The 11th International Congress
Motor Vehicles & Motors 2026
Kragujevac, Serbia
October 8th-9th, 2026

Motor Vehicles & Motors

MVM2026-035
https://doi.org/10.24874/CMVM2026.035
Original scientific paper

Marijana Jeremić[1]
Lazar Krstić[2]
Miloš Ivanović[3]
Mihailo Lazarević[4]
Milan Matijević[5]

DIGITAL TWIN DESIGN FOR TORSIONAL STIFFNESS ESTIMATION IN MOTOR-LOAD OSCILLATORY SYSTEMS USING PHYSICS-INFORMED NEURAL NETWORKS

ABSTRACT: The modern development of electric vehicles (EVs) faces significant challenges driven by the high power density and rapid torque response of electric motors. These characteristics frequently excite undesirable torsional oscillations in the drivetrain (motor-load coupling), adversely affecting ride comfort, efficiency, and the structural longevity of mechanical components. Real-time monitoring of dynamic stiffness within these systems is crucial for predictive maintenance and adaptive control. However, direct stiffness measurement requires expensive and fragile physical sensors. This paper proposes a Digital Twin framework for monitoring the stiffness of a torsional oscillatory system using Physics-Informed Neural Networks (PINNs). Unlike purely data-driven black-box AI models, the proposed PINN approach embeds the governing physics of the mechanical system directly into the neural network's loss function. This allows the time-varying stiffness parameter to be accurately estimated even with sparse training data, relying only on standard vehicle sensor inputs like speed encoders. Simulation results demonstrate that the developed Digital Twin successfully identifies stiffness degradation in real-time, providing a robust baseline or advanced vibration damping and control strategies in modern electric vehicles.

KEYWORDS: digital twin, electric vehicles, torsional oscillations, physics-informed neural networks, stiffness estimation

[1] *Marijana Jeremić, Faculty of Engineering at University of Kragujevac, Sestre Janjić 6, 34000 Kragujevac, Serbia, marijanajeremic002@gmail.com, https://orcid.org/0009-0002-5620-1445*
[2] *Lazar Krstić, University of Kragujevac, Faculty of Science, Radoja Domanovica 12, Kragujevac, 34000, Serbia, lazar.krstic@pmf.kg.ac.rs, https://orcid.org/0000-0002-4703-2291*
[3] *Miloš Ivanović, Faculty of Information Studies, Ljubljanska cesta 31a, 8000 Novo Mesto, Slovenia, milos.ivanovic@fis.unm.si, https://orcid.org/0000-0002-8974-2267*
[4] *Mihailo Lazarević, Faculty of Mechanical Engineering at University of Belgrade, Kraljice Marije 16, 11000 Beograd, Serbia, mlazarevic@mas.bg.ac.rs, https://orcid.org/0000-0002-3326-6636*
[5] *Milan Matijević, Faculty of Engineering at University of Kragujevac, Sestre Janjić 6, 34000 Kragujevac, Serbia, matijevic@kg.ac.rs, https://orcid.org/0000-0001-5307-0027*

PROJEKTOVANJE DIGITALNOG BLIZANCA ZA PROCENU TORZIONE KRUTOSTI U OSCILATORNIM SISTEMIMA MOTOR-OPTEREĆENJE KORIŠĆENJEM NEURONSKIH MREŽA PODRŽANIH FIZIČKIM ZAKONIMA

REZIME: Savremeni razvoj električnih vozila (EV) suočava se sa značajnim izazovima koji proizilaze iz velike gustine snage i brzog odziva obrtnog momenta električnih motora. Ove karakteristike često izazivaju neželjene torzione oscilacije u pogonskom sistemu (sprega motor–opterećenje), što negativno utiče na udobnost vožnje, energetsku efikasnost i dugotrajnost mehaničkih komponenti. Praćenje dinamičke krutosti ovih sistema u realnom vremenu od ključnog je značaja za prediktivno održavanje i adaptivno upravljanje. Međutim, direktno merenje krutosti zahteva primenu skupih i osetljivih fizičkih senzora. U ovom radu predlaže se okvir digitalnog blizanca za praćenje krutosti torzionog oscilatornog sistema korišćenjem neuronskih mreža vođenih fizičkim zakonima (Physics-Informed Neural Networks – PINNs). Za razliku od klasičnih pristupa zasnovanih isključivo na podacima i modelima veštačke inteligencije tipa „crne kutije“, predloženi PINN pristup ugrađuje fizičke zakonitosti mehaničkog sistema direktno u funkciju greške neuronske mreže. Na taj način omogućena je precizna procena vremenski promenljive krutosti čak i u uslovima ograničene količine podataka za obuku, oslanjajući se samo na standardne senzorske informacije dostupne u vozilu, kao što su enkoderi brzine. Rezultati simulacija pokazuju da razvijeni digitalni blizanac uspešno identifikuje degradaciju krutosti u realnom vremenu, čime predstavlja pouzdanu osnovu za razvoj naprednih strategija prigušenja vibracija, upravljanja i prediktivnog održavanja u savremenim električnim vozilima.

KLJUČNE REČI: digitalni blizanac, električna vozila, torzione oscilacije, neuronske mreže informisane fizikom (pinns), procena krutosti

INTRODUCTION

Electrified drivetrains combine high power density with rapid torque transients, making flexible shafts, couplings, gears, and connected inertias susceptible to torsional oscillation. These oscillations affect vibration, acoustic behavior, transmitted loads, and component durability (Hua, X. & Gandee, E. 2021). Torsional stiffness is therefore both a structural parameter and a potential condition indicator: a progressive decrease may reflect coupling degradation, loosening, fatigue, or another loss of mechanical integrity (Li, W. et al. 2016; Park, Y.-H. et al. 2023). Because direct time-resolved stiffness measurements are rarely available in standard drive systems, its evolution must be inferred from accessible excitation and encoder information (Li, W. et al. 2016).

Digital twins integrate a physical representation of an asset with operating information and an updating computational mode. For predictive maintenance, the useful output is not merely a fitted signal but a physically interpretable parameter that can be tracked as degradation develops. In a motor–load system, electromagnetic torque and motor/load encoder speeds provide practical inputs, while the torsional stiffness remains latent. The resulting inverse problem is challenging because errors in reconstructed relative motion can be absorbed by the stiffness estimate, particularly when high-order derivatives are obtained from a flexible neural approximation. Physics-informed neural networks incorporate governing equations into learning objectives and have been established for forward and inverse identification problems (Ivanović, M. 2023; Ivanović, M. et al. 2026). The broader physics-informed machine-learning framework offers a route to combine partial observations, prior physical knowledge, and constrained parameter estimation (Ivanović, M. 2023). In the present problem, however, a pointwise second-order residual can produce an apparently accurate state fit while retaining a biased stiffness estimate. The proposed formulation mitigates the identifiability difficulty by predicting relative angle and relative speed as separate states and enforcing the motor–load dynamics through overlapping weak integral residuals that avoid second derivatives.

The resulting digital twin couples a first-order RelativeStateNet with a physically bounded, monotonically decreasing sigmoid stiffness model. It is evaluated through constant-stiffness controls, clean and noisy time-varying degradation, reduced sensor supervision with 751 sensor-labelled instants and dense unlabeled physics collocation, and causal warm-started online updates. The study uses the time vector and electromagnetic-torque profile provided in the jera1.mat dataset; the motor and load encoder-speed responses are synthetic signals generated by the reference ODE model. Accordingly, the contribution is presented as a simulation validation of time-varying stiffness identification and causal near-real-time monitoring on the tested CPU, while experimental drivetrain validation remains future work. However, the combination of time-varying torsional-stiffness identification,

weak first-order physics enforcement, reduced sensor supervision, and causal degradation monitoring has received considerably less attention in flexible motor–load digital twins.

Related Work

Reduced-order torsional models remain useful because they preserve dominant drivetrain modes while keeping parameter identification tractable (Ivanović, M. et al. 2026). Detailed electric-drive studies show that torque ripple, gear excitation, stiffness, and inertial coupling materially affect drivetrain vibration (Hua, X. & Gandee, E. 2021). Online parameter estimators with limited sensors have traditionally used recursive or adaptive schemes, whereas predictive-maintenance studies emphasize traceable health indicators and uncertainty-aware validation (Li, W. et al. 2016; Park, Y.-H. et al. 2023).

Physics-Informed Neural Networks (PINNs) combine data fitting with governing differential-equation constraints through a composite loss function. In inverse problems, unknown physical parameters can be estimated jointly with latent system states by minimizing both data-mismatch and physics-residual terms (Ivanović, 2023). The present work extends this physics-informed principle by replacing the conventional pointwise second-order formulation with weak integral constraints over temporal windows. The flexible motor–load reference implementation that motivated the physical structure is available as a software repository. Here the weak form is applied to a two-inertia ODE and is used specifically to identify a bounded time-varying torsional stiffness from encoder-derived relative quantities. In contrast to approaches focused primarily on forward state prediction or control-input reconstruction, the present work formulates torsional stiffness itself as the latent health-related parameter and investigates its evolution under clean, noisy, reduced-supervision, and causal monitoring conditions. A closely related study applied PINNs to inverse modelling of a flexible motor–load system by embedding the coupled torsional dynamics directly into the neural-network loss function (Ivanović, M. et al. 2026).

MOTOR–LOAD TORSIONAL MODEL

The considered drivetrain is represented by an equivalent two-inertia motor–load system connected through a flexible coupling (Ivanović, M. et al. 2026). The model consists of the motor inertia ($\mathrm{J_m}$), load inertia ($\mathrm{J_l}$), viscous torsional damping ($\mathrm{b_v}$), and time-varying torsional stiffness ($\mathrm{k(t)}$). The motor is excited by the electromagnetic torque ($\mathrm{M_{em}(t)}$), while no independent external load torque is considered in the present simulation study. This reduced-order representation captures the dominant torsional dynamics of the flexible drivetrain while retaining a sufficiently simple structure for inverse parameter identification.

The motor- and load-side angular positions and angular velocities are denoted by ($\mathrm{\theta_m(t)}$), ($\mathrm{\theta_l(t)}$), ($\mathrm{\omega_m(t)}$), and ($\mathrm{\omega_l(t)}$), respectively. The equations of motion are

$$\dot{\theta}_m(t) = \omega_m(t),]M_{em}(t) b_v[\omega_m(t) - \omega_l(t)]k(t)[\theta_m(t) - \theta_l(t)], \tag{1}$$

$$\dot{\theta}_l(t) = \omega_l(t),]b_v[\omega_m(t) - \omega_l(t)] + k(t)[\theta_m(t) - \theta_l(t)]. \tag{2}$$

The equal and opposite coupling torque transmitted between the two inertias is therefore

$$b_v[\omega_m(t) - \omega_l(t)] + k(t)[\theta_m(t) - \theta_l(t)]. \tag{3}$$

To formulate the inverse problem in variables that directly describe the deformation of the flexible coupling, the relative angular displacement and relative angular velocity are introduced as

$$\delta(t) = \theta_m(t) - \theta_l(t), v_\delta(t) = \omega_m(t) - \omega_l(t). \tag{4}$$

Consequently,

$$\dot{\delta}(t) = v_\delta(t). \tag{5}$$

Subtracting the motor- and load-side acceleration equations yields the first-order relative dynamics

$$\dot{v}_\delta(t) + b_v \Gamma v_\delta(t) + k(t) \Gamma \delta(t) - \frac{M_{em}(t)}{J_m} = 0. \tag{6}$$

Where

$$\Gamma = \frac{1}{J_m} + \frac{1}{J_l}. \tag{7}$$

This relative-state formulation is particularly suitable for torsional-stiffness identification because the unknown stiffness enters the dynamics directly through the coupling term $(\mathrm{k(t)\delta(t)})$. The quantities $(\delta(\mathrm{t}))$ and $(\mathrm{v_\delta(t)})$ therefore contain the information most directly associated with deformation and relative motion of the flexible transmission.

The nominal physical parameters used in the simulations are summarized in Table 1. The healthy and degraded stiffness values define the reference degradation range used for performance evaluation. Importantly, the reference time-varying stiffness trajectory is not supplied to the PINN as a training target and is used only after optimization for quantitative assessment of the estimated stiffness.

Table 1 Physical parameters of the two-inertia motor–load model

Parameter	Symbol	Value	Unit	Role
Motor inertia	(J_m)	6.2×10^{-4}	kg·m²	Known model parameter
Load inertia	(J_l)	2.2×10^{-4}	kg·m²	Known model parameter
Viscous torsional damping	(b_v)	0.004	N·m·s/rad	Known model parameter
Healthy stiffness	(k_0)	350	N·m/rad	Evaluation reference
Degraded stiffness	(k_f)	245	N·m/rad	Evaluation reference

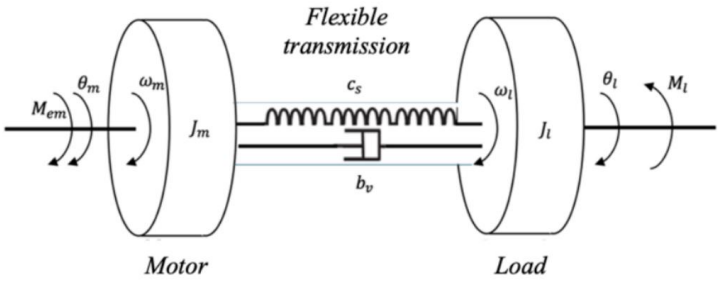


Figure 1 Equivalent two-inertia motor–load system with flexible coupling. In the present formulation, the coupling stiffness $\boldsymbol{c_s}$ shown in the conventional schematic is generalized to the time-varying stiffness $\boldsymbol{k(t)}$, while the independent load torque $\boldsymbol{M_l}$ is set to zero. The motor is excited by the electromagnetic torque $\boldsymbol{Mem(t)}$, and the motor- and load-side angular velocities $\boldsymbol{\omega_m(t)}$ and $\boldsymbol{\omega_l(t)}$ represent the encoder-accessible quantities used in the proposed digital twin.

PROPOSED WEAK FIRST-ORDER PHYSICS-INFORMED MODEL

The proposed digital twin combines a first-order neural representation of the relative torsional states with a bounded model of the time-varying stiffness. The architecture is designed to avoid the numerical sensitivity associated with second-order neural derivatives while retaining direct physical dependence between the estimated stiffness and the deformation of the flexible coupling.

Relative-state neural representation

The RelativeStateNet predicts two separately normalized outputs, $\hat{\delta}(\mathrm{t}), \hat{\mathrm{v}}_\delta(\mathrm{t})$ representing the relative angular displacement and relative angular velocity, respectively.

The relative angular velocity available from the encoder signals is defined as

$$v_{\delta,\mathrm{meas}}(t) = \omega_m(t) - \omega_l(t). \tag{8}$$

Since direct relative-angle measurements are not available, an encoder-derived angular quantity is obtained by cumulative numerical integration,

$$\delta_{\mathrm{int}}(t) = \int_0^t v_{\delta,\mathrm{meas}}\,(\tau)\,d\tau. \tag{9}$$

with $\delta(0) = 0$.

Importantly, δ_{int} does not represent an additional physical sensor measurement. It is derived entirely from the available motor- and load-side encoder-speed information.

The state network uses a Fourier representation of time to improve its ability to reproduce the oscillatory torsional dynamics. During state pretraining, the network is constrained using the available relative-state supervision together with the first-order kinematic consistency condition

$$\frac{d\hat{\delta}}{dt} - \hat{v}_{\delta} = 0. \tag{10}$$

Time-varying stiffness representation

After state pretraining, the unknown torsional stiffness is represented by a bounded monotone sigmoid function,

$$\hat{k}(t) = k_{\text{low}} + \frac{k_{\text{high}} - k_{\text{low}}}{1 + exp(wt - t_c)}. \tag{11}$$

where $\mathrm{k_{high}}$ and $\mathrm{k_{low}}$ denote the upper and lower stiffness levels, $\mathrm{t_c}$ determines the centre of the degradation transition, and w controls its temporal width.

The parameterization constrains the stiffness estimate to a physically meaningful range while describing the progressive degradation scenario considered in this study. The initialization of the stiffness parameters is independent of the reference degradation trajectory, and the true stiffness is not used as a training target.

Weak physics-informed residuals

Instead of enforcing the relative dynamics through pointwise second-order derivatives, the proposed formulation uses weak integral residuals over overlapping time windows $[\mathrm{t_a}, \mathrm{t_b}]$.

The weak kinematic residual is

$$r_{\text{kin}} = \hat{\delta}(t_b) - \hat{\delta}(t_a) - \int_{t_a}^{t_b} \hat{v}_{\delta}\,(t)\,dt. \tag{12}$$

Defining

$$\Gamma = \frac{1}{J_m} + \frac{1}{J_l}. \tag{7}$$

the weak dynamic residual becomes

$$= \hat{v}_{\delta}(t_b) - \hat{v}_{\delta}(t_a) + b_v \Gamma \int_{t_a}^{t_b} \hat{v}_{\delta}\,(t)\,dt + \Gamma \int_{t_a}^{t_b} \hat{k}\,(t)\hat{\delta}(t)\,dt - \frac{1}{J_m}\int_{t_a}^{t_b} M_{\text{em}}\,(t)\,dt. \tag{13}$$

All integrals are evaluated using trapezoidal quadrature on the actual temporal coordinates.

The main advantage of this formulation is that the dynamic constraint does not require direct evaluation of $\frac{\mathrm{d}^2\hat{\delta}}{\mathrm{dt}^2}$.

or equivalently high-order derivatives of the neural approximation. This reduces the sensitivity of the inverse stiffness-identification problem to local differentiation errors.

Training objective

The complete objective function combines four components,

$$L_{\text{total}} = \lambda_{\text{data}} L_{\text{data}} + \lambda_{\text{kin}} L_{\text{kin}} + \lambda_{\text{dyn}} L_{\text{dyn}} + \lambda_{\text{IC}} L_{\text{IC}}. \tag{14}$$

Here,

- $\mathrm{L_{data}}$ provides supervision from encoder-derived relative quantities;
- $\mathrm{L_{kin}}$ enforces first-order kinematic consistency;
- $\mathrm{L_{dyn}}$ imposes the weak integral form of the motor–load dynamics;
- $\mathrm{L_{IC}}$ constrains the initial condition.

For the constant-stiffness verification experiments, the pretrained state representation is fixed while a single stiffness parameter is identified. For the main time-varying experiment, the bounded stiffness parameters are

optimized using the weak physics-informed objective. In the reduced-supervision configuration, sensor-labelled observations are reduced while the dense unlabeled physics-collocation grid is retained.

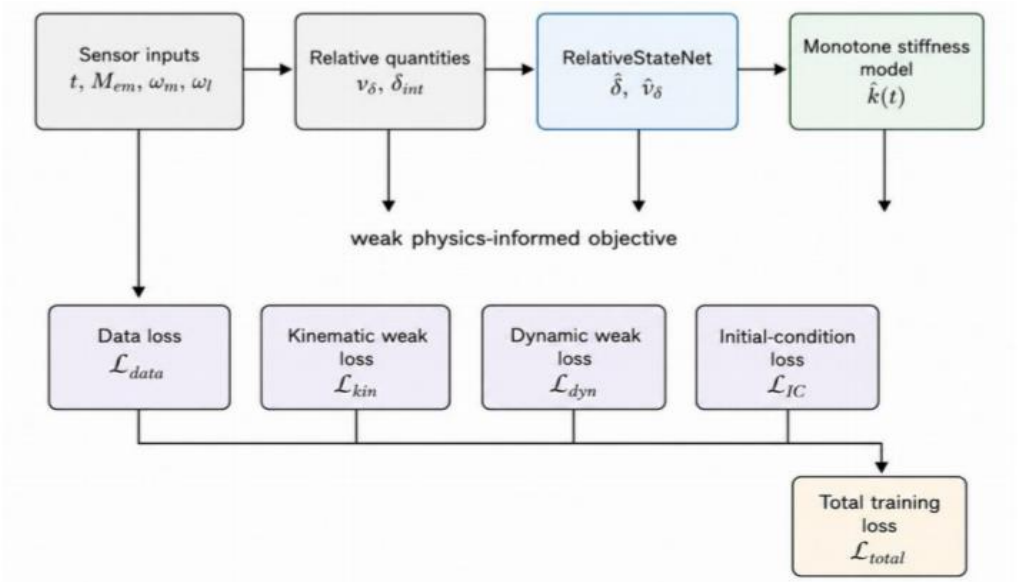


Figure 2 Proposed weak first-order PINN architecture for torsional-stiffness estimation. Motor- and load-side encoder information is transformed into relative quantities and processed by the RelativeStateNet to estimate $\widehat{\boldsymbol{\delta}}(\boldsymbol{t})$ and $\widehat{\boldsymbol{v}}_{\boldsymbol{\delta}}(\boldsymbol{t})$. The time-varying stiffness $\widehat{\boldsymbol{k}}(\boldsymbol{t})$ is represented by a bounded monotone model, while sensor supervision, weak kinematic and dynamic residuals, and the initial-condition constraint are combined in the total training objective.

SIMULATION AND EVALUATION SETUP

The proposed weak first-order physics-informed digital twin is evaluated through a sequence of simulation experiments designed to assess inverse stiffness-identification accuracy, robustness to measurement noise, performance under reduced sensor supervision, and computational feasibility for causal online monitoring. The evaluation begins with controlled constant-stiffness cases and is then extended to a time-varying degradation scenario, noisy encoder-derived observations, and progressively reduced temporal sampling. A second-order physics-informed formulation is retained as a baseline for comparison. All experiments use the same two-inertia motor–load model and physical parameters introduced in MOTOR–LOAD TORSIONAL MODEL, while the reference stiffness trajectory is used exclusively for post-training performance assessment and is never supplied to the neural network as a training target. The following subsections describe the data and sampling configuration, validation scenarios, evaluation metrics, and constant-stiffness verification procedure.

Data and sampling configuration

The analysed time interval extends from (0) to (0.75) s and contains 1501 uniformly spaced samples, corresponding to a sampling frequency of (2) kHz and a sampling interval of (0.5) ms. The electromagnetic-torque profile $(\mathrm{M_{em}(t)})$ and the time vector are provided by the available jera1.mat dataset. Since experimental motor- and load-side encoder trajectories are not available, the corresponding angular-speed responses are generated using a high-accuracy reference ODE model driven by the supplied torque profile.

The weak physics-informed formulation is evaluated over overlapping temporal windows containing 101 samples. Each window therefore spans approximately (50) ms. Consecutive windows are shifted by 25 samples, corresponding to a stride of (12.5) ms. This configuration retains several oscillatory cycles within each weak integration interval while providing substantial overlap between neighbouring windows.

Validation scenarios

The proposed inverse formulation is evaluated through a sequence of increasingly demanding validation scenarios. First, three constant-stiffness control cases are considered, $[\mathrm{k} = 350, 300, 245\ \mathrm{N, m/rad,}]$ representing healthy, intermediate, and degraded mechanical conditions. For each case, a RelativeStateNet is pretrained using the corresponding simulated state response and subsequently fixed while the scalar stiffness parameter is optimized. All stiffness-identification runs start from the same initial estimate, $[\mathrm{k_{init}} = 288.75\ \mathrm{N, m/rad,}]$ thereby providing a consistent comparison between the three operating conditions. The main experiment considers a continuously

varying monotone stiffness degradation. The same formulation is subsequently evaluated with (0.3%) noise introduced into the encoder-derived velocity supervision.

A reduced-supervision experiment is also performed using 751 uniformly distributed sensor-labelled samples over the complete (0)–(0.75) s interval. In this case, the physics-informed constraints are still evaluated at all 1501 time points. The additional physics-collocation points do not contain true stiffness, relative-angle, or relative-speed labels; they provide only the time coordinates, known torque excitation, and governing physical equations. The term reduced sensor supervision refers specifically to the reduction of labelled encoder-derived observations; it does not imply a reduction of the physics-collocation grid.

Finally, strongly reduced sampling configurations containing 121 and 401 sensor-labelled points are considered to investigate the limits imposed by insufficient temporal resolution.

Evaluation criteria

Stiffness-estimation performance is quantified using the root-mean-square error (RMSE), relative RMSE, coefficient of determination (R2), mean absolute error (MAE), maximum absolute error, and the initial and final stiffness errors.

For the constant-stiffness verification cases, relative identification error is used as the primary metric. For the time-varying experiments, both trajectory-level error and (R2) are reported in order to evaluate not only the overall magnitude of the error but also the ability of the estimator to reproduce the temporal evolution of degradation.

State-reconstruction quality is evaluated separately from stiffness-identification quality. This distinction prevents an accurate stiffness estimate from obscuring inadequate reconstruction of the auxiliary relative states.

Model selection is performed exclusively using training-objective information. The reference stiffness trajectory is not used for checkpoint selection, early stopping, or optimization and is introduced only after training for quantitative evaluation.

Constant-stiffness verification

Before addressing time-varying degradation, the inverse formulation is verified using the three constant-stiffness cases summarized in Table 2.

Table 2 Constant-stiffness inverse-identification results

Reference stiffness, (k) [N·m/rad]	Initial estimate [N·m/rad]	Estimated stiffness [N·m/rad]	Absolute error [N·m/rad]	Relative error [%]
350	288.75	351.21	1.208	0.345
300	288.75	300.24	0.240	0.080
245	288.75	245.23	0.234	0.096

Despite the common initialization, the estimated stiffness converges close to the corresponding reference value in all three cases. The largest relative error is approximately (0.35%), observed for the healthy $(350\ \mathrm{N,m/rad})$ condition, while the intermediate and degraded cases remain below (0.10%). These results provide a controlled verification of the inverse objective before the stiffness model is extended to the time-varying degradation scenario.

Table 3 Data provenance in the simulation study

Quantity	Source	Used as
(t)	jera1.mat	Time coordinate
$(M_{em}(t))$	jera1.mat	Known excitation / physics input
$THref$	jera1.mat	Available, not used in stiffness loss
$(\omega_m(t))$	Reference ODE	Synthetic encoder response
$(\omega_l(t))$	Reference ODE	Synthetic encoder response
$(v_\delta(t))$	Derived from encoders	State supervision
$(\delta_{\mathrm{int}}(t))$	Integrated encoder difference	Derived supervision
$(k_{\mathrm{ref}}(t))$	Simulation reference	Evaluation only

RESULTS AND DISCUSSION

The proposed weak first-order physics-informed formulation was evaluated with respect to three central objectives: accurate reconstruction of the time-varying torsional stiffness, robustness to measurement noise and reduced sensor supervision, and improvement over the conventional second-order inverse formulation. The results show that the proposed approach consistently follows the imposed stiffness degradation while retaining physically meaningful estimates under all successful validation scenarios.

Time-Varying Torsional-Stiffness Identification

Figure 3 compares the reference time-varying torsional stiffness with the estimates obtained using the proposed weak first-order formulation under three different sensor conditions and with the conventional second-order baseline. The reference trajectory represents a monotone degradation from the healthy stiffness region toward the degraded operating condition and is used exclusively for post-training evaluation.

The full-rate clean estimate closely follows the reference stiffness trajectory over the complete analysed interval. The proposed estimator achieves a relative stiffness RMSE of (2.8717%) and a coefficient of determination of $(\mathrm{R}_{\mathrm{k}}^2 = 0.95245)$. Although a small initial overestimation is present, the estimated trajectory accurately reproduces the principal degradation transition and converges to approximately $(244.3\ \mathrm{N,m/rad})$, which is very close to the degraded reference value of $(245\ \mathrm{N,m/rad})$.

The 0.3% noise case exhibits behaviour very similar to the clean experiment. The corresponding relative stiffness RMSE is (2.9032%), while $(\mathrm{R}_{\mathrm{k}}^2 = 0.9514)$. The close overlap between the clean and noisy curves indicates that the proposed weak integral formulation retains its stiffness-identification capability under the investigated level of encoder-speed perturbation. This behaviour is consistent with the use of temporally integrated physical constraints, which reduce the sensitivity of the inverse estimation process to small local disturbances in the sensor-derived signals.

The reduced-supervision estimate, shown for the configuration with 751 sensor-labelled instants and 1501 unlabeled physics-collocation points, also reproduces the main degradation trend. In this case, the relative stiffness RMSE increases to (3.7483) and the coefficient of determination decreases to $(\mathrm{R}_{\mathrm{k}}^2 = 0.91898)$. The largest deviation is visible primarily around the onset and timing of the stiffness transition. Nevertheless, the overall shape and final degraded stiffness remain well represented. Since the number of sensor-labelled observations is reduced by approximately (49.97%), this result demonstrates that the physical constraints can partially compensate for reduced measurement supervision when they are enforced on a sufficiently dense collocation grid.

The distinction between sensor-labelled observations and physics-collocation points is important. The additional collocation points do not contain reference values of stiffness or hidden relative states and therefore do not represent additional sensor measurements. They provide only the known time coordinates, electromagnetic-torque excitation, and locations at which the governing equations are enforced.

The second-order baseline, in contrast, follows a substantially different trajectory. Its relative stiffness RMSE is approximately (10.87), with $(\mathrm{R}_{\mathrm{k}}^2 \approx 0.319)$. Although the baseline retains a smooth decreasing trend, it fails to reproduce both the timing and magnitude of the reference degradation and remains significantly above the degraded stiffness level at the end of the interval.

This comparison provides an important motivation for the proposed formulation. In the second-order inverse approach, small errors in the neural approximation can be amplified by higher-order differentiation and subsequently compensated through an incorrect stiffness estimate. The first-order weak formulation avoids direct evaluation of second derivatives and instead imposes the governing dynamics through integral constraints over temporal windows. The resulting inverse problem therefore exhibits substantially improved stiffness-identification performance in the investigated system.

Overall, Figure 3 demonstrates three central properties of the proposed digital twin: accurate reconstruction of the time-varying stiffness trajectory under clean conditions, robustness to moderate encoder-speed noise, and preservation of useful identification accuracy under substantially reduced sensor supervision. The simultaneous comparison with the second-order baseline further indicates that the improvement is associated with the proposed first-order weak physics-informed formulation rather than merely with the flexibility of the neural-network approximation.

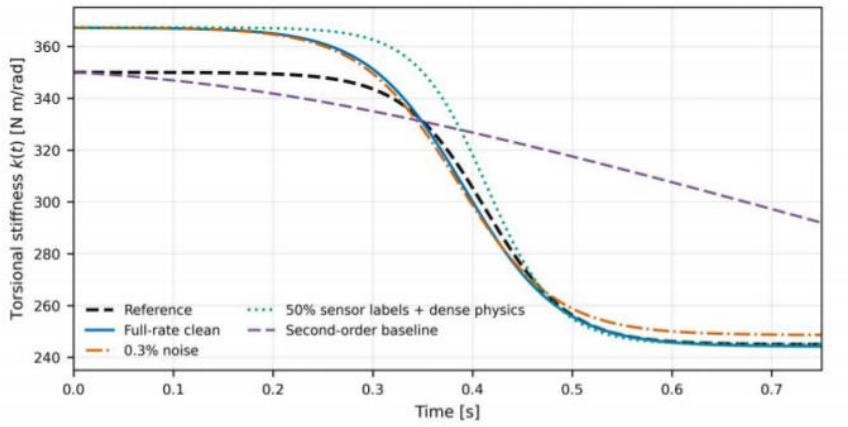


Figure 3 Time-varying torsional-stiffness identification under different validation conditions. The reference degradation trajectory is compared with the proposed weak first-order estimates obtained using full-rate clean data, 0.3% encoder-speed noise, and 751 sensor-labelled samples with dense unlabeled physics collocation, together with the conventional second-order baseline.

Relative-state reconstruction was additionally evaluated as a physical-consistency check for the proposed inverse model. In the full-rate clean experiment, the RelativeStateNet achieved $(\mathrm{R}^2 = 0.9702)$ for the relative angular displacement $(\delta(\mathrm{t}))$ and $(\mathrm{R}^2 = 0.9982)$ for the relative angular velocity $(\mathrm{v}_\delta(\mathrm{t}))$. These results indicate that the accurate stiffness estimate is accompanied by a physically consistent reconstruction of the underlying torsional dynamics, rather than being obtained solely through compensation within the stiffness parameter. Under reduced sensor supervision, the relative-angle reconstruction remained acceptable, whereas the relative-speed reconstruction degraded more noticeably, showing that accurate identification of the lower-dimensional stiffness trajectory does not necessarily imply equally accurate reconstruction of all auxiliary states.

Reduced Sensor Supervision and Sampling Limitations

An important objective of the proposed physics-informed digital twin is to reduce the amount of sensor-labelled information required for stiffness identification while preserving the physical resolution needed to describe the dominant torsional dynamics. For this purpose, three reduced-sampling configurations containing 121, 401, and 751 sensor-labelled time instants were investigated. Figure 4 compares their influence on the estimation of the time-varying torsional stiffness.

In the most strongly reduced configuration, only 121 uniformly distributed samples are retained over the complete (0)–(0.75) s interval. This corresponds to an effective sampling frequency of approximately $\mathrm{f_s} = 160\ \mathrm{Hz}$.

Since the dominant torsional frequency of the considered system is approximately $\mathrm{f_{tor}} = 230.5\ \mathrm{Hz}$, this sampling rate is insufficient to satisfy the classical Nyquist condition. The configuration provides only approximately 0.69 samples per dominant torsional period. Consequently, the oscillatory dynamics cannot be represented without aliasing, and the resulting stiffness estimate fails to reproduce the reference degradation trajectory.

Increasing the number of uniformly distributed samples to 401 raises the effective sampling frequency to approximately $\mathrm{f_s} = 533.3\ \mathrm{Hz}$, with a corresponding Nyquist frequency of approximately. $\mathrm{f_N} = 266.7\ \mathrm{Hz}$. The dominant torsional component is therefore formally below the Nyquist limit. Nevertheless, only approximately 2.31 samples are available per dominant oscillation period. Although this temporal resolution is sufficient for basic spectral separation, it remains inadequate for accurate reconstruction of the oscillatory waveform and for reliable numerical evaluation of the weak integral constraints. The corresponding stiffness trajectory consequently remains substantially biased.

A different behaviour is obtained when 751 sensor-labelled samples are retained. This configuration corresponds to approximately 4.34 samples per dominant torsional period. More importantly, only the amount of sensor supervision is reduced: the weak physical constraints continue to be evaluated on the complete grid of 1501 unlabeled physics-collocation points. Thus, the number of sensor-labelled observations is reduced from 1501 to 751, corresponding to a reduction of approximately 49.97%.Despite this reduction, the estimated stiffness retains good agreement with the reference degradation trajectory. The relative stiffness RMSE is $\mathrm{R_k^2} = 0.91898$. his result demonstrates that the number of sensor-labelled observations can be substantially reduced without a proportional loss of stiffness-identification accuracy, provided that the governing dynamics continue to be enforced on a sufficiently dense temporal grid.

A key distinction must therefore be made between sensor-labelled samples and physics-collocation points. The additional collocation points do not contain reference values of the stiffness (k(t)), relative angular displacement ($\delta(t)$), or relative angular velocity ($v_\delta(t)$). They do not represent additional physical measurements. Instead, they provide only the time coordinates, the known electromagnetic-torque excitation, and locations at which the governing equations are enforced. Their role is to constrain the inverse solution through the physical model rather than through additional labelled data.

The results in Figure 4 also show that satisfying the classical Nyquist criterion alone is not sufficient for reliable physics-informed inverse identification. The 401-sample configuration formally satisfies the Nyquist condition but still fails to provide sufficient temporal resolution for accurate reconstruction and weak-form integration. Reliable identification therefore depends not only on spectral admissibility but also on the number of samples available within each dominant torsional period and on the accuracy with which the integral physics constraints can be evaluated.

Overall, the comparison indicates that reducing both the measurement and physics grids simultaneously leads to a substantial deterioration of the inverse solution, whereas reducing only the sensor-labelled supervision while retaining dense unlabeled physics collocation provides a considerably more effective strategy. This finding is particularly relevant for practical condition-monitoring applications, where labelled sensor information may be limited while the governing physical model remains available.

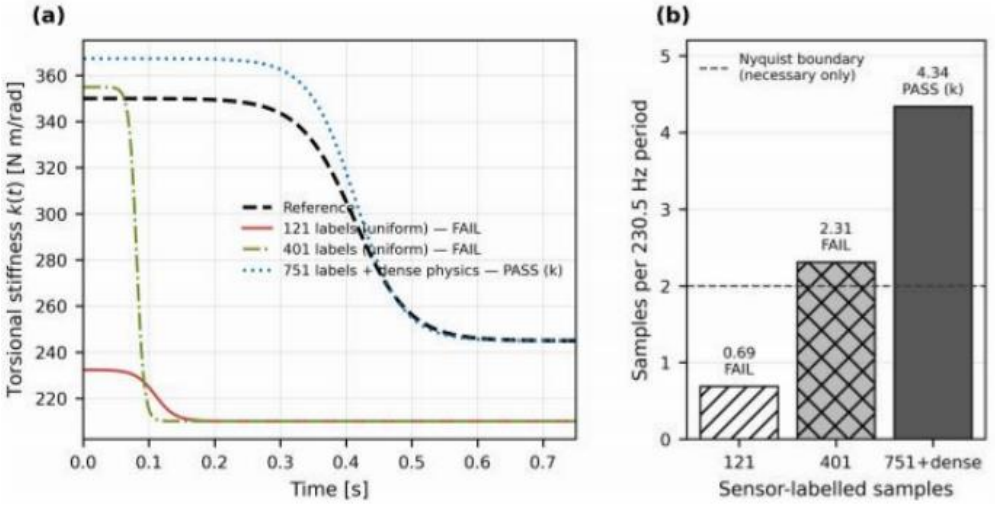


Figure 4 Influence of sensor-sampling density on time-varying torsional-stiffness identification. (a) Reference stiffness trajectory and estimates obtained using 121, 401, and 751 sensor-labelled samples, with the complete 1501-point unlabeled physics-collocation grid retained for the 751-label configuration. (b) Number of samples per dominant torsional period for the investigated sampling configurations. The comparison demonstrates that satisfaction of the classical Nyquist condition alone does not guarantee reliable inverse reconstruction.

Causal Online Degradation Monitoring

An addition to offline stiffness identification, the proposed digital twin was evaluated in a causal online configuration to investigate its suitability for condition-monitoring applications. At each update instant, the estimator uses only information available at the current and previous time instants; no future samples are included in the estimation window. The RelativeStateNet remains fixed after offline pretraining, while the parameters of the bounded stiffness model are updated within each causal window. The resulting procedure should therefore be interpreted as a warm-started causal estimator, rather than as a cold-start online learning algorithm.

Figure 5(a) compares the reference stiffness trajectory with the corresponding causal online estimate. A degradation threshold of $k_{th} = 315\ \mathrm{N,m/rad}$ is introduced as an illustrative condition-monitoring limit. The estimated stiffness crosses this threshold at approximately $t_{det} = 0.300\ s$, whereas the reference trajectory crosses the same threshold at approximately $t_{ref} = 0.3839\ s$. The corresponding threshold-crossing offset is therefore approximately $t_{det} - t_{ref} - 83.9\ ms$.The negative value indicates an early threshold crossing rather than a detection delay. This result demonstrates that the estimator responds to the developing degradation before the reference trajectory reaches the selected threshold. However, the early crossing should not be interpreted as an exact estimate of fault-onset time, since the online stiffness trajectory remains less accurate than the corresponding offline estimate. For the selected five-step online configuration, the relative stiffness RMSE is 6.7267%, with $R_k^2 = 0.7394$. The accuracy is lower than in the offline full-rate experiment, which is expected because each causal update has access only to past and present information and is restricted to a small number of optimization steps. Nevertheless, the estimated trajectory preserves the overall degradation trend and provides a usable condition-monitoring signal.

Computational performance was evaluated independently for three online configurations using 3, 4, and 5 Adam optimization steps per update. Each configuration was repeated ten times, resulting in 290 measured update latencies per configuration. The five-step configuration was selected because it provided the lowest stiffness-estimation error while still remaining well below the (25) ms computational target in the final repeated benchmark.

For the selected configuration, the measured mean update latency was $2.5410\ \mathrm{ms}$, with a 95th percentile of $4.0309\ \mathrm{ms}$, a 99th percentile of $6.4649\ \mathrm{ms}$, and a maximum measured latency of $6.8162\ \mathrm{ms}$. No deadline misses were observed among the 290 measured updates in the final repeated test. Figure 5(b) shows the individual measured latencies together with the (25) ms computational target.

These results demonstrate that the proposed estimator can perform causal stiffness updates at near-real-time rates on the tested CPU while maintaining useful degradation-tracking capability. The benchmark does not constitute hard-real-time certification or a formal worst-case execution-time guarantee, since operating-system scheduling and hardware conditions may influence individual update times. The result is therefore interpreted as evidence of causal near-real-time computational feasibility rather than deterministic real-time performance.

Overall, the online experiment complements the offline identification results by showing that the proposed physics-informed digital twin can provide sequential estimates of torsional stiffness without access to future measurements. This capability represents an important step toward future integration with condition-monitoring, predictive-maintenance, and adaptive-control systems.

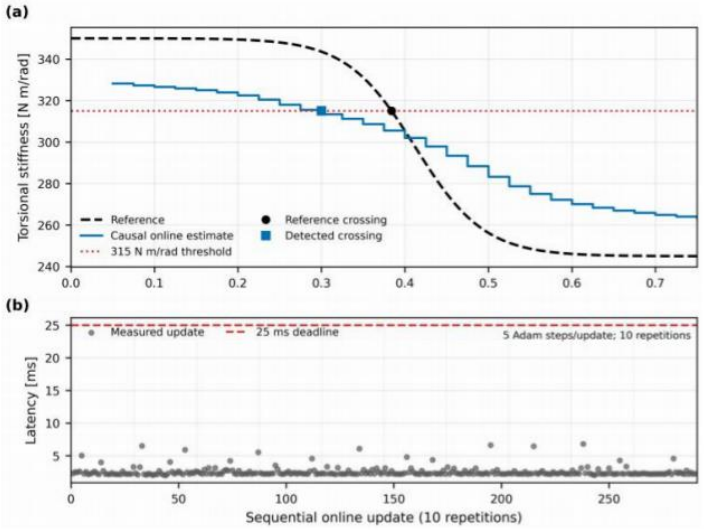


Figure 5 Causal online torsional-stiffness monitoring and computational performance. (a) Reference and warm-started causal stiffness estimates together with the illustrative 315 Nm/rad degradation threshold and corresponding threshold crossings. (b) Measured update latencies for the selected five-Adam-step configuration over ten repeated runs (290 updates), compared with the 25 ms computational target.

Table 4. Repeated causal online benchmark for different numbers of optimization steps per update

Adam steps/update	Mean [ms]	P95 [ms]	P99 [ms]	Max [ms]	Deadline misses	Relative (k)-RMSE [%]
3	1.9662	3.7804	4.8888	5.8089	0/290	6.8711
4	2.5400	4.6835	6.1140	9.0540	0/290	6.8105
5	**2.5410**	**4.0309**	**6.4649**	**6.8162**	**0/290**	**6.7267**

LIMITATIONS AND FUTURE WORK

The present study provides a simulation-based validation of the proposed physics-informed digital twin. The time vector and electromagnetic-torque profile are obtained from the available jera1.mat dataset, whereas the motor- and load-side encoder-speed responses are generated using the reference ODE model. Consequently, the current results do not yet account for effects that may arise in a physical drivetrain, including encoder quantization, sensor bias, temperature variation, backlash, friction nonlinearities, structural uncertainties, and unmodelled disturbances. The present formulation also neglects an independently varying external load torque. Future studies should investigate joint or robust identification under unknown load disturbances in order to distinguish stiffness degradation from variations in external mechanical loading.

The physical parameters (J_m), (J_l), and (b_v) are assumed to be known, while the stiffness evolution is represented by a bounded monotone sigmoid function. This parameterization is suitable for the progressive degradation scenario considered in this work, but it cannot describe stiffness recovery, intermittent behaviour, multiple abrupt faults, or more complex non-monotone degradation trajectories. Future work should therefore investigate more flexible stiffness representations and quantify the sensitivity of the inverse estimate to uncertainty in the assumed mechanical parameters.

Although the reduced-supervision experiment demonstrates that the number of sensor-labelled observations can be decreased by approximately (49.97%), the dense physics-collocation grid is retained. The result should therefore be interpreted as a reduction in measurement supervision rather than a reduction in the number of physics evaluations. Further research could investigate adaptive collocation strategies in which physics points are concentrated in dynamically informative regions, potentially reducing the computational burden without sacrificing identification accuracy.

The online estimator is causal within each temporal window, but it uses a RelativeStateNet pretrained offline. A fully adaptive cold-start implementation using only an incoming sensor stream has not yet been evaluated. Future developments should therefore investigate continual or online state adaptation, embedded implementation, and hardware-in-the-loop validation. Particular attention should be given to deterministic scheduling and worst-case execution-time analysis before the method is considered for hard-real-time deployment.

Another important extension concerns the transition from condition monitoring to active control and prognostics. The estimated stiffness ($k(t)$) could be used as an online scheduling or adaptation variable for vibration-suppression, feedforward, model-predictive, or adaptive-control strategies. Similarly, the estimated degradation trajectory could provide an input to higher-level prognostic models for fault-probability assessment and remaining-useful-life estimation.

Finally, experimental validation using measured motor- and load-side encoder signals represents the most important next step. A laboratory motor–load test bench with controlled stiffness variation would enable direct evaluation of the method under realistic noise, modelling uncertainty, disturbances, and repeatability conditions. Multi-seed training studies and uncertainty-aware confidence intervals should also be included in future validation in order to quantify the statistical robustness of the identified stiffness trajectory. As a longer-term research direction, hybrid quantum–classical learning architectures may also be investigated for physics-constrained inverse parameter identification, provided that their computational or representational advantages over the present classical PINN formulation can be demonstrated quantitatively.

CONCLUSION

This paper presented a weak first-order physics-informed digital twin for inverse estimation of time-varying torsional stiffness in a flexible motor–load system. By reconstructing relative torsional states and enforcing the governing dynamics through weak integral residuals, the proposed formulation avoids direct dependence on second-order neural derivatives and improves the robustness of the inverse stiffness-identification problem.

The proposed method was first verified using constant-stiffness control cases and subsequently evaluated under time-varying clean, noisy, and reduced-supervision conditions. In the clean full-rate experiment, the relative stiffness RMSE was (2.8717) with ($R_k^2 = 0.95245$), while the (0.3%) noise experiment retained a similar accuracy with a relative RMSE of (2.9032%). When the number of sensor-labelled observations was reduced from 1501 to 751 while preserving the dense 1501-point physics-collocation grid, the relative stiffness RMSE remained (3.7483%) with ($R_k^2 = 0.91898$). The strongly reduced 121- and 401-sample configurations further demonstrated that satisfaction of the classical Nyquist criterion alone is not sufficient for reliable physics-informed inverse reconstruction.

A causal warm-started online implementation additionally demonstrated the feasibility of sequential stiffness monitoring on the tested CPU. The selected five-step configuration achieved a relative online stiffness RMSE of (6.7267%), while all 290 measured updates in the final repeated benchmark remained below the (25) ms computational target.

Overall, the results demonstrate that the proposed weak first-order physics-informed formulation can accurately identify a hidden time-varying torsional stiffness while remaining robust to moderate measurement noise and substantially reduced sensor supervision. The developed framework therefore provides a promising methodological basis for future experimental digital twins aimed at drivetrain condition monitoring, predictive maintenance, and adaptive vibration-control applications.

ACKNOWLEDGMENTS

This research was supported by the Science Fund of the Republic of Serbia, Grant No. 214, within the project "Parametric and Nonparametric Methods for Data-Driven Optimal Control of Mechatronic Systems with Flexible Coupling - DDOC".

REFERENCES

Hua, X., & Gandee, E. (2021). Vibration and dynamics analysis of electric vehicle drivetrains. *Journal of Low Frequency Noise,* Vibration and Active Control, 40(3), 1241–1251. https://doi.org/10.1177/1461348420979204

Ivanović, M. (2023). *Neuronske mreže podržane fizičkim zakonima – Praktikum.* Prirodno-matematički fakultet, Univerzitet u Kragujevcu. https://scidar.kg.ac.rs/bitstream/123456789/18602/5/pinn-skripta-lat.pdf

Ivanović, M., Matijević, M., Krstić, L., & Jeremić, A. (2026). I*nverse modeling of flexible rotational systems via hybrid physics-informed and data-driven learning*

Jeremić, M., Krstić, L., Ivanović, M., Lazarević, M., & Matijević, M. (2026). *Physics-informed digital twin for torsional stiffness estimation* [Computer software]. GitHub. https://github.com/poni0107/pinn-torsional-stiffness-digital-twin

Li, W., Chai, Z., Wang, M., Hu, X., & Guo, Y. (2016). *Online identification and verification of the elastic coupling torsional stiffness. Shock and Vibration, 2016,* Article 2016432. https://doi.org/10.1155/2016/2016432

Park, Y.-H., Lee, H.-B., & Kim, G.-W. (2023). Crack monitoring in rotating shaft using rotational speed sensor-based torsional stiffness estimation with adaptive extended Kalman filters. *Sensors, 23(5)*, 2437. https://doi.org/10.3390/s23052437