As tendons traverse multiple chamfered guides, configuration-dependent contact, accumulated friction, and viscoelastic hysteresis introduce persistent mismatch between actuator-side measurements and joint-side quantities [7]. Direction-dependent friction and stick–slip behaviors across numerous guides further challenge purely analytical models [8], [9]. This work addresses an instrumentation-limited setting: actuator-side tensions are measurable for all tendons, but joint-side tensions are unobserved at intermediate joints and measurable only at the terminal joint. Terminal joint-side measurements provide the sole supervision signal for learning, while runtime estimation relies on actuator-side sensing alone. This setting is typical for lightweight or slender manipulators where embedding distributed force sensors is spatially prohibitive.

Existing transmission models fall into physics-based and data-driven categories. Physics-based pipelines accumulate calibration errors along the routing path, with residual error scaling unfavorably in robot length [10], [11]. Black-box learners remain ill-posed under terminal-only joint-side supervision, since infinite internal tension combinations can map to the same terminal output. Section II details both threads.

To bridge this gap, a physics-guided, weakly supervised framework is developed in this letter, specifically tailored to discrete-joint CDHRRs. Identical per-guide predictors, each regularized by a Capstan-inspired cosine-affine friction prior, are composed into a *multiplicative* cascade with parameters shared across all modules, alleviating the ill-posedness inherent in terminal-only supervision and enabling robust end-to-end training from terminal joint-side measurements alone.

This letter makes three contributions:

2377-3766 © 2026 IEEE. All rights reserved, including rights for text and data mining, and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

Authorized licensed use limited to: Shanghai Jiaotong University. Downloaded on July 13,2026 at 06:30:19 UTC from IEEE Xplore. Restrictions apply.

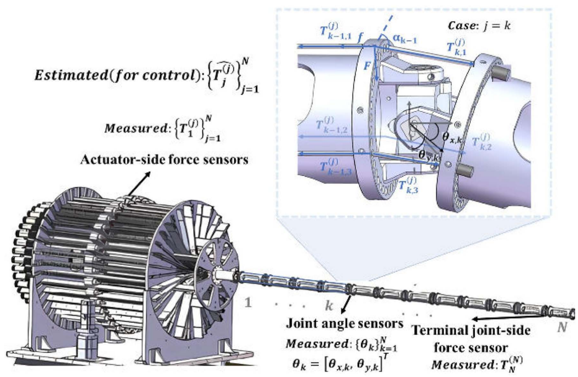

Fig. 1. Problem definition and sensing configuration of the instrumentation-limited CDHRR. The actuator-side tendon triplets $\{\mathbf{T}_1^{(j)}(t)\}_{j=1}^N$ and joint angles $\{\boldsymbol{\theta}_k(t)\}_{k=1}^N$ are measurable, while only the terminal joint-side tension $\mathbf{T}_N^{(N)}(t)$ is available for supervision/evaluation. The objective is to estimate joint-side tendon tensions at their own joints, $\{\hat{\mathbf{T}}_j^{(j)}(t)\}_{j=1}^N$.

1) A physics-guided cascaded learning framework for *discrete-joint* CDHRRs is proposed. By improving practical identifiability via shared-parameter, element-wise positive ratio predictors, the framework learns from *terminal-only joint-side measurements*, addressing the inability to instrument intermediate joints.
2) A physics-regularized training objective is introduced. Intermediate tension-ratio predictions are regularized with an Euler–Eytelwein capstan prior in the log domain and a spatial consistency term, while a causal temporal residual captures hysteretic effects, improving stability and data efficiency under terminal-only joint-side measurements.
3) The approach is deployed on a 12-module CDHRR. In open-loop logged-data evaluation at the terminal joint, joint-side tension RMSE is reduced by 70.2% relative to physics-only baselines [5], [12]; single-passage validation shows lower RMSE and higher R^2 than learning-only baselines [13]. In a closed-loop setting, settling time is reduced by 44.4%–50.8% relative to the Capstan baseline estimator [3].

The remainder of this letter is organized as follows. Section II reviews the existing literature on transmission modeling and physics-guided learning for cable-driven robots. Section III details the proposed methodology, mathematically formulating the PG-CHTC architecture and providing the identifiability analysis for the shared-parameter cascade. Section IV presents the experimental validation, encompassing open-loop logged-data estimation evaluation and ablation studies, followed by estimator-in-the-loop closed-loop control experiments. Finally, Section V concludes this letter and outlines directions for future research.

II. Related Work

Transmission modeling for discrete-joint CDHRRs has predominantly relied on analytical physics derived from serial tendon–guide routing. At the guide scale, capstan relations are employed to map tensions via wrap angles and friction coefficients. Extensions to this formulation have introduced bending rigidity and angle-dependent, nonsmooth friction to account for stick–slip phenomena [8], [12]. These local models are computationally lightweight and theoretically composable along the chain. At the global routing scale, statics and dynamics models offering actuator-to-terminal mappings have been developed, with recent multi-section variants embedding friction to reduce estimation bias [5], [10], [11]. Nonetheless, on long CDHRRs, identification becomes brittle due to the multiplicity of friction, contact, and routing parameters, alongside unmodeled viscoelastic effects. Local mismatches inevitably accumulate across guides, causing both identification complexity and residual error to scale unfavorably with path length [14].

To address the brittleness of analytical parameter identification, data-driven methods have been explored for unmodeled friction and compliance: compact feedforward regressors, recurrent models, hysteresis compensation, vision-based shape estimation, and attention-based time-series regression [13], [15], [16], [17], [18]. However, under terminal-only joint-side supervision, purely data-driven cascades are underdetermined: multiple latent tension trajectories can yield identical terminal labels, causing intermediate estimates to collapse without structural constraints.

Physics-informed neural networks (PINNs) [19], [20] embed governing equations into the training loss, with robotics applications spanning continuum-robot Cosserat-rod approximation [21], friction-inclusive rigid-body inverse dynamics [22], hybrid residual dynamics [23], [24], and physics-informed multi-agent reinforcement learning [25]. For tendon-driven systems, modern friction-compensation models [14], [17], [18] and capstan and friction generalizations beyond constant-friction assumptions [8], [12] relax the static capstan assumption. The framework developed here instead targets cascade identifiability under terminal-only joint-side supervision by injecting a domain-specific capstan prior as a soft regularizer within a shared-parameter multi-guide cascade.

III. Methodology

The core idea is to decompose the ill-posed terminal-only estimation problem into a cascade of identical, physics-regularized single-passage units. By sharing parameters across all units and injecting a capstan-based prior, the degrees of freedom are reduced and practical identifiability is improved.

A. Problem Statement and Preliminaries

A class of discrete-joint, N-module CDHRRs is considered. Each module possesses two degrees of freedom (DoF) and is actuated by three tendons. Under a serial pass-through routing architecture, tendons actuating joint j pass through the co-located guide disks of all preceding joints $k < j$. This serial structure creates a compounding transmission dependency. The problem is formulated under a specific sensing configuration: actuator-side tensions are measurable, whereas joint-side tensions remain unobserved except at the terminal joint, where joint-side measurements are available and used as the sole supervision signal. Accordingly, the measurable signals are the actuator-side tendon triplets $\{\mathbf{T}_1^{(j)}(t)\}_{j=1}^N$ and joint angles $\{\boldsymbol{\theta}_k(t)\}_{k=1}^N$, equivalently $\mathbf{q}(t)$, while only the terminal joint-side tension $\mathbf{T}_N^{(N)}(t)$ is available for supervision and evaluation. The estimation target is the set of joint-side triplets at their own joints, $\{\hat{\mathbf{T}}_j^{(j)}(t)\}_{j=1}^N$. The sensing configuration is summarized in Fig. 1.

Authorized licensed use limited to: Shanghai Jiaotong University. Downloaded on July 13,2026 at 06:30:19 UTC from IEEE Xplore. Restrictions apply.

Let $j \in \{1, \ldots, N\}$ index the actuated joint and k index the guide location. The tendon index within each triplet is denoted by $i \in \{1, 2, 3\}$. Due to the serial routing, the tendon tension $\mathbf{T}_k^{(j)}(t) \in \mathbb{R}^3$ is defined only for the kinematic path $1 \leq k \leq j$. The system state comprises these tension triplets and the joint angles $\boldsymbol{\theta}_k(t) \in \mathbb{R}^2$. Accordingly, two configuration vectors are defined to represent the robot geometry: the full configuration $\mathbf{q}(t)$ and the local geometry $\mathbf{q}_k(t)$ for the adjacent pair $(k, k+1)$:

$$\mathbf{q}(t) = \left[\boldsymbol{\theta}_1^\top(t),\, \boldsymbol{\theta}_2^\top(t),\, \ldots,\, \boldsymbol{\theta}_N^\top(t)\right]^\top \in \mathbb{R}^{2N}, \tag{1}$$

$$\mathbf{q}_k(t) = \left[\boldsymbol{\theta}_k^\top(t),\, \boldsymbol{\theta}_{k+1}^\top(t)\right]^\top \in \mathbb{R}^4, \quad k = 1, \ldots, N-1. \tag{2}$$

The transmission modeling task is formulated as a causal, weakly supervised temporal regression. Given the causal history of observable inputs $\mathbf{X}(t) = [\mathbf{T}_1^{(1)}(t), \ldots, \mathbf{T}_1^{(N)}(t), \mathbf{q}(t)] \in \mathbb{R}^{5N}$, a map $\hat{\mathbf{Y}}(t) = \mathcal{F}(\mathbf{X}(0{:}t))$ is sought to estimate the controller-relevant joint-side tensions at their own joints, defined as $\mathbf{Y}(t) = [\mathbf{T}_1^{(1)}(t),\, \mathbf{T}_2^{(2)}(t),\, \ldots,\, \mathbf{T}_N^{(N)}(t)] \in \mathbb{R}^{3N}$. For $j = 1$, actuator-side and joint-side tensions coincide because routing starts at the first joint. For the tendon triplet actuating joint j, $\mathbf{T}_1^{(j)}(t)$ denotes the actuator-side measurement entering the routing chain, whereas $\mathbf{T}_j^{(j)}(t)$ denotes the joint-side tension at joint j.

B. Capstan Friction Model

The local geometry maps the adjacent-joint angles $\mathbf{q}_k(t)$, defined in (2), to a wrap/contact angle

$$\alpha_k(t) = \mathrm{Geom}(\mathbf{q}_k(t))\,, \qquad \alpha_k(t) \in [0, \pi], \tag{3}$$

where $\mathrm{Geom}(\cdot)$ is the geometry at joint/guide k as in [1]. Here the index k is the same joint/guide index defined in Section III-A.

Transmission is modeled across each joint/guide using a per-tendon Coulomb capstan relation with an angle-dependent friction law validated by recent analyses and experiments [3], [12]. For tendon $i \in \{1, 2, 3\}$, the friction coefficient under wrap angle α is modeled as

$$\mu_i(\alpha) = A_i \cos\alpha + B_i, \tag{4}$$

where A_i and B_i are per-tendon friction-law parameters. The capstan exponent at guide k for tendon i is then

$$\phi_{k,i}(t) = \mu_i(\alpha_k(t))\ \alpha_k(t). \tag{5}$$

Here $\phi_{k,i}$ is the product of the friction coefficient and the wrap angle; it determines the exponential tension attenuation in (6). Then the capstan update for the triplet that actuates joint j is written *componentwise for each tendon* i as

$$T_{k+1,i}^{(j)}(t) = \exp(-\,\beta_k(t)\,\sigma_{k,i}(t)\,\phi_{k,i}(t))\ T_{k,i}^{(j)}(t), \quad i = 1, 2, 3. \tag{6}$$

Contact and nonsmooth states: $\beta_k(t) \in \{0, 1\}$ indicates tendon–guide contact at joint/guide k, and $\sigma_{k,i}(t) \in \{-1, 0, +1\}$ encodes reverse slip / stick / forward slip for tendon i. The standard capstan hysteresis and stick bounds as formulated for cable–hole friction in [12] are followed. This analytical model serves two roles in the proposed framework: as a physics prior for regularizing the learned transmission ratios (Section III-C), and as an experimental baseline for quantifying improvement (Section IV).

TABLE I
SINGLE-PASSAGE TEST (AVERAGED OVER TENDONS). CAPSTAN BASELINES ARE CALIBRATED ON THE TRAINING/VALIDATION SPLIT BY CONSTRAINED NONLINEAR LEAST-SQUARES UNDER $\mu_i \geq 0$; HELD-OUT TEST TRAJECTORIES ARE NOT USED. NEGATIVE R^2 VALUES INDICATE THAT CALIBRATED STATIC CAPSTAN VARIANTS REMAIN STRUCTURALLY INSUFFICIENT UNDER HYSTERESIS-DOMINATED TRANSMISSION.

Method	Type	RMSE [N]↓	R^2↑	Max Err. [N]↓
Proximal	Measurement	36.046	−2.809	46.010
Capstan-simple [5]	Physics	47.673	−2.746	58.260
Capstan-contact [10]	Physics	47.022	−2.748	57.882
Capstan-nonsmooth [12]	Physics	36.478	−2.626	51.670
Transformer [13]	Data	8.862	0.492	14.596
LSTM [16]	Data	8.184	0.553	15.183
MLP [15]	Data	6.260	0.708	**11.099**
PG-CHTM	Physics-guided	**4.628**	**0.929**	11.835

C. Physics-Guided Cable-Hole Transmission Model

The PG-CHTM constitutes the atomic building block of the cascade. As illustrated in the detailed inset of Fig. 2, the unit comprises preprocessing, prediction, and output processing modules. Physics is explicitly injected at two stages: at the input via geometric conditioning, and in the loss via a parallel regularization branch.

1) Input Feature Construction: Raw sensory and kinematic streams are first buffered to maintain a sliding history window of length h. Within the *Preprocessing* block, these signals undergo filtering and channel-wise normalization. For an actuating tendon triplet j at guide k, the normalized feature vector $\tilde{\mathbf{X}}_k^{(j)}(t)$ is constructed by augmenting the upstream tension with the trigonometric encoding of the local geometry $\mathbf{q}_k(t)$, defined in (2), and the wrap angle:

$$\tilde{\mathbf{X}}_k^{(j)}(t) = \left[\sin(\mathbf{q}_k(t))^\top,\, \cos\alpha_k(t),\, \tilde{\mathbf{T}}_k^{(j)\top}(t)\right]^\top. \tag{7}$$

Here, $\sin(\cdot)$ denotes element-wise operation. The wrap angle is encoded as $\cos\alpha_k$ because the capstan exponent $\mu\alpha$ is approximately linear in α for small angles; $\cos\alpha$ provides a bounded, smooth representation without discontinuities at $\alpha = 0$.

2) Hybrid Ratio Prediction: The *Prediction* module maps the feature history to a positive tension ratio. It integrates a static MLP backbone f_θ with a strictly causal Transformer-based *Temporal Residual* g_ϕ. Let $\hat{r}^{\mathrm{MLP}}$ and r^{TR} denote their respective outputs. The data-driven ratio is synthesized via multiplicative composition:

$$\hat{r}_{k,i}^{\mathrm{data},(j)}(t) = \left(1 + r_{k,i}^{\mathrm{TR},(j)}(t)\right)\ \hat{r}_{k,i}^{\mathrm{MLP},(j)}(t). \tag{8}$$

Strict positivity is ensured by parameterizing $\hat{r}^{\mathrm{MLP}} = \mathrm{softplus}(\cdot) + \epsilon$ with $\epsilon = 0.01$, and bounding $r^{\mathrm{TR}} \in (-0.5, 0.5)$ via a scaled $\tanh$ output. Multiplicative composition aligns with the exponential capstan decay $T_{\mathrm{out}} \propto T_{\mathrm{in}}\, e^{-\mu\alpha}$ and preserves the zero-tension boundary.

A strictly causal Transformer is adopted for g_ϕ. Masked attention provides stable gradients across the multiplicative cascade where LSTM gradients would vanish exponentially, is computationally efficient on the short window $h \leq 20$, and naturally focuses on reversal-adjacent frames to capture hysteresis.

In parallel, the *Capstan Prior* block derives a physics-based reference ratio $\hat{\boldsymbol{r}}_k^{\mathrm{phys}}(t)$ by evaluating the analytical friction law

Authorized licensed use limited to: Shanghai Jiaotong University. Downloaded on July 13,2026 at 06:30:19 UTC from IEEE Xplore. Restrictions apply.

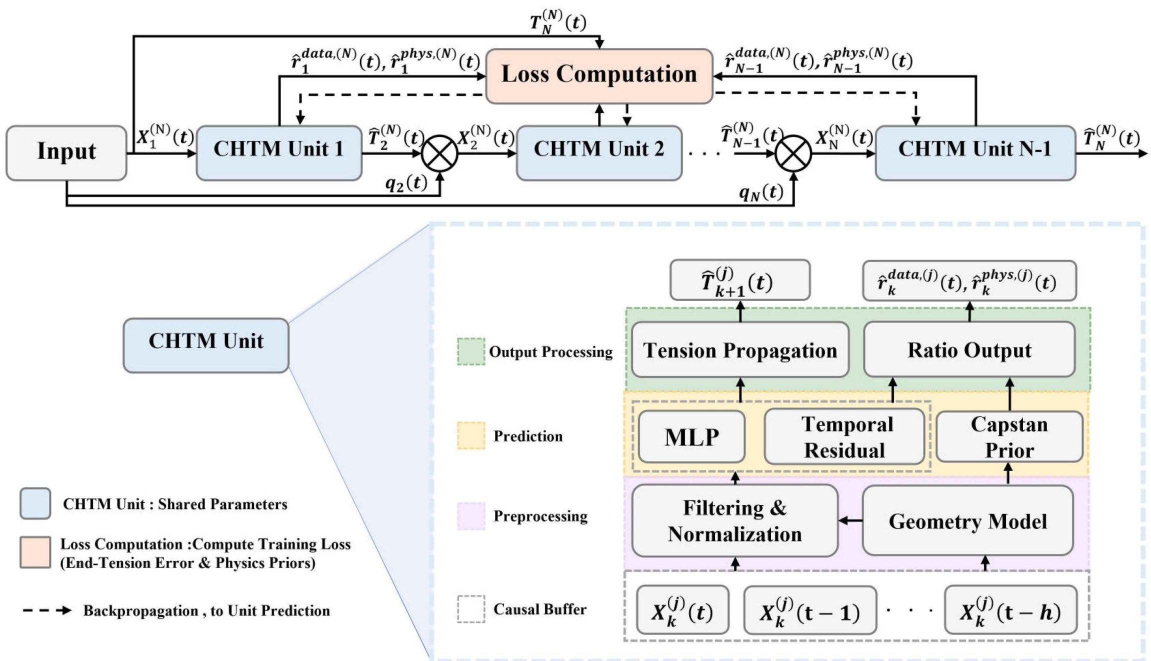

Fig. 2. PG-CHTC (Physics-Guided Cable-Hole Transmission Cascade) architecture. The shared-parameter cascade consists of identical PG-CHTM (Physics-Guided Cable-Hole Transmission Model) units trained with terminal joint-side supervision. Each unit (detailed inset) processes local geometry and upstream tension through prediction modules (f_θ, g_ϕ) with capstan regularization. Parameters are shared across all stages to improve practical identifiability. Dashed lines indicate training-only signals.

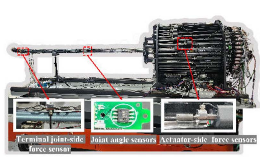

(a) Platform and sensors.

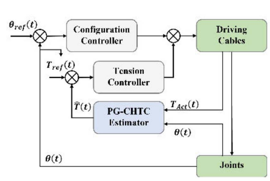

(b) Estimator-in-the-loop framework.

Fig. 3. Experimental platform and estimator-in-the-loop deployment. (a) Proximal in-line load cells, joint-angle encoders, and terminal joint-side load cells for evaluation. (b) The PG-CHTC estimator provides online joint-side tension estimates to the hybrid controller.

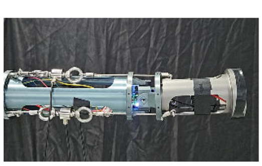

(a) Terminal-module-only

(b) Posterior-3

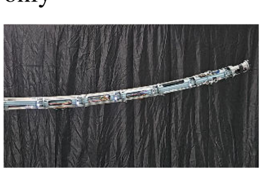

(c) Posterior-6

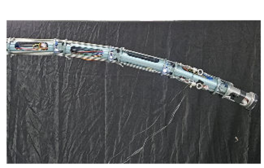

(d) Planned path

Fig. 4. Representative robot configurations across the four acquisition protocols defined in Section IV-A. "Posterior-n" denotes a step-and-hold protocol exciting only the distal n modules while all proximal joints remain fixed.

using the friction-law parameters $(\mathbf{A}, \mathbf{B})$:

$$\hat{\boldsymbol{r}}_k^{\text{phys}}(t) = \exp(-\boldsymbol{\mu}(\alpha_k(t))\,\alpha_k(t))\,. \tag{9}$$

This contact-agnostic form is used only as a soft regularizer, while nonsmooth/contact effects are captured by the data-driven branch. The β/σ states are not enforced to avoid brittle contact-state estimation, and nonsmoothness is learned by the residual branch. This reference is routed to the loss computation block in Fig. 2, top, to serve as a regularization target.

3) Output Processing and Propagation: In the *Output Processing* stage, the estimated ratio scales the upstream tension to propagate the state to the subsequent guide:

$$\hat{T}_{k+1,i}^{(j)}(t) = \hat{r}_{k,i}^{\text{data},(j)}(t)\,\hat{T}_{k,i}^{(j)}(t). \tag{10}$$

The updated tension $\hat{\mathbf{T}}_{k+1}^{(j)}(t)$ is then combined with the subsequent local geometry $\mathbf{q}_{k+1}(t)$ to refresh the causal buffer for the next unit.

D. Physics-Guided Cable-Hole Transmission Cascade

To address the identifiability constraints inherent in terminal joint-side supervision, the PG-CHTM units are arranged in a serial chain over $k = 1, \ldots, N-1$, as depicted in the main view of Fig. 2. A joint-level contact gate $\beta_k(t)$ is assumed shared by the three tendons at each guide. Crucially, the model parameters $\Theta = \{\theta, \phi, \mathbf{A}, \mathbf{B}\}$ are strictly shared across all units.

Assumption (Tribological homogeneity): All guide disks share similar tendon–guide interface characteristics, so a tied per-passage ratio law is a reasonable first-order approximation. Under terminal-only sensing, per-guide homogeneity is not

Authorized licensed use limited to: Shanghai Jiaotong University. Downloaded on July 13,2026 at 06:30:19 UTC from IEEE Xplore. Restrictions apply.

directly identifiable; no training signature attributable to substantial mismatch is observed in the present data. Severe wear or damage would violate this assumption, motivating the relaxation discussed in Section V.

Assumption (Quasi-static operation): Transmission is treated as quasi-static, so the static tension-ratio law is used as a first-order approximation. Joint angular velocity stays below $9^\circ\,\mathrm{s}^{-1}$ for 99% of the recorded dataset, consistent with the quasi-static treatments in [3], [12]. The short-memory residual g_ϕ accounts for hysteretic departures from this approximation.

Assumption (No-slack operation): Slack handling is not within the present scope. The low-level puller-follower controller [26] antagonistically maintains non-zero tension, and the estimator is deployed only within this taut regime; across 7164 closed-loop frames, the minimum tendon tension is 62.45 N and the fifth percentile is 77.2 N, above the 50 N pretension threshold. The zero-propagation property of (10) additionally avoids additive-offset artifacts when measured upstream tension approaches zero.

End-to-end training is governed by the *Loss Computation* block. The multi-objective loss $\mathcal{L}$ aggregates signals from all stages to minimize terminal discrepancy while regularizing latent ratios:

$$\mathcal{L} = \sum_t \left[\mathcal{L}_{\text{data}}(t) + \lambda\,\mathcal{L}_{\text{phys}}(t) + \lambda_s\,\mathcal{L}_{\text{smooth}}(t)\right]. \tag{11}$$

The components are formulated in the logarithmic domain to linearize the multiplicative transmission chain, thereby stabilizing gradients and framing the loss as a relative error:

$$\mathcal{L}_{\text{data}}(t) = \left\| \ln \hat{\mathbf{T}}_N^{(N)}(t) - \ln \mathbf{T}_N^{(N)}(t) \right\|_2^2, \tag{12}$$

$$\mathcal{L}_{\text{phys}}(t) = \sum_{k=1}^{N-1} \left\| \ln \hat{\boldsymbol{r}}_k^{\text{data},(N)}(t) - \ln \hat{\boldsymbol{r}}_k^{\text{phys},(N)}(t) \right\|_2^2, \tag{13}$$

$$\mathcal{L}_{\text{smooth}}(t) = \sum_{k=1}^{N-2} w_k(t) \left\| \ln \hat{\boldsymbol{r}}_k^{\text{data},(N)}(t) - \ln \hat{\boldsymbol{r}}_{k+1}^{\text{data},(N)}(t) \right\|_2^2. \tag{14}$$

Each term serves a distinct physical purpose. The primary data term, $\mathcal{L}_{\text{data}}$, anchors the total product of the cascade to the terminal joint-side labels. However, $\mathcal{L}_{\text{data}}$ alone cannot solve the identifiability problem. The *Internal Physics Regularization*, $\mathcal{L}_{\text{phys}}$, acts as a per-guide regularizer, providing a physics-informed soft regularization target for each unobserved intermediate ratio to constrain the solution space to physically plausible manifolds. Finally, $\mathcal{L}_{\text{smooth}}$ imposes a *Spatial Physical Consistency* constraint. The adaptive weight $w_k(t)$ is rigorously defined as a Gaussian kernel:

$$w_k(t) = \exp\left(-\eta_w \|\alpha_k(t) - \alpha_{k+1}(t)\|^2\right), \tag{15}$$

where the *smoothness sensitivity* $\eta_w > 0$ governs how sharply the weight decays with the angular difference between adjacent guides: a larger η_w confines the smoothness penalty to guide pairs whose configurations are nearly identical, while a smaller value extends the penalty to geometrically dissimilar pairs. The value $\eta_w = 2.0$ is selected via the one-at-a-time sensitivity sweep summarized in Section IV-B, balancing inter-guide regularization strength against flexibility at high-curvature inflection points.

Algorithm 1: Causal Cascade Forward Pass at Time t.

Input : Actuator-side tensions $\{\mathbf{T}_1^{(j)}(t)\}_{j=1}^N$, angles $\mathbf{q}(t)$, buffer $\mathcal{B}$
Params : Shared PG-CHTM predictors f_θ, g_ϕ
Output : Joint-side tensions $\{\hat{\mathbf{T}}_j^{(j)}(t)\}_{j=1}^N$

for $j \leftarrow 1$ **to** N **do**
  $\hat{\mathbf{T}}_1^{(j)}(t) \leftarrow \mathbf{T}_1^{(j)}(t)$
for $k \leftarrow 1$ **to** $N-1$ **do**
  $\mathbf{q}_k(t) \leftarrow$ `ExtractLocalGeometry` $(\mathbf{q}(t), k)$
  $\alpha_k(t) \leftarrow \mathrm{Geom}(\mathbf{q}_k(t))$
  for $j \leftarrow k+1$ **to** N **do**
    $\tilde{\mathbf{X}}_k^{(j)}(t) \leftarrow$ `BuildUnitInputVector` $(\mathbf{q}_k, \alpha_k, \hat{\mathbf{T}}_k^{(j)})$
    `UpdateCausalBuffer` $(\mathcal{B}_{j,k}, \tilde{\mathbf{X}}_k^{(j)}(t))$
    $\hat{\boldsymbol{r}}_k^{\text{data},(j)}(t) \leftarrow$ `PGCHTMUnit` $(f_\theta, g_\phi, \mathcal{B}_{j,k})$
    $\hat{\mathbf{T}}_{k+1}^{(j)}(t) \leftarrow \hat{\boldsymbol{r}}_k^{\text{data},(j)}(t) \odot \hat{\mathbf{T}}_k^{(j)}(t)$

Identifiability Analysis: Here, 'practical identifiability' refers to the determinacy of the shared ratio law from terminal-only supervision under diverse geometry excitation, rather than formal structural identifiability of all latent states. The necessity of parameter sharing is justified by examining the cumulative transmission in the logarithmic domain:

$$\ln \hat{\mathbf{T}}_N^{(N)} = \ln \mathbf{T}_1^{(N)} + \sum_{k=1}^{N-1} \ln \hat{\boldsymbol{r}}_k(\tilde{\mathbf{X}}_k^{(N)}; \Theta_k), \tag{16}$$

if each guide employed independent parameters Θ_k, infinitely many local ratio combinations could yield the same terminal sum, rendering intermediate states unobservable. The shared constraint $\Theta_k = \Theta$ reduces degrees of freedom, and practical identifiability is improved provided that the training data exhibits sufficient geometric diversity. This does not imply unique recovery of guide-specific friction parameters or contact/slip states; the learned object is a shared ratio law constrained by terminal supervision and the capstan prior.

At runtime, the recursive estimation procedure is executed as summarized in Algorithm 1.

IV. Experiments

The proposed framework is evaluated in two phases: open-loop logged-data estimation evaluation and estimator-in-the-loop closed-loop control experiments. Section IV-A introduces the experimental platform, sensing protocol, and dataset construction. Section IV-B reports open-loop estimation performance, model selection, and ablations under terminal joint-side supervision. Section IV-C evaluates the deployed estimator in a 10 Hz closed-loop tracking task.

A. Experimental Setup and Dataset Construction

Validation is performed on the SJTU-Snake III [3], a 12-module CDHRR platform implementing the kinematics described in Section III-A. As shown in Fig. 3(a), joint angles are measured by 12-bit magnetic rotary encoders, and tendon tensions by in-line load cells with 500 N range; load cells of the

Authorized licensed use limited to: Shanghai Jiaotong University. Downloaded on July 13,2026 at 06:30:19 UTC from IEEE Xplore. Restrictions apply.

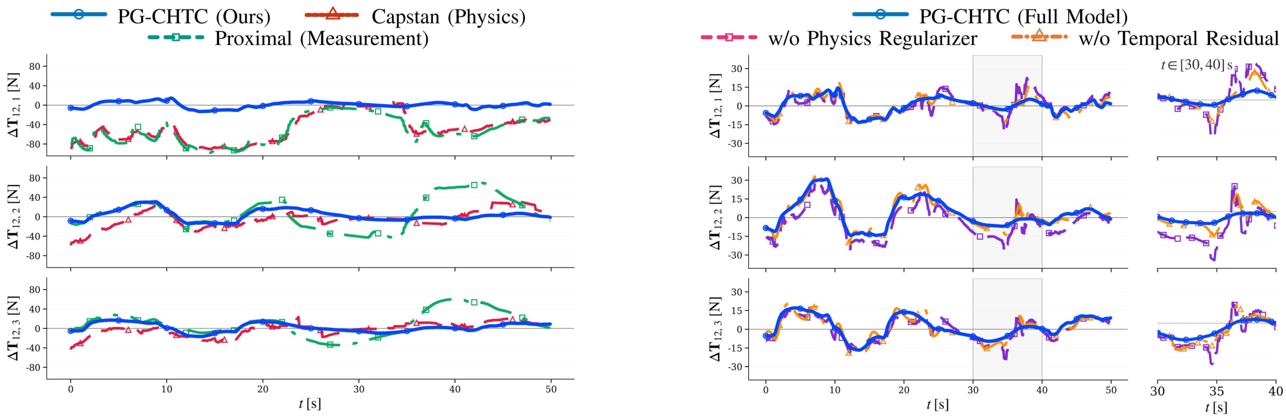

(a) Baselines comparison.

(b) Ablation study.

Fig. 5. Posterior-6 (distal six modules active) terminal-joint estimation errors $\Delta \mathbf{T} = \hat{\mathbf{T}}_N^{(N)} - \mathbf{T}_N^{(N)}$. Ablation variants: *w/o Physics Regularizer* sets $\lambda = \lambda_s = 0$; *w/o Temporal Residual* removes g_ϕ. Full protocol metrics for all datasets are in Table II.

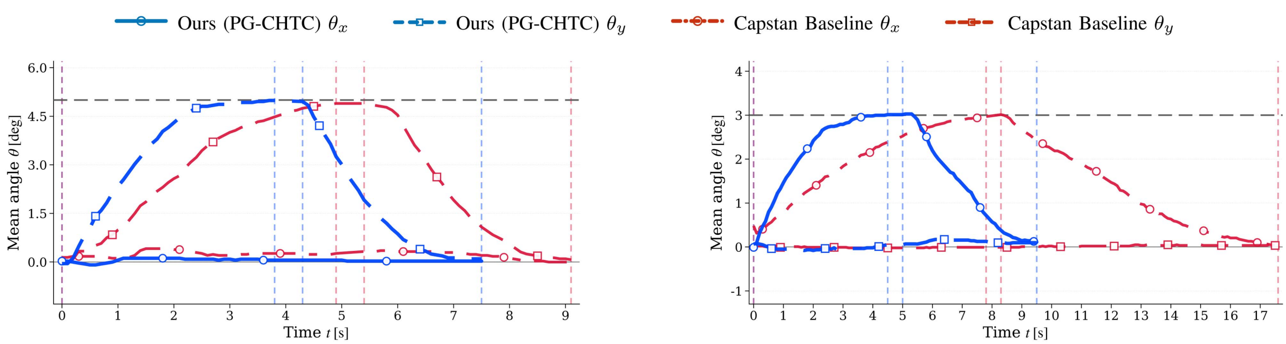

(a) **Posterior-3**: 0° → 5° → 0°.

(b) **Posterior-6**: 0° → 3° → 0°.

Fig. 6. Closed-loop tracking at 10 Hz. PG-CHTC is drawn in blue and the Capstan baseline in red; linestyles and markers distinguish θ_x and θ_y. Dashed lines mark the commanded amplitude and rise/settling times on the ±0.5° band; aggregated metrics in Table III.

TABLE II
OPEN-LOOP LOGGED-DATA EVALUATION OF CASCADED ESTIMATION (AVERAGED OVER TENDONS). CAPSTAN BASELINES USE THE SAME CALIBRATION PROTOCOL AS TABLE I.

	Posterior-3 (T=800)		Posterior-6 (T=500)		Planned path (T=1260)		Overall (frame-weighted)	
Method	RMSE [N]↓	R^2↑	RMSE [N]↓	R^2↑	RMSE [N]↓	R^2↑	RMSE [N]↓	R^2↑
Reference baselines								
Proximal measurement	48.522	−0.655	73.075	−2.747	32.699	0.486	45.530	−0.502
Physics Baseline (Capstan-nonsmooth) [3], [12]	37.225	0.475	60.480	−0.181	28.769	0.476	37.605	0.347
Ablations of PG-CHTC								
w/o Physics Regularizer	19.940	0.674	18.097	0.468	9.531	0.893	14.457	0.742
w/o Temporal Residual	11.907	0.882	17.631	0.585	9.924	0.896	12.049	0.831
Proposed								
PG-CHTC (Full Model)	**11.230**	**0.906**	**16.256**	**0.707**	**9.145**	**0.908**	**11.185**	**0.868**

TABLE III
CLOSED-LOOP CONTROL METRICS. VALUES FROM LOOP 1 (STEP) AND LOOP 2 (RETURN); SETTLING COMPUTED ON THE ±0.5° BAND. PAIRS WITH FIG. 6.

Test Regime / Estimator		**Step to target (loop 1)**			**Return to zero (loop 2)**		
Regime	Estimator Model	Rise [s]↓	Settle [s]↓	RMS [°]↓	Rise [s]↓	Settle [s]↓	RMS [°]↓
Posterior-3 (0° → 5° → 0°)	Ours (PG-CHTC)	1.86	2.50	0.19	1.78	2.50	0.21
	Capstan Baseline [3]	3.28	4.50	0.68	2.01	3.10	0.29
Posterior-6 (0° → 3° → 0°)	Ours (PG-CHTC)	2.21	3.10	0.72	2.84	3.90	0.83
	Capstan Baseline [3]	5.14	6.30	1.01	6.45	8.00	1.03

Authorized licensed use limited to: Shanghai Jiaotong University. Downloaded on July 13,2026 at 06:30:19 UTC from IEEE Xplore. Restrictions apply.

same type are mounted at the terminal joint to provide ground truth. All sensor streams are synchronized at 10 Hz, filtered, and normalized. Data splits are partitioned by trajectory to strictly prevent temporal leakage.

To capture configuration-dependent transmission effects, four datasets were acquired (Fig. 4):

i) *Terminal-module-only:* 50000 frames; quasi-static terminal-joint scans to maximize wrap-angle coverage and induce reversals.
ii) *Posterior-3:* 5000 frames; step-and-hold on modules 10–12 with $\pm 3^\circ/\pm 5^\circ$ amplitudes, proximal joints fixed.
iii) *Posterior-6:* 3000 frames; extended steps on modules 7–12 with $\pm 3^\circ$ amplitudes.
iv) *Planned paths:* 10000 frames; randomized, collision-free joint-space trajectories with multi-joint motion.

The *Terminal-module-only* set supports the unit-level analysis in Table I; *Posterior-3/6* with *Terminal-module-only* subsets are used for cascade training and evaluation in Section IV-B; *Planned paths* and step protocols are used for estimator-in-the-loop experiments in Section IV-C.

B. Open-Loop Logged-Data Estimation Evaluation and Ablations

This replay-based assessment is treated as open-loop logged-data evaluation because estimator outputs do not affect robot motion. The PG-CHTC cascade is evaluated under the strict constraint of terminal joint-side supervision. Training utilizes a mixed subset of the Terminal-module-only, Posterior-3, and Posterior-6 datasets (23000 frames after downsampling, with a 70/30 train/validation split), while testing is conducted on independent Posterior-3, Posterior-6, and Planned-path trajectories.

The PG-CHTM combines a static MLP backbone f_θ (hidden widths $512/512/256$, GELU, BatchNorm1d) with a 3-layer causal Transformer residual g_ϕ (d_{model}=128, 4 heads, history window $h = 20$). Both stages use AdamW with weight decay 10^{-4} and cosine learning-rate annealing: f_θ is pretrained via the cascade-end loss against terminal-joint targets (lr 10^{-5}, batch 128, up to 1500 epochs, patience 100), then frozen while g_ϕ is trained on sequence chunks (lr 10^{-3}, chunk length 1024, up to 120 epochs, patience 15). Loss weights are set to $\lambda = 1.0$ and $\lambda_s = 0.1$. Data-driven baselines (MLP, LSTM, and Transformer) are sized to match the $\sim 400\,\mathrm{k}$ parameter count of f_θ for fair comparison.

On an Intel i9-13900HX CPU, the 12-module cascade runs at 23.48 ms per frame, comfortably within the 100 ms control period; the MLP backbone accounts for 14.28 ms, the causal Transformer residual for 8.75 ms, and the friction-prior evaluation and buffer updates for the remaining 0.45 ms. Further software-level speedup is available through quantization or backbone-width reduction.

The operating point is selected on the validation split at $\lambda = 1.0$, $\lambda_s = 0.1$, $\eta_w = 2.0$, and $h = 20$ ($h = 20$ corresponds to 2 s at 10 Hz). A one-at-a-time validation sweep over $\lambda \in [0.03, 30]$, $\lambda_s \in [0.003, 3]$, $\eta_w \in [0.25, 16]$, and $h \in [5, 80]$ shows that performance remains stable around this setting, with worst-to-best RMSE ratios of 1.18 (λ), 1.22 (λ_s), 1.30 (η_w), and 1.12 (h).

1) Unit-Level Validation: All Capstan baselines implement the analytical friction model of (4)–(6), and instantiate $(\beta_k, \sigma_{k,i})$ in (6) differently: Capstan-simple takes $\beta_k \equiv 1$, $\sigma_{k,i} \equiv 1$; Capstan-contact derives β_k from the tendon–hole contact geometry with $\sigma_{k,i} \equiv 1$; Capstan-nonsmooth combines this β_k with $\sigma_{k,i}$ from the hysteresis formulation of Section III-B. Their per-tendon coefficients (A_i, B_i) are fitted on the training/validation split by constrained nonlinear least-squares under $\mu_i(\alpha) \geq 0$; held-out test trajectories are not used for calibration. The identified per-tendon coefficients are $(A_i, B_i) = (0.30, 0.09)$, $(0.32, 0.08)$, $(0.28, 0.10)$ for the three tendons.

The architectural choice of the PG-CHTM unit is verified prior to cascade analysis. As reported in Table I, the proposed hybrid unit outperforms both analytical models and pure data-driven baselines on the held-out single-passage task. The PG-CHTM achieves an RMSE of 4.628 N, significantly lower than the LSTM baseline at 8.184 N and the static MLP at 6.260 N. This result validates the design premise that combining a static backbone with a temporal residual offers a superior accuracy-efficiency trade-off compared to fully recurrent architectures.

2) Cascade Performance and Ablation Studies: Cascade estimation performance is detailed in Table II, with representative baseline and ablation comparisons shown in Fig. 5. Relative to the physics baseline, compounding drift is compensated for and errors are kept near the zero-error reference by the PG-CHTC estimator.

The sources of these gains are isolated in ablation studies. When physics regularization is removed, significant accuracy loss is observed, with RMSE increasing from 11.185 N to 14.457 N. This empirical evidence supports the theoretical argument in Section III-D: the physics regularization constrains the solution space under terminal-only joint-side supervision and reduces persistent terminal-joint bias.

The temporal residual module g_ϕ further improves dynamic fidelity, especially around motion reversals. When omitted, sharp error spikes appear at reversals (Fig. 5); the full model suppresses these transients, indicating the residual's role in capturing history-dependent viscoelastic effects beyond the static backbone.

C. Estimator-in-The-Loop Closed-Loop Control Experiments

The trained estimator is deployed within a 10 Hz hybrid tension and configuration control framework adapted from [3], executing synchronously with the control loop (Fig. 3(b)). Evaluation on *Posterior-3* and *Posterior-6* step protocols applies standardized rise-time and settling metrics to compare the controller driven by the analytical Capstan baseline against the one driven by the proposed PG-CHTC estimator (Table III and Fig. 6).

For the *Posterior-3* step-and-return task, the PG-CHTC-driven controller reduces settling time by 44.4% (from 4.50 s to 2.50 s) and RMS error by 72.0% (from 0.68° to 0.19°) relative to the baseline. In Fig. 6, the baseline system exhibits sluggish response and steady-state offsets due to friction-induced estimation bias, whereas the PG-CHTC system converges rapidly with minimal overshoot.

V. Conclusion

This letter presents PG-CHTC, a physics-guided cascade framework for estimating joint-side tendon tensions in discrete-joint CDHRRs from actuator-side measurements. By recursively applying a shared tension-ratio unit with capstan-based regularization and a causal temporal residual, the multiplicative architecture enforces physical consistency and improves practical identifiability under terminal-only joint-side measurements. Validation on a 12-module CDHRR shows that the estimator

Authorized licensed use limited to: Shanghai Jiaotong University. Downloaded on July 13,2026 at 06:30:19 UTC from IEEE Xplore. Restrictions apply.

mitigates the compounding friction errors of analytical baselines [5], [12] while remaining stable relative to purely data-driven approaches [13]. The estimator sustains 10 Hz closed-loop control and serves as a virtual sensor for systems with limited terminal instrumentation.

The current formulation assumes tribological homogeneity, quasi-static operation, and a taut-tendon regime maintained by the low-level controller; wear, fast motions, or slack incursions fall beyond this scope. Future work will relax strict weight tying via low-dimensional per-guide adaptation and study how sparse intermediate sensing can calibrate spatial variations while preserving practical identifiability.

References

[1] W. Xu, T. Liu, and Y. Li, "Kinematics, dynamics, and control of a cable-driven hyper-redundant manipulator," *IEEE/ASME Trans. Mechatron.*, vol. 23, no. 4, pp. 1693–1704, Aug. 2018.

[2] Z. Chen, X. Zhang, M. Liu, J. Huo, H. Zhang, and L. Tan, "Review of cable-driven hyper-redundant robotic arms," in *Proc. IEEE Int. Conf. Robot., Artif. Intell. Intell. Control*, 2024, pp. 110–119.

[3] Z. Zhu, Z. Li, Z. Peng, C. Liu, and G. Gu, "Hybrid tension and configuration control of cable-driven hyper-redundant robots for high accuracy and stability," *IEEE Robot. Automat. Lett.*, vol. 10, no. 6, pp. 5601–5608, Jun. 2025.

[4] B. Liang, W. Li, B. Lin, L. Yan, X. Wang, and W. Xu, "Hybrid tension-position control of cable-driven serial manipulator to improve accuracy and avoid cable slack and breakage," *IEEE/ASME Trans. Mechatron.*, vol. 30, no. 2, pp. 1072–1083, Apr. 2025.

[5] D. C. Rucker and R. J. Webster III, "Statics and dynamics of continuum robots with general tendon routing and external loading," *IEEE Trans. Robot.*, vol. 27, no. 6, pp. 1033–1044, Dec. 2011.

[6] Z. Li, S. Zheng, Y. Zhang, H. Li, Z. Jiang, and P. Shi, "Multilevel motion control of cable-driven hyper-redundant manipulators," *IEEE/ASME Trans. Mechatron.*, vol. 30, no. 6, pp. 5364–5375, Dec. 2025.

[7] Y. Shi et al., "Human-in-the-loop modeling and control of an upper limb exosuit with tendon-sheath actuation," *IEEE Robot. Automat. Lett.*, vol. 9, no. 6, pp. 5919–5926, Jun. 2024.

[8] J. H. Jung, N. Pan, and T. J. Kang, "Capstan equation including bending rigidity and non-linear frictional behavior," *Mechanism Mach. Theory*, vol. 43, no. 6, pp. 661–675, 2008.

[9] D. B. Camarillo, C. F. Milne, C. R. Carlson, M. R. Zinn, and J. K. Salisbury, "Mechanics modeling of tendon-driven continuum manipulators," *IEEE Trans. Robot.*, vol. 24, no. 6, pp. 1262–1273, Dec. 2008.

[10] H. Yuan, L. Zhou, and W. Xu, "A comprehensive static model of cable-driven multi-section continuum robots considering friction effect," *Mechanism Mach. Theory*, vol. 135, pp. 130–149, 2019.

[11] W. Li, X. Huang, L. Yan, H. Cheng, B. Liang, and W. Xu, "Force sensing and compliance control for a cable-driven redundant manipulator," *IEEE/ASME Trans. Mechatron.*, vol. 29, no. 1, pp. 777–788, Feb. 2024.

[12] S. Zhang et al., "Quasi-static modeling of a cable-driven continuum manipulator considering non-smooth cable-hole friction and experimental verification," *Mechanism Mach. Theory*, vol. 204, 2024, Art. no. 105856.

[13] T.-C. Çallar and S. Böttger, "Hybrid learning of time-series inverse dynamics models for locally isotropic robot motion," *IEEE Robot. Automat. Lett.*, vol. 8, no. 2, pp. 1061–1068, Feb. 2023.

[14] H.-M. Kim, D. Kim, S. H. Park, and S. Jin, "Hysteresis compensation of tendon-sheath mechanism using nonlinear programming based on Preisach model," *IEEE Robot. Automat. Lett.*, vol. 9, no. 6, pp. 5246–5253, Jun. 2024.

[15] Y. Wang et al., "Using neural networks to model hysteretic kinematics in tendon-actuated continuum robots," in *Proc. Int. Symp. Med. Robot.*, 2024, pp. 1–7.

[16] D. Wu, Y. Zhang, M. Ourak, K. Niu, J. Dankelman, and E. V. Poorten, "Hysteresis modeling of robotic catheters based on long short-term memory network for improved environment reconstruction," *IEEE Robot. Automat. Lett.*, vol. 6, no. 2, pp. 2106–2113, Apr. 2021.

[17] Q. Gao, G. Ji, M. Sun, Y. Xiao, H. Rao, and Z. Sun, "Dynamic hysteresis compensation for tendon-sheath mechanism in flexible surgical robots without distal perception," *IEEE Trans. Robot.*, vol. 41, pp. 3703–3721, 2025.

[18] J. Hong, D. Hong, C. Kim, and S. Won, "On-line shape estimation for hysteresis compensation in tendon-sheath mechanisms using endoscopic camera," *IEEE Robot. Automat. Lett.*, vol. 10, no. 6, pp. 5201–5208, Jun. 2025.

[19] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *J. Comput. Phys.*, vol. 378, pp. 686–707, 2019.

[20] G. E. Karniadakis, I. G. Kevrekidis, L. Lu, P. Perdikaris, S. Wang, and L. Yang, "Physics-informed machine learning," *Nature Rev. Phys.*, vol. 3, no. 6, pp. 422–440, 2021.

[21] M. Bensch, T.-D. Job, T.-L. Habich, T. Seel, and M. Schappler, "Physics-informed neural networks for continuum robots: Towards fast approximation of static Cosserat rod theory," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2024, pp. 17293–17299.

[22] H. Hu, Z. Shen, and C. Zhuang, "A PINN-based friction-inclusive dynamics modeling method for industrial robots," *IEEE Trans. Ind. Electron.*, vol. 72, no. 5, pp. 5136–5144, May 2025.

[23] G. Shi et al., "Neural lander: Stable drone landing control using learned dynamics," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2019, pp. 9784–9790.

[24] X. Huang, Y. Rong, and G. Gu, "High-precision dynamic control of soft robots with the physics-learning hybrid modeling approach," *IEEE/ASME Trans. Mechatron.*, vol. 30, no. 3, pp. 1658–1669, Jun. 2025.

[25] E. Sebastián, T. Duong, N. Atanasov, E. Montijano, and C. Sagüés, "Physics-informed multiagent reinforcement learning for distributed multirobot problems," *IEEE Trans. Robot.*, vol. 41, pp. 4499–4517, 2025.

[26] Y. Zheng et al., "Design and validation of cable-driven hyper-redundant manipulator with a closed-loop puller-follower controller," *Mechatronics*, vol. 78, 2021, Art. no. 102605.

Authorized licensed use limited to: Shanghai Jiaotong University. Downloaded on July 13,2026 at 06:30:19 UTC from IEEE Xplore. Restrictions apply.