Check for updates

Earthquake Engineering & Structural Dynamics

WILEY

RESEARCH ARTICLE OPEN ACCESS

Learning Rocking Dynamics From Sparse Shake-Table Data With Interpretable Physics-Informed Neural Networks

Shirley Shen | Christian Málaga-Chuquitaype

Department of Civil and Environmental Engineering, Imperial College London, London, UK

Correspondence: Christian Málaga-Chuquitaype (c.malaga@imperial.ac.uk)

Received: 29 January 2026 | **Revised:** 21 June 2026 | **Accepted:** 23 June 2026

Keywords: computational structural dynamics | experimental data | explainable AI | physics-informed neural networks | rocking structures | seismic response prediction

ABSTRACT

We present a hybrid interpretable Physics-Informed Neural Network Long-Short Term Memory (Hybrid PINN LSTM) framework for predicting the seismic response of rocking blocks. Existing analytical models rely on uncertain idealizations, while purely data-driven and machine-learning approaches lack physical consistency and interpretability. Our framework addresses these limitations by integrating the governing physics of rocking motion with a limited number of shake-table data. Taking the ground acceleration series as input, our model predicts the rotational response of rocking blocks and its associated peak quantities with lower error and higher accuracy than classical rocking theories, while predictions of angular velocity and angular acceleration remain more challenging. Physical consistency is enforced through mechanics-based loss terms, while data-driven losses are calibrated using a small set of laboratory measurements obtained with "off-the-shelf" experimental equipment. Besides, to enhance engineering insight, explainable AI tools at both the surrogate and sequence levels are used to interpret the learned relationships and identify the structural and ground-motion parameters, and excitation time windows governing the rocking response. Our explainability analysis shows that for high-intensity earthquakes, velocity-based measures and related displacement measures dominate overturning behavior, whereas for moderate intensities, the response becomes more interactive, with temporal characteristics, such as duration, playing a comparable role. For non-overturning cases, large safe rocking amplitudes are governed jointly by velocity, displacement, dominant period, and duration of the ground motion. Overall, this work pioneers an accurate, physically consistent and interpretable tool for seismic rocking response prediction. It could be used for assessing the seismic performance of critical equipment on the basis of limited experimental data without relying on current idealistic assumptions on key parameters like the coefficient of restitution.

1 | Introduction

Estimating the seismic response of free-standing rocking structures represents one of the classical yet most challenging nonlinear problems in earthquake engineering. Since the seminal work of Housner [1], rocking systems have been recognized as inherently non-smooth dynamical systems governed by alternating phases of uplift, pivoting about successive corners, occasional sliding or free-flight rotation, and abrupt impacts accompanied by energy loss. These mechanisms introduce strong nonlinearities arising from geometric effects, discontinuous changes in stiffness, and lead to high sensitivity to excitation characteristics. As a result, the seismic response of rocking structures can vary dramatically even under seemingly similar ground motions, making its predictive modelling an enduring research challenge.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.
© 2026 The Author(s). *Earthquake Engineering & Structural Dynamics* published by John Wiley & Sons Ltd.

Building on Housner's original formulation [1], several analytical and semi-analytical investigations have deepened the understanding of rocking dynamics [2–4]. The dynamically equivalent rocking oscillator proposed by DeJong and Dimitrakopoulos [5] clarified how slenderness, mass distribution, and the frequency parameter collectively govern small-amplitude rocking oscillations. Subsequent studies have incorporated additional complexities, including interface friction and sliding [6], impact restitution and contact modelling [7], geometric imperfections and surface roughness, and multi-impact phenomena such as chatter [8].

Shaking table experiments have further revealed that rocking response often has limited repeatability [9–11]. And more recent datasets [11, 12] have demonstrated that this variability persists across different block geometries and excitation levels, highlighting the importance of response models that explicitly account for experimental uncertainty, noise, and guard against unjustified idealizations. Classical deterministic or purely analytical rocking models, although physically insightful, struggle to reproduce these experimentally observed variability.

In parallel with advances in experimental capabilities, the rapid growth of global strong-motion databases [13] and the emergence of extensive ML-compatible datasets have stimulated increasing interest in data-driven methods for predicting seismic response. Although neural networks have been used to predict seismic responses for a long time [14], to date, only a few machine learning (ML) models have been deployed to classify rocking initiation and overturning, predict peak rotation, and identify governing intensity measures (IMs) [15, 16]. Interpretable ML frameworks [17–19] have shown that seismic rocking response is strongly influenced by nonlinear combinations of velocity-, duration-, and frequency-related IMs. However, conventional ML methods, including support vector machines, feed-forward neural networks [20], random forests, and generic surrogate models [21]; remain fundamentally black-box. As a result, their extrapolation abilities are limited, their interpretability remains partial, and their predictions can become unreliable.

The emergence of physics-informed neural networks (PINNs) [22, 23] offers a promising avenue for bridging data-driven learning with the governing physical laws. PINNs incorporate differential equations as soft constraints within the loss function, hence guiding the network towards physically consistent solutions even in the presence of sparse training data. In structural dynamics, extensions such as PhyLSTM [24] and physics-guided convolutional networks [25] have demonstrated superior extrapolation performance relative to purely data-driven networks. The recently proposed constitutive-model-constrained PINNs (CM–PINNs) [26] further highlight the strength of embedding constitutive laws and adaptive loss weighting, for enhanced stability, accuracy, and physical interpretability in nonlinear structural systems. To our knowledge, the first and only application of PINNs to rocking structures was done by Shen and Málaga-Chuquitaype [15]. They showed, on the basis of numerical simulations only, that embedding rigid-body mechanics into deep learning models, supported by clustering analysis and limited simulation data, can effectively predict non-overturning rocking responses while mitigating overfitting and reducing data requirements.

However, applying PINNs to realistic rocking dynamics remains highly challenging. Unlike smooth dynamics, rocking is characterized by unilateral contact, sign-dependent response, and impact events involving instantaneous velocity changes. The state-dependent governing equations of rocking dynamics are difficult to incorporate directly into standard PINN formulations. Moreover, available experimental data exhibits measurement noise, non-repeatability, and subtle imperfections that challenge purely physics-based models. Existing PINNs also generally assume smoothness and differentiability that do not hold in impact-driven rocking systems.

In this context, recurrent neural network (RNN) architectures, particularly long short-term memory (LSTM) networks [27–30], appear well suited to capture long-range temporal dependencies and the multi-stage temporal nature of rocking response. However, in their conventional form, LSTMs do not enforce physical constraints, conservation principles, or the peculiarities of rigid-body motion. Thus, while they can model complex time-series patterns, they lack physical consistency and robustness outside the training range. Likewise, explainable AI (XAI) has emerged as a complementary tool, providing insights into the behavior and decision-making processes of ML models. SHAP, LIME, and partial dependence plots have been used in earthquake engineering to identify dominant IMs [31], visualize nonlinear interactions [32], and uncover causal mechanisms in ML-based response prediction [16, 33]. In rocking analysis, XAI offers the ability to connect learned model behavior with known mechanics.

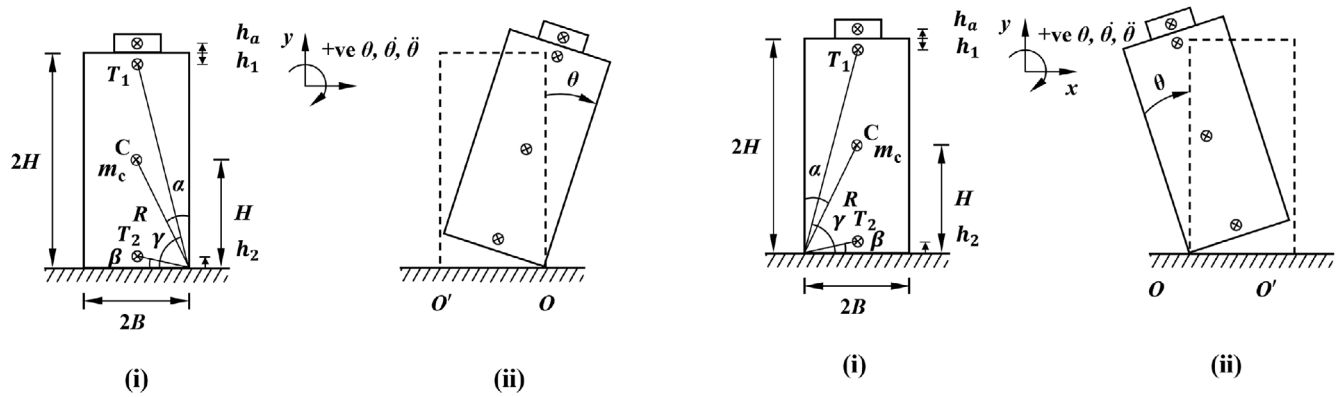
Motivated by these opportunities and limitations, the present study develops a novel Hybrid PINN framework for predicting the seismic rocking response of free-standing rigid blocks. The framework pioneers the integration of: (i) Experimental acceleration-rotation sequences acquired from shaking-table tests, (ii) Physics-based residuals derived from the full rigid-body equations of motion and energy-loss, enforcing conservation principles, (iii) A deep FC-LSTM architecture capable of modelling long-range temporal dependencies across uplift, impact, and post-impact phases; and (iv) XAI-based diagnostics to interpret the learned relationships and validate the proposed framework's alignment with established mechanics.

To our knowledge, this approach is the first to simultaneously enforce the governing physical laws, leverage limited experimental data (in the order of a few tens rather than the several thousand required routinely for training neural networks), and preserve interpretability; all of which are not jointly achieved by any of the existing analytical, numerical, or purely data-driven models on their own. The new framework is validated against unseen experimental data, and its performance is assessed in terms of predictive accuracy, generalization to ground motions beyond the training dataset, robustness to measurement noise, and adherence to physical consistency.

2 | Dynamics of Free Rigid Rocking Bodies

This section summarizes the basics of rigid body rocking motion setting up the nomenclature and parameters used in subsequent parts of the paper. First, let us consider a rigid block of mass m, width $2B$, and height $2H$, as illustrated schematically in Figure 1.

10969845, 2026, 12, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/eqe.70239 by Imperial College London, Wiley Online Library on [12/10/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License



(a) Case 1. **(b)** Case 2.

FIGURE 1 | Schematic representation of a rigid rocking block subjected to base excitation under two configurations of motion, Case 1 and Case 2, rocking to the right and left, respectively.

The block is characterized by a size parameter R and a slenderness angle α. For simplicity, it is assumed that the coefficient of friction between the block and the supporting surface is sufficiently large to prevent sliding, so that the motion is restricted to uplift and rotation about the block's lower corners, an assumption that has been corroborated by the experimental data generated in the following section. The size parameter R is defined in Equation 1, and the slenderness about the potential pivot points O and O' is defined in Equation 2.

$$R = \sqrt{H^2 + B^2} \tag{1}$$

$$\alpha = a\tan\left(\frac{B}{H}\right) \tag{2}$$

where, H is the location of the center of gravity above the base of the block and B is the distance of the center of gravity from the side for the 2-dimensional block, a is the angle the block semi-diagonal made with the vertical when the block is at rest (in [rad]), and the angle θ is measured to quantify the block's inclination from the vertical.

When subjected to a horizontal ground acceleration $a_g(t)$, the block remains in full contact with the ground until the inertial overturning moment overcomes the stabilizing moment generated by gravity [34]. $\theta(t)$ is taken as positive when rocking about the right corner and negative when rocking about the left one. The block has a moment of inertia about the pivot equal to: $I_0 = \frac{4}{3}mR^2$.

Following Housner's seminal work [1], and later generalizations [35], the rocking response consists of two smooth phases depending on the sign of $\theta(t)$:

$$I_0\ddot{\theta}(t) + mgR\sin[-\alpha - \theta(t)] = -mR\,a_g(t)\cos[-\alpha - \theta(t)], \quad \theta(t) < 0, \tag{3}$$

$$I_0\ddot{\theta}(t) + mgR\sin[\alpha - \theta(t)] = -mR\,a_g(t)\cos[\alpha - \theta(t)], \quad \theta(t) > 0. \tag{4}$$

These expressions represent the rotational dynamic equilibrium between gravitational restoring forces, inertial forces induced by the ground motion, and the rotational inertia of the body.

Normalizing Equations (3)–(4) by I_0 yields a compact sign-dependent equation of motion:

$$\ddot{\theta}(t) + p^2\sin[\alpha\,\mathrm{sgn}(\theta(t)) - \theta(t)] = \frac{a_g(t)}{\Gamma}\cos[\alpha\,\mathrm{sgn}(\theta(t)) - \theta(t)], \tag{5}$$

where:

$$p = \sqrt{\frac{3g}{4R}}, \qquad \Gamma = \frac{I_0}{mR} = \frac{4}{3}R, \tag{6}$$

A key feature of rocking dynamics is the occurrence of impacts when the block transitions from one pivot corner to the other. These impacts introduce energy dissipation that is typically represented analytically by a coefficient of restitution e, defined through:

$$\dot{\theta}_{\text{after}} = e\,\dot{\theta}_{\text{before}}, \qquad 0 < e < 1. \tag{7}$$

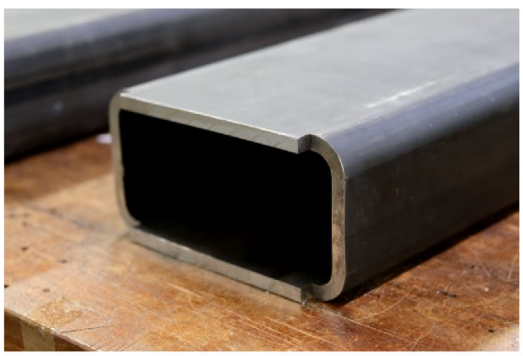
Although Housner's model provides a closed-form expression for e based solely on the block's geometry, real rocking systems exhibit additional energy losses related to material deformation, surface roughness, and micro-slip.

To integrate the rocking dynamics into a differentiable deep-learning framework, it is necessary to regularize the non-differentiable sign function in Equation (5). This is achieved by approximating $\mathrm{sgn}(\theta)$ by $\mathrm{sgn}(\theta) \approx \tanh(\kappa\theta)$ with $\kappa = 30$. The smoothing parameter (κ) controls the sharpness of the regularized sign function used in the physics-informed residual. A larger value gives a sharper approximation of the discontinuous switching behavior associated with rocking, while a smaller value gives a smoother but less abrupt transition [36]. The selected value of 30 was used after a preliminary sensitivity analysis and represents a practical compromise between preserving the switching behavior near zero rotation and maintaining sufficient smoothness for stable gradient-based training [37]. Very small values may over-smooth the transition and weaken the representation of the rocking switch, whereas much larger values may

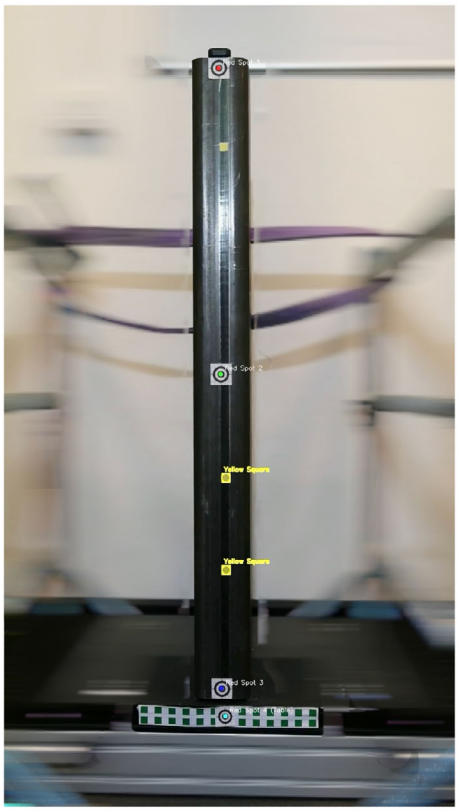
10969845, 2026, 12, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/eqe.70239 by Imperial College London, Wiley Online Library on [12/10/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License



(a) View of some of the specimens used as rocking blocks.



(b) Close-up view of the specimen cross-section and footing.



(c) Full view of a specimen (Specimen 4) and its corresponding optical markers for tracking.

FIGURE 2 | Experimental setup.

make the residual numerically stiff and more difficult to optimize. Therefore:

$$\phi(t) = \alpha \tanh(\kappa\theta(t)) - \theta(t), \tag{8}$$

and the governing equation of motion becomes:

$$\ddot{\theta}(t) + p^2 \sin\phi(t) = \frac{a_g(t)}{\Gamma}\cos\phi(t), \tag{9}$$

which retains the essential nonlinear behavior of rocking while enabling backpropagation through the physics constraint. The differentiable form in Equation (9) forms the basis of the physics-informed loss used in the Hybrid PINN model described later in this paper.

3 | Experimental Input and Dataset Generation

The experimental database employed in this study was developed from a series of shake-table tests on free-standing steel rocking blocks, conducted at the Emerging Structural Technologies Research Group at Imperial College London. To this end, we used a novel hybrid sensing approach, fusing off-the-shelf standard-resolution camera video tracking with low-cost high-frequency accelerometers to resolve the full kinematic state vector $\{\theta, \dot{\theta}, \ddot{\theta}\}$ which proved to have sufficient precision to calculate residuals for physics-informed loss functions like the one trialed in this paper.

The experiments included six free-standing steel columns, specifically dimensioned to investigate a broad range of slenderness ratios and size effects, which are known to govern the stochasticity of the rocking response. The specimens were fabricated from hot rolled 7.5 mm thick hollow steel box sections made of S275 structural steel. To mitigate out-of-plane behavior, the specimens were designed with a specific width-to-depth aspect ratio of 2:1, effectively creating a preferred direction of motion and an energetic barrier against twisting. By aligning the excitation strictly with the weaker axis, out-of-plane torsional coupling was minimized; post-test video analysis confirmed that transverse rotation remained negligible, thereby justifying the use of the planar rigid body assumption. All specimens were milled to form 10 mm feet at their bottom edges (Figure 2). Furthermore, to favor the rocking mode and rigorously inhibit basal sliding, high-friction sandpaper was applied to the interface between the steel base and the shaking table platen.

Table 1 summarizes the geometric and mass properties of the specimens tested. The characteristic length R ranged from approximately 120–710 mm, corresponding to block slenderness angles (α in Figure 2a) between 0.06 and 0.35 rad. Specimen masses varied from 5.4 to 35.7 kg. Post-test 3D surface scans revealed negligible plate warping induced by welding, with out-of-plane deviations of the order of a few tenths of a mm; however, these imperfections introduced only minor out-of-plane motion, and the rocking response remained essentially planar. Figure 2 illustrates a typical test configuration, showing the hollow steel block, the optical markers

10969845, 2026, 12, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/eqe.70239 by Imperial College London, Wiley Online Library on [12/10/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

TABLE 1 | Geometric and mass properties of the rigid block specimens.

ID	R [mm]	α [rad]	B [mm]	D [mm]	H [mm]	M [kg]
1	118.79	0.348	160.53	81.03	219.61	5.43
2	156.90	0.261	160.56	80.86	299.46	7.46
3	180.61	0.22	160.34	80.79	348.42	8.71
4	461.04	0.088	160.29	80.83	912.07	22.89
5	669.74	0.060	160.35	80.81	1331.67	33.39
7	712.96	0.057	160.46	80.87	1420.90	35.73

and the calibrated ruler that defined the world coordinate system.

All tests were conducted on a low-cost servo-hydraulic single-axis shaking table with plan dimensions of 870 × 500 × 130 mm. The input ground motions were derived from recorded earthquake accelerograms available through the NGA PEER ground motion database [38]. Given the table's maximum displacement capacity of 10 cm, each accelerogram was preprocessed through a rigorous two-stage correction pipeline. First, low-order polynomial trends were removed to eliminate long-period drifts. Second, an FFT-based, Butterworth-equivalent high-pass filter of order 2 was applied. The cutoff frequencies were iteratively adjusted to satisfy the dual constraints of PGA $\geq$ 0.5g and PGD $\leq$ 0.10 m, ensuring the tests maximized excitation intensity without exceeding the physical limits of the shaking table.

As detailed later in this paper, K-means clustering, with K = 4, was applied to the full set of records to identify distinct excitation regimes. An 80%–20% stratified split was then performed within each cluster, resulting in around 42 training and 10 test records per specimen. This low number of tests is advantageous and stands in stark contrast to the several tens of thousands usually required for training neural networks.

As previously mentioned, the block kinematics were captured using a hybrid measurement system combining a standard camera video tracking system with low-cost inertial sensing. Each test configuration included at least three circular red markers: two attached to the top and bottom faces of the column and one mounted on the shaking-table (Figure 2). The motion was recorded by a standard fixed-position digital camera aligned with the plane of rocking, operating at 60 frames per second. The field of view was calibrated with a diving ruler and a yellow checkered scale, yielding a precise pixel-to-metric conversion factor of 20 pixels/cm. The global coordinate origin was defined at the bottom-left corner of the ruler and remained fixed across all tests. A custom motion-tracking algorithm based on dynamic color thresholding, morphological filtering and circular feature detection was developed in Python to automatically identify the centroids of all markers in each frame. The resulting pixel coordinates were transformed into physical coordinates in the world reference frame, and subsequently rotated into a local reference frame aligned with the specimen geometry. This allowed the direct computation of the rocking angle $\theta(t)$, the instantaneous rotation center, and the ground displacement history $x_g(t)$ from the marker trajectories.

TABLE 2 | Specifications of the low-cost WT9011DCL-BT50 IMU sensors.

Parameter	Specification	Note
Output frequency	0.2–200 Hz	Configurable sampling rate
Acceleration range	±16 g	3-axis MEMS accelerometer
Angular velocity	±2000°/s	3-axis MEMS gyroscope
Resolution	0.5 mg/LSB	16-bit internal ADC
Accuracy	0.01 g, 0.2°/s	After calibration

In parallel, inertial measurements were obtained using WT9011DCL-BT50 AHRS IMU sensors [39]. Two units were mounted on each specimen, one near the top plate and one near the base, while a third sensor was attached to the table surface to capture the input acceleration history. The sensors communicated via Bluetooth and provided tri-axial acceleration, angular velocity and orientation at a sampling frequency of 50 Hz. Key specifications of the IMUs are summarized in Table 2. The raw IMU acceleration data were post-processed using baseline correction and digital filtering to suppress low-frequency drift and high-frequency noise, respectively. All IMUs operated with synchronized timestamps so that the inertial measurements of the specimen could be temporally aligned with the optical data. The inertial signals provided independent acceleration and angular velocity histories that complemented the optical measurements and allowed consistency checks across all recorded channels.

The recorded videos and IMU time histories were processed into structured numerical arrays suitable for use in the Hybrid-PINN framework. The data processing pipeline is schematically illustrated in Figure 3. Optical marker trajectories were extracted using a custom-made algorithm based on dynamic color thresholding and circular feature detection, which successfully identified four distinct markers per frame. Coordinate transformations were applied to align the raw pixel data with the physical model; this included a rotational correction (typically < 0.2°) to align the camera axes with the specimen's vertical axis. From these corrected trajectories, the rocking angle $\theta(t)$ and ground displacement $x_g(t)$ were computed, and the table-mounted IMU provided an independent estimate of the ground acceleration $\ddot{x}_g(t)$.

The resulting time series were smoothed using a Savitzky–Golay moving-average filter with edge handling to suppress optical noise while preserving the essential rocking dynamics. Numerical differentiation with a consistent time step Δt was used to obtain velocity and acceleration histories, including $\dot{x}_g(t)$ and $\ddot{x}_g(t)$, as well as $\dot{\theta}(t)$ and $\ddot{\theta}(t)$. All signals were subsequently resampled to a uniform time step of $\Delta t = 0.01$ s, ensuring precise phase alignment between $\theta(t)$ and $\ddot{x}_g(t)$. Each test was stored as a structured time series containing the variables $\{t, \theta(t), \dot{\theta}(t), \ddot{\theta}(t), \ddot{x}_g(t)\}$, together with metadata such as specimen ID, test index, peak ground acceleration (PGA), and excitation duration. What's more, all quantities were normalized via z-score transformation using the statistics of the training subset to ensure numerical stability and balanced gradient magnitudes.

10969845, 2026, 12, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/eqe.70239 by Imperial College London, Wiley Online Library on [12/10/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

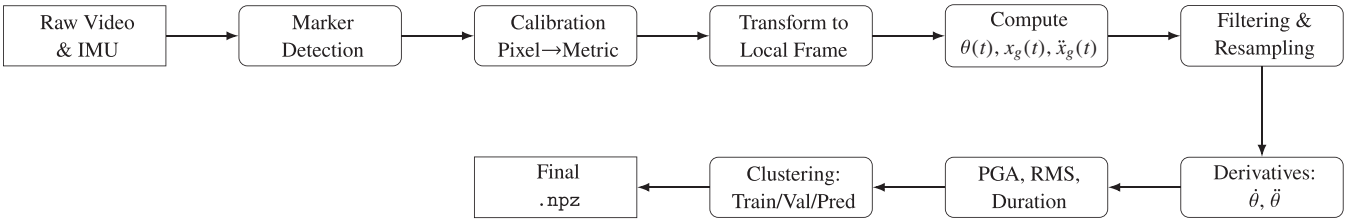


FIGURE 3 | Data-processing pipeline from raw video and IMU recordings to structured experimental datasets.

4 | Physics-Informed Deep LSTM Metamodel

Previous sections have stressed the inherently nonlinear and non-smooth response that characterizes rocking motion. Transitions between contact points induce discontinuous changes in the governing equations and angular velocity, while the amplitude-dependent natural frequency and impact-induced energy loss produce response regimes that vary markedly with excitation intensity, rendering the input-output mapping highly non-linear and non-scalable. Consequently, purely data-driven sequencing models without explicit physical structure tend to misrepresent sustained rocking oscillations, struggle to capture regime transitions, and exhibit poor generalization when evaluated outside the intensity range observed during training.

To address these limitations, we propose a physics-informed recurrent learning framework that couples a deep LSTM architecture with the equation of rocking motion embedded as a differentiable constraint. This hybridization leverages the temporal representation capability of LSTMs while enforcing consistency with the underlying mechanics through physics-based residuals. This section summarizes the nonlinear governing equations, outlines the core LSTM formulation, and details the proposed Hybrid PINN-LSTM architecture, including its physics-derived loss functions and curriculum-based training strategy designed to enhance stability and generalization across intensity levels.

4.1 | Deep LSTM Network

The modelling of highly nonlinear dynamical systems, particularly those exhibiting non-smooth, path-dependent behavior such as the rocking response of rigid blocks, necessitates a computational architecture capable of capturing complex temporal dependencies and history-dependent state transitions. While classical feedforward neural networks can approximate universal functions, they lack the internal memory mechanisms required to process sequential data where the current state is functionally dependent on the entire preceding history. Recurrent Neural Networks (RNNs) address this by introducing cyclic connections, allowing information to persist. However, classical RNNs suffer severely from vanishing and exploding gradient phenomena during the training phase. This instability arises during Backpropagation Through Time (BPTT), where error gradients are propagated backward through the unrolled network. Mathematically speaking, if the spectral radius of the recurrent weight matrix deviates even slightly from unity, the gradients will either decay exponentially (vanishing) or grow exponentially to infinity (exploding) as the sequence length increases. This prevents standard RNNs from learning dependencies that span time horizons longer than a few dozen steps, rendering them unsuitable for rocking dynamics.

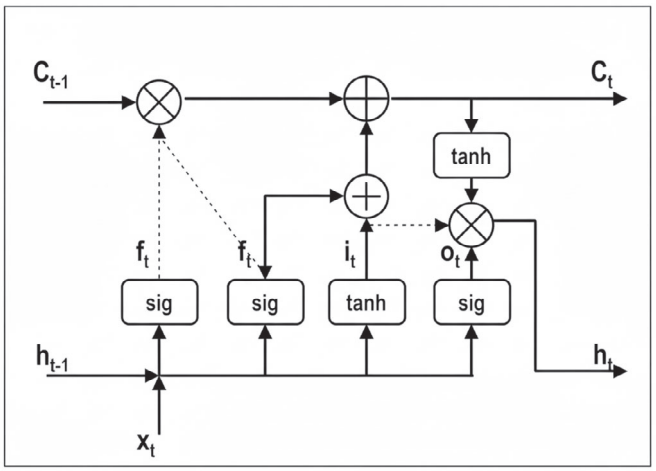
To overcome these fundamental limitations, we employ LSTM networks. First introduced by Hochreiter and Schmidhuber [40], their specialized recurrent architecture is explicitly designed to maintain constant error flow, overcoming the vanishing- and exploding-gradient phenomena mentioned before [41], enabling them to learn long-term temporal dependencies. Subsequent refinements by Gers et al. [42, 43] introduced forget gates and peephole connections, enhancing the ability of LSTMs to discard irrelevant historical information while selectively incorporating new temporal patterns. These innovations made LSTMs the key architecture for sequence modelling in a wide variety of domains, including speech recognition [44] and natural language processing [45]. Within the structural engineering and applied mechanics fields, recurrent neural networks have been widely applied to modelling material hysteresis, forecasting structural response time histories, and representing systems governed by complex path-dependent behavior [24, 46]. The capacity of recurrent networks to capture long-range temporal dependencies is essential for rocking systems with significant non-smoothness and long-memory characteristics that shallow feedforward architectures are fundamentally unable to represent.

The core innovation of the LSTM is the introduction of a memory cell, c_t, which acts as a conveyor belt for information, running straight down the entire chain with only minor linear interactions, see Figure 4. This cell state is regulated by a sophisticated gating mechanism comprising a forget gate, an input gate, and an output gate. These gates serve as differentiable regulators that learn to selectively retain or discard information. In mathematical terms, the operation of a single LSTM cell at time step t is governed by a specific set of equations that regulates the flow of information based on the current input vector $x_t \in \mathbb{R}^d$ and the previous hidden state $h_{t-1} \in \mathbb{R}^h$. The first component is the forget gate f_t, which determines what information from the previous cell state c_{t-1} becomes irrelevant given the new context. For a rocking block, this is analogous to "forgetting" the pre-impact velocity once an impact has occurred and energy has been dissipated. The forget gate computes a value between 0 and 1 using a sigmoid activation function σ:

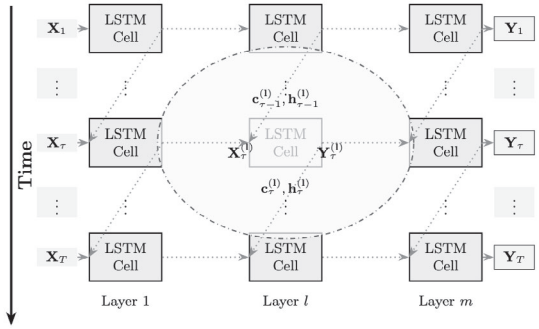
$$f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f) \tag{10}$$

where W_f and b_f are the learnable weight matrix and bias vector, respectively. A value of 1 represents total retention, while 0 represents complete forgetting. Simultaneously, the network must decide what new information to store. This is handled by the input gate i_t and a candidate cell state generator $\tilde{c}_t$. The input

10969845, 2026, 12, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/eqe.70239 by Imperial College London, Wiley Online Library on [12/10/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License



(a) Traditional internal structure of an LSTM cell with gates and state flows. (Adapted from[47]).



(b) Proposed deep LSTM network architecture unrolled in time, showing multiple stacked layers (Layer 1 through Layer m) processing input sequence $\mathbf{X}_1$ to $\mathbf{X}_T$ to produce outputs $\mathbf{Y}_1$ to $\mathbf{Y}_T$. The vertical flow is the passing of hidden layer output $\mathbf{Y}_\tau^{(l)}$ as the input $\mathbf{X}_\tau^{(l+1)}$ to the next layer. The time direction is handled by the recurrence, where the hidden and cell states $(\mathbf{h}_\tau^{(l)}, \mathbf{c}_\tau^{(l)})$ are passed horizontally to the same cell in the next time step..

FIGURE 4 | Traditional LSTM internal cell structure and Deep LSTM network architecture with multiple stacked layers used in this study.

gate decides which values will be updated, while the candidate layer creates a vector of new candidate values using a hyperbolic tangent (tanh) activation function:

$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i), \quad \tilde{c}_t = \tanh(W_c \cdot [h_{t-1}, x_t] + b_c) \quad (11)$$

Each step involves combining these gates to update the cell state. The old cell state c_{t-1} is multiplied, element-wise, by f_t, and the new values $\tilde{c}_t$ are scaled by i_t. This additive update structure is the key theoretical advantage of LSTMs. The gradient of the sum includes a term that is simply the forget gate vector, the gradient can flow backwards through time without implicit multiplication by a weight matrix that would cause decay. Therefore, the update equation is:

$$c_t = f_t \odot c_{t-1} + i_t \odot \tilde{c}_t \quad (12)$$

where, $\odot$ denotes the Hadamard (element-wise) product. Finally, the output gate o_t controls what part of the cell state is exposed as the output hidden state, h_t. The cell state is passed through a tanh to push values between -1 and 1, and then multiplied by the output gate:

$$o_t = \sigma(W_o \cdot [h_{t-1}, x_t] + b_o), \quad h_t = o_t \odot \tanh(c_t) \quad (13)$$

The Deep LSTM network is constructed by stacking multiple LSTM layers (Figure 4). In this hierarchical architecture, the hidden state output $h_t^{(l)}$ of layer l serves as the input $x_t^{(l+1)}$ to layer $l+1$. This depth enables the model to learn representations at varying levels of abstraction. Lower layers typically capture high-frequency fluctuations, noise characteristics, and immediate impact responses, effectively acting as a learning filter. Deeper layers, having a larger effective receptive field, encode slower, long-term oscillatory modes, global energy decay trends, and the cumulative effects of energy dissipation. This multi-scale feature extraction is particularly critical for rocking blocks, where the response is a complex superposition of rapid impacts (high frequency) and slow, free-vibration rocking phases (low frequency). The network parameters $\theta_{LSTM} = \{W_f, U_f, b_f, \ldots\}$ are optimized using the Adam algorithm, an adaptive learning rate method that combines the advantages of AdaGrad and RMSProp, ensuring stable convergence even in the presence of the sparse gradients associated with impact events. Recent physics-informed LSTM studies [24] have demonstrated that embedding physical constraints further enhances the ability of deep LSTMs to generalize across varying excitation intensities and ground-motion characteristics.

4.2 | Hybrid PINN-LSTM Architecture

Our Hybrid PINN-LSTM architecture is developed to learn the nonlinear mapping from the horizontal ground acceleration series, $a_g(t)$, to the resulting rocking motion, while remaining rigorously constrained by the governing rocking dynamics. Unlike a purely data-driven recurrent model which does not enforce mechanics, frequently violating basic energetic principles leading to predictions that gain or lose energy in a physically inconsistent manner, our model incorporates the dynamics of the system and benefits from real data. Similarly, in contrast with yielding or hysteretic structural systems, where the response depends on latent internal variables such as plastic strains, slip, degradation, or damage; for rigid-body rocking, no such internal variables exist. The dynamics of rocking are fully determined by the observable state $(\theta, \dot{\theta})$ together with discrete contact conditions such as lift-off, pivot switching, and impact. There are no hidden material states or internal evolution laws beyond these directly observable quantities. Therefore, introducing additional hidden states cannot be interpreted as adding meaningful physical variables, instead this would only introduce unnecessary architectural complexity, increasing the risk of overfitting without improving predictive capability.

For this reason, the proposed Hybrid PINN architecture consists of a unified deep LSTM network with an interconnected physics module to enforce dynamic admissibility. By constraining the recurrent network with the governing mechanics, the proposed

10969845, 2026, 12, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/eqe.70239 by Imperial College London, Wiley Online Library on [12/10/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Hybrid PINN-LSTM avoids drifting predictions, respects energetic consistency, and generalizes more reliably across different excitation intensities.

The input ground motion is first normalized and lifted to a higher-dimensional feature space through a SiLU-activated fully connected layer. The SiLU activation ($x \cdot \sigma(x)$) provides the smoothness necessary for higher-order derivative calculations in the physics loss, avoiding the "dead neuron" issues of ReLU in physics-informed contexts. The encoded features are then processed by a stacked LSTM, which produces a nonlinear temporal hidden representation H_n at each time step. This backbone provides a nonlinear temporal representation that captures uplift, free rotation, and impact-induced state transitions.

A multi-head linear projection, augmented by a residual skip connection from the input acceleration, maps the hidden state H_n to three standardized output predictions at each time step: standardized angle $\theta_{\text{std}}(t_n)$, velocity$\dot{\theta}(t_n)$, and directly predicted angular acceleration $\ddot{\theta}_{\text{pred}}(t_n)$. The residual skip connection preserves the near-linear dependence of the angular acceleration on $a_g(t)$ away from impacts. This shared representation ensures kinematic coherence across all output quantities and avoids the autoregressive error accumulation typical of single-output sequence models.

To enforce physical consistency, the raw network outputs are subjected to a series of differentiable operators. Kinematic consistency is imposed by passing the predicted angle through centered finite-difference (FD) operators to obtain reference angular velocity $\dot{\theta}_{\text{FD}}$ and angular acceleration $\ddot{\theta}_{\text{FD}}$ fields. The physical angle is recovered through de-standardization using dataset statistics. These serve as kinematic targets during training, penalizing discrepancies between the network's direct predictions (from the multi-head decoder) and the derivatives of its angle prediction.

We propose a training objective formed by a weighted combination of data fidelity, kinematic consistency, physics enforcement, and regularization terms, which are optimized via a staged learning curriculum. The total loss function $\mathcal{L}_{\text{total}}$ is defined as:

$$\mathcal{L}_{\text{total}} = \alpha\mathcal{L}_{\text{data}} + \beta_1\mathcal{L}_{\text{pv}} + \beta_2\mathcal{L}_{\text{dd}} + \gamma\mathcal{L}_{\text{PE}} + \delta\mathcal{L}_{\text{int}} + \epsilon\mathcal{L}_{\text{reg}} \tag{14}$$

where α, β_1, β_2, γ, δ, and ϵ are time-varying weight coefficients that follow the learning curriculum schedule. The primary data-driven component of the loss function is the data loss ($\mathcal{L}_{\text{data}}$), which enforces fidelity to the ground-truth rocking angle $\theta_{\text{std,true}}(t_n)$:

$$\mathcal{L}_{\text{data}} = \sum_{n=1}^{N_t} m_n \left\| \theta_{\text{std}}(t_n) - \theta_{\text{std,true}}(t_n) \right\|_2^2 \tag{15}$$

A blended acceleration, combining finite-difference and network-predicted accelerations with an epoch-dependent blending factor, stabilizes the physics constraints optimization during early training. The dynamics are enforced through a differentiable residual based on the smoothed pivot-switching equation of motion, while an integration constraint reconstructs rotation from the predicted velocities to suppress long-term drifting. Spectral and temporal gradient regularization further mitigates spurious oscillatory artefacts and strengthens physical plausibility. To stabilize optimization in the early epochs, the physics module introduces a blended acceleration, written as:

$$\ddot{\theta}_{\text{blend}}(t_n) = (1-\lambda)\,\ddot{\theta}_{\text{FD}}(t_n) + \lambda\,\ddot{\theta}_{\text{pred}}(t_n) \tag{16}$$

where λ increases gradually with the epoch index. Substituting $\ddot{\theta}_{\text{blend}}$ into the smoothed rocking equation of motion yields the residual:

$$r(t_n) = \ddot{\theta}_{\text{blend}}(t_n) + p^2 \sin\phi(t_n) + \frac{a_g(t_n)}{\Gamma}\cos\phi(t_n) \tag{17}$$

where the term $\phi(t)$ has been defined in Equation 8. The corresponding physics loss is defined in variational form as:

$$\mathcal{L}_{\text{PE}} = \sum_{n=1}^{N_t} m_{2,n} \left\| r(t_n) \right\|_2^2, \tag{18}$$

with $m_{2,n}$ masking boundary and invalid points. The angle-regression loss, $\mathcal{L}_{\text{data}}$, has been defined before; while the kinematic consistency losses enforce kinematic coherence by penalizing discrepancies between the network's direct predictions and the FD-derived reference values. The velocity kinematic loss ($\mathcal{L}_{\text{pv}}$) is:

$$\mathcal{L}_{\text{pv}} = \sum_{n} m_{1,n} \left\| \dot{\theta}(t_n) - \dot{\theta}_{\text{FD}}(t_n) \right\|_2^2 \tag{19}$$

and the acceleration kinematic loss ($\mathcal{L}_{\text{dd}}$) is:

$$\mathcal{L}_{\text{dd}} = \sum_{n} m_{2,n} \left\| \ddot{\theta}_{\text{pred}}(t_n) - \ddot{\theta}_{\text{FD}}(t_n) \right\|_2^2 \tag{20}$$

A further integration consistency loss ($\mathcal{L}_{\text{int}}$) is applied, penalizing the difference between the network-predicted angle of rotation $\theta(t_n)$ and the reconstructed angle (θ_{rec}) via numerical integration of the predicted velocities:

$$\theta_{\text{rec}}(t_n) = \theta(t_0) + \sum_{k=1}^{n} \dot{\theta}(t_k)\Delta t, \qquad \mathcal{L}_{\text{int}} = \sum_{n} m_n \left\| \theta_{\text{rec}}(t_n) - \theta(t_n) \right\|_2^2. \tag{21}$$

To enhance smoothness and frequency fidelity, the model incorporates a Regularization Loss term ($\mathcal{L}_{\text{reg}}$), a spectral loss comparing the amplitude spectra of predicted and true standardized rotations, together with a temporal-gradient loss penalizing discrepancies in first-order discrete derivatives. All components are combined into a weighted objective (Equation 14) whose main coefficients follow a curriculum schedule, transitioning from predominantly data-driven regression in the initial phase to physics-enforced training once the recurrent backbone has been stabilized.

Importantly, since rocking dynamics comprises binary regime shifts between; (1) non-overturning, for which the rigid-body rocking equation of motion is valid, and (2) overturning events; an auxiliary classifier is used to predict the probability of overturning based on the hidden sequence. The classifier is trained using a cross-entropy loss, and its output is used to down-weight the physics residuals $\mathcal{L}_{\text{PE}}$ near predicted overturning events, thereby preventing enforcement of the equation of motion beyond its physical validity. This joint regression classification formulation substantially improves the robustness and generalization of the

10969845, 2026, 12, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/eqe.70239 by Imperial College London, Wiley Online Library on [12/10/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

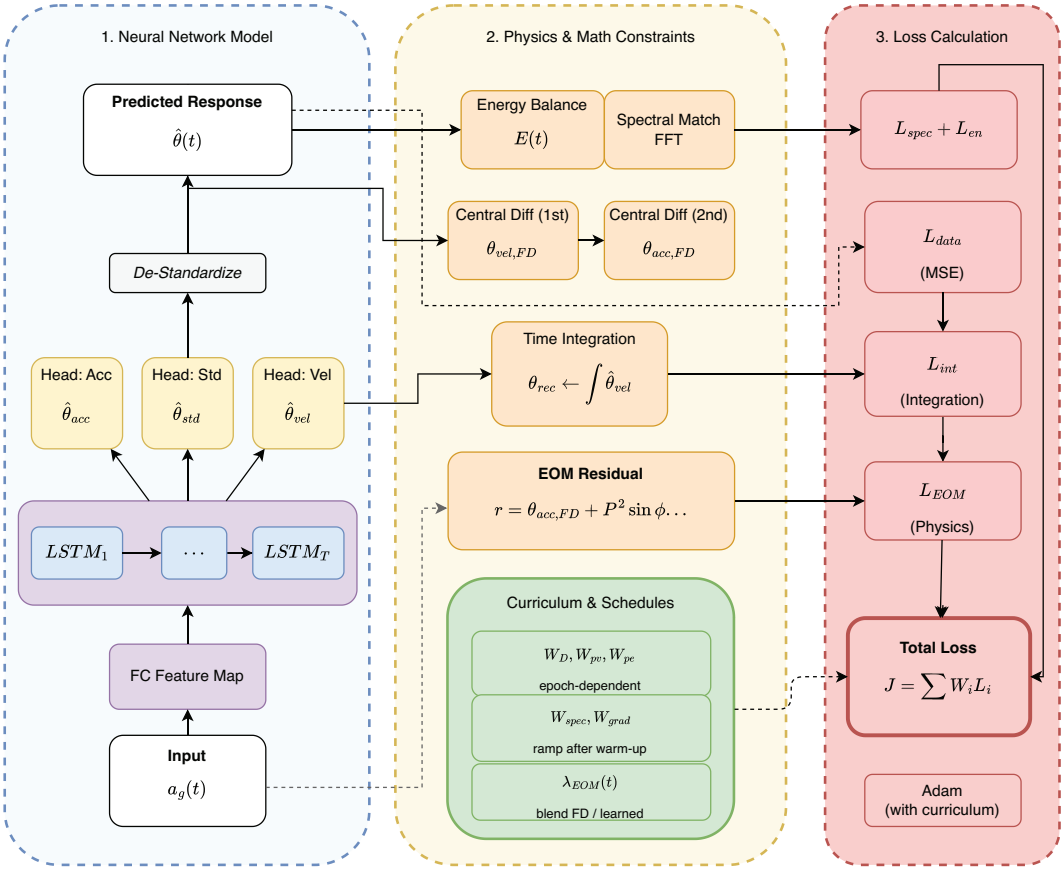


FIGURE 5 | Architecture of the newly proposed Hybrid PINN model. A shared FC-LSTM backbone encodes the input ground motion and outputs standardized angle, velocity, and acceleration through a multi-head prediction layer. Physics modules (finite differences, equation of motion residual, spectral and energy consistency, and integration) generate multiple losses, which are combined through epoch-dependent curriculum weights to train the network.

Hybrid PINN-LSTM across a wide range of ground motion characteristics and intensity levels.

The overall optimization follows a staged learning curriculum, controlled by the loss weights and the blend factor λ in three phases (see Figure 5): (1) An initial phase where data fidelity (α) dominates, with $\lambda \approx 0$, allowing the recurrent core to establish a stable representation; (2) A transitional phase, where the weights shift to increase the influence of physics constraints ($\gamma, \beta_1, \beta_2, \delta$), while the blend factor λ gradually transitions towards 1; (3) a final regime, where the model is supervised by a balanced combination of data and physics. As training proceeds, increasing emphasis is placed on the equation-of-motion residuals, finite-difference consistency, and integration constraints, while the blend factor transitions from purely finite-difference to mixed FD-network accelerations. Hence, the model is supervised by a balanced combination of data and physics, enabling it to produce dynamically admissible rocking and robust generalization across a broad range of excitation intensities.

In summary, the proposed methodology, summarized in Figure 6, integrates experimental mechanics with scientific ML through a four-stage framework designed to predict and analyze the rocking response of rigid blocks. The framework initiates with Stage 1: Data Generation, comprising a targeted shaking-table experimental campaign to capture acceleration and video data based on low-fidelity measurements. These measurements allow for the collection of rocking histories, denoted as $\dot{u}_g(t)$ and $\theta(t)$. This feeds into Stage 2: Data and Physics Preprocessing, where the ground-motion IM are calculated and a video tracking algorithm converts displacement $x(t)$ into rotation $\theta(t)$. Crucially, this stage also computes physics residuals derived from the rigid-rocking Equation of Motion to construct a hybrid training set $\{\dot{u}_g, \theta\}$ that bridges data-driven and physics-based domains.

The computational core of our framework is implemented in Stage 3: Hybrid PINN Model, which uses a Hybrid PINN employing a Fully Connected Long Short-Term Memory (FC-LSTM) architecture which has been described in detail in this section. This framework processes input sequences of ground acceleration ($\dot{u}_g$) to generate predicted rocking responses. To ensure transparency and engineering relevance, the workflow concludes with Stage 4: Interpretable (SHAP), which applies SHapley Additive exPlanations at both Surrogate and IM levels. This provides critical insights into the physical determinants of the response.

The proposed model uses a two-layer unidirectional LSTM, predicting each time step from preceding inputs only, consistent with forward response prediction. The LSTM output feeds a multi-head layer estimating rotation and auxiliary kinematic quantities, with physics-informed losses enforcing

10969845, 2026, 12, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/eqe.70239 by Imperial College London, Wiley Online Library on [12/10/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

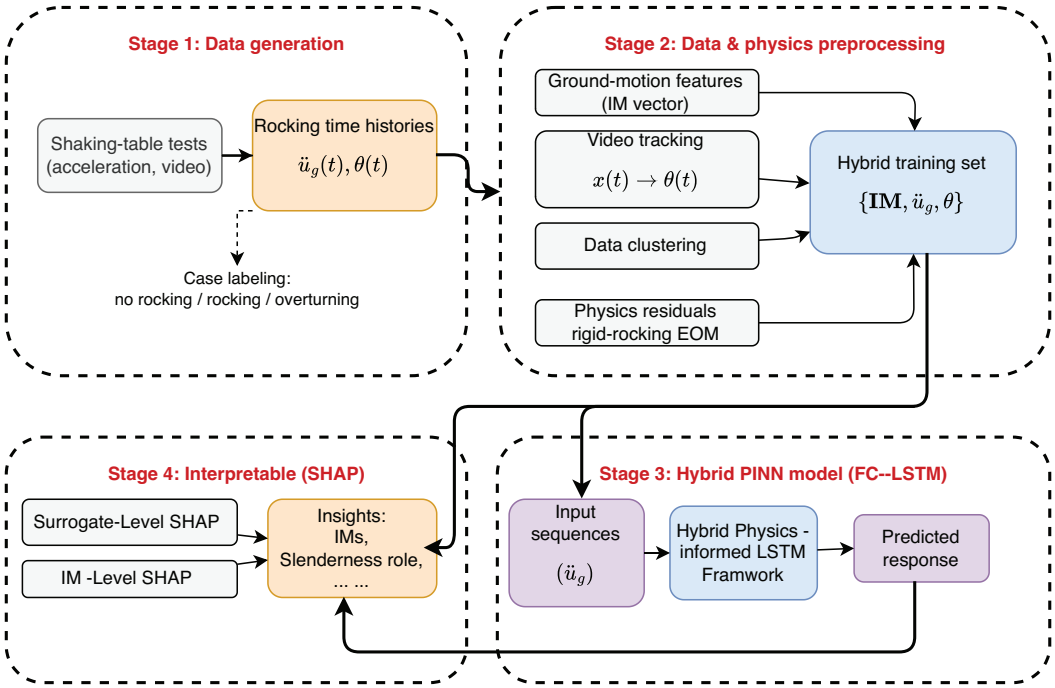


FIGURE 6 | Four-stage Hybrid PINN framework for free rigid rocking. Stage 1 collects rocking data from a small number of shaking-table experiments on the object of interest. Stage 2 constructs a hybrid training set by combining video-based tracking, data clustering, intensity measures, and physics residuals. Stage 3 trains a Hybrid PINN with an FC–LSTM backbone constrained by rigid-rocking equations of motion. Stage 4 employs SHAP-based interpretability to explain how structural and ground-motion parameters influence the predicted rocking responses.

consistency with the governing rocking mechanics during training.

To facilitate reproducibility and support future application to other rocking systems, the implementation details of the proposed framework are summarized in Appendix A. This appendix reports the sequence-model architecture, input and output variables, preprocessing of the experimental time-series data, normalization procedure, train-validation-test split, hybrid loss formulation, loss weighting strategy, smoothing parameters used in the regularized governing equations, and training protocol. These details are provided to allow technically competent researchers to implement, retrain, and adapt the proposed framework for new rocking configurations.

5 | Model Validation

The predictive performance of the proposed Hybrid PINN-LSTM model is validated in this section using the experimental database of rigid-block rocking responses described previously in Section 2. We assess the capability of the model to reproduce the rotation $\theta(t)$, angular velocity $\hat{\theta}(t)$, and angular acceleration $\ddot{\theta}(t)$ over a selected range of ground motion intensities. The validation covers both response-history accuracy and statistical agreement with experimental measurements. The dataset was randomly partitioned into 80% for training and 20% for testing.

To rigorously assess the generalization capabilities of the Hybrid PINN-LSTM framework, it is insufficient to treat the experimental records as a homogeneous dataset. This is because, as shown by various researchers [19, 47, 48], the rocking response of a free-standing block is highly sensitive to the frequency, pulse characteristics, duration, and energy persistence of the ground motion. To capture this variability within a small set of experiments, a K-means clustering analysis was performed on the input ground-motion dataset using a feature space defined by the Peak Ground Acceleration (PGA) and Peak Ground Velocity (PGV), as shown in Figure 7. This unsupervised procedure revealed four statistically distinct excitation regimes (Figure 7).

Figures 8 and 9 compare the measured rotation histories against the Hybrid PINN-LSTM predictions for six representative ground-motion records that are not part of the training/testing dataset for Specimens 3 and 4, respectively, as an example. Similar results are obtained for other specimens. It is evident from these figures that the model captures the key dynamical features of the rocking response across all cases, including the amplitude modulation induced by the nonstationary ground-motion components, the block's motion associated with its geometric nonlinearity, and the rapid decay following impacts. The comparison demonstrates that our newly proposed PINN-LSTM model accurately reproduces both the timing and magnitude of the rocking responses, with minimal cumulative drift and correct representation of impact-induced velocity changes. The consistent performance across all records highlights the model's ability to generalize to diverse ground motions.

Figure 10 presents a typical comparison of experimental rotation response, PINN-LSTM predictions, and a classical Housner-type rocking solution (MATLAB, ODE45) for two ground-motion records on Specimen 4 in terms of rotation histories. The Housner model's coefficient of restitution was estimated experimentally, though this estimate is sensitive to measurement noise and

10969845, 2026, 12, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/eqe.70239 by Imperial College London, Wiley Online Library on [12/10/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

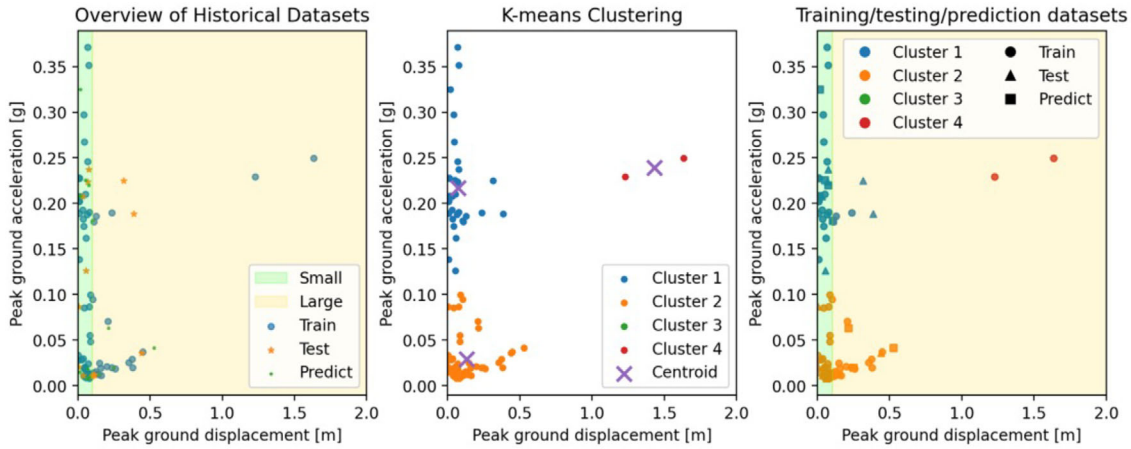


FIGURE 7 | K-means Clustering on the ground-motion dataset used is performed to ensure the heterogeneity of the dataset is reflected at different stages of the training and testing processes.

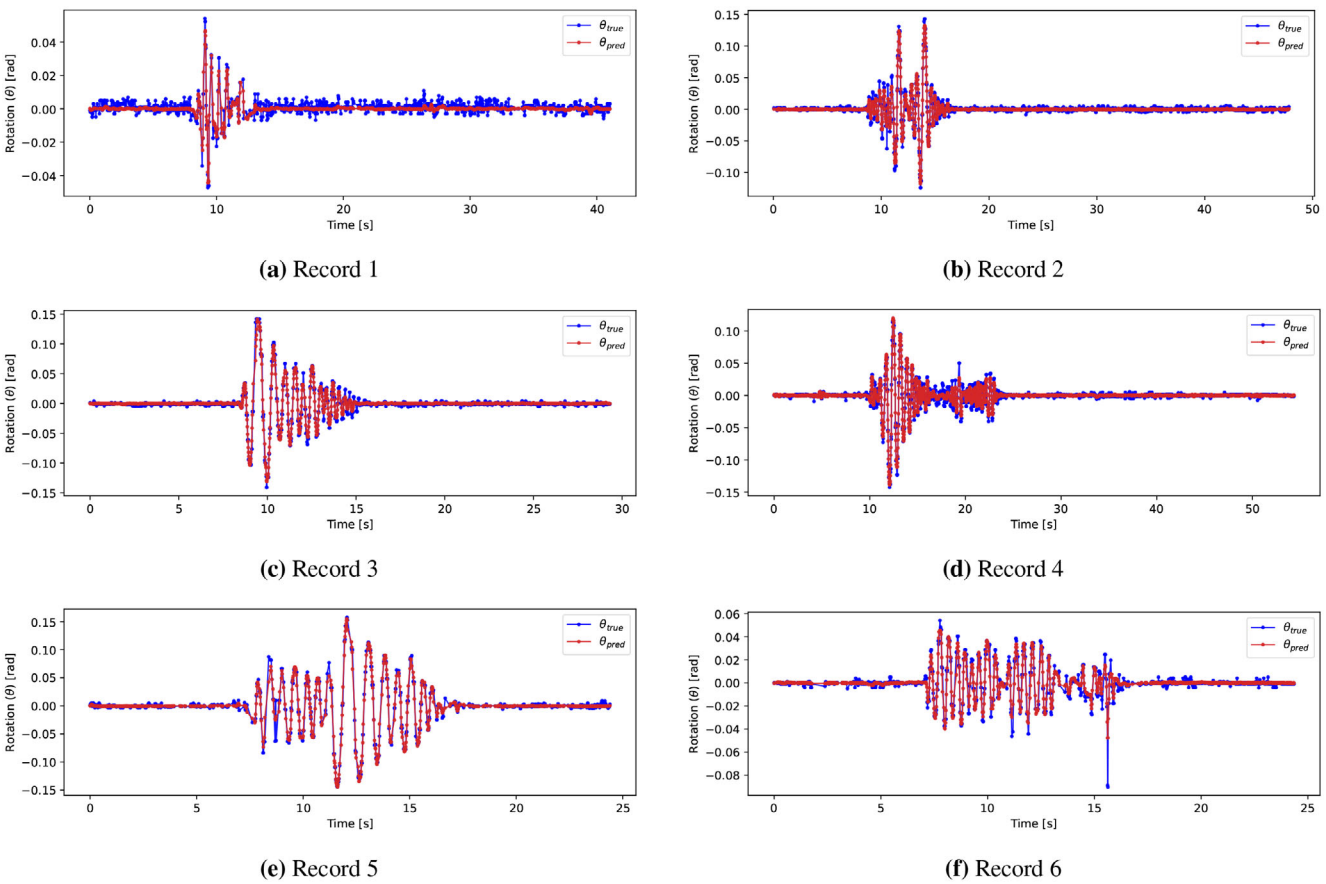


FIGURE 8 | Comparison of measured and Hybrid PINN-LSTM predicted rotation histories $\theta(t)$ for six representative ground-motion records. Panels (a)–(f) correspond to Records 1 to 6 on Specimen 3, respectively. The model accurately reproduces the amplitude modulation and decay patterns of the rocking response across a wide range of excitation characteristics.

differentiation choices and should be read as representative rather than calibrated. The Housner model captures the general rocking mechanism but shows amplitude and phase discrepancies, particularly when the same restitution coefficient is applied across records, whereas the PINN-LSTM provides a closer match to experimental data in both cases, underscoring the value of combining governing-equation priors with data-driven learning.

To further evaluate the model predictions, Figure 11 presents a representative comparison between the predicted and measured responses for Record 4. Again, the Hybrid PINN-LSTM

10969845, 2026, 12, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/eqe.70239 by Imperial College London, Wiley Online Library on [12/10/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

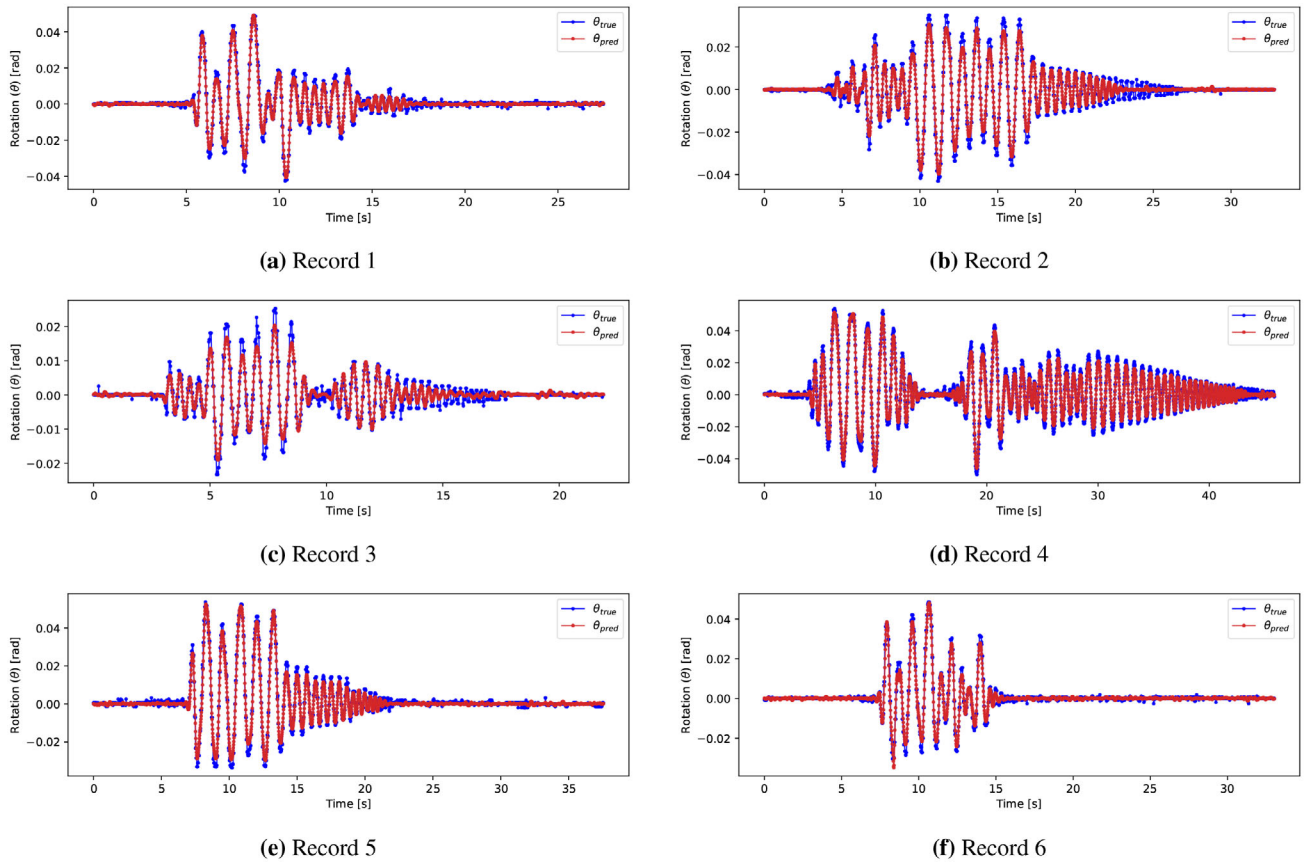


FIGURE 9 | Comparison of measured and Hybrid PINN-LSTM predicted rotation histories $\theta(t)$ for six representative ground-motion records. Panels (a)–(f) correspond to Records 1 to 6 on Specimen 4, respectively. The model accurately reproduces the amplitude modulation and decay patterns of the rocking response across a wide range of excitation characteristics.

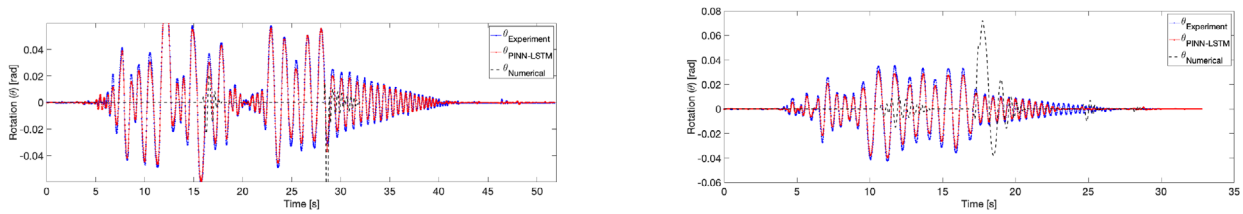


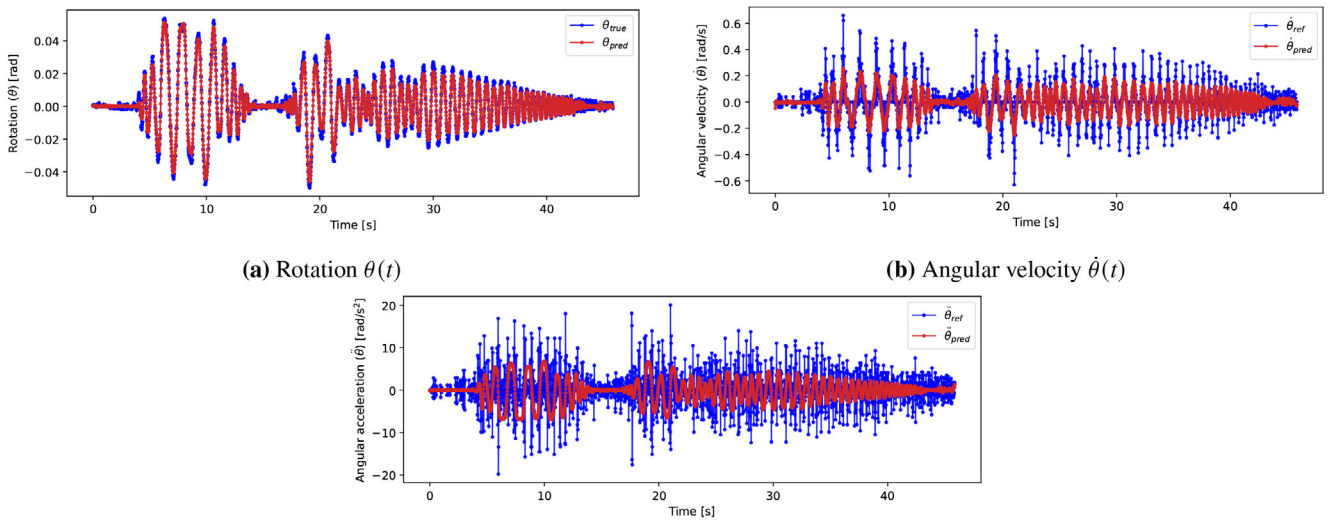
FIGURE 10 | Measured versus predicted rotation histories for two ground-motion records: PI-LSTM, Housner-type model (MATLAB ODE45), and experiment. PI-LSTM shows closer agreement, while the Housner model exhibits amplitude and phase discrepancies.

accurately reproduces the structural response of the rocking motion. The rotation response $\theta(t)$ shows near-perfect phase alignment and close amplitude matching throughout the whole duration. The angular velocity $\dot{\theta}(t)$ remains well captured, with slightly increased scatter near high-velocity impacts. However, and as expected, the prediction of angular acceleration $\ddot{\theta}(t)$ is the most challenging due to the severe spikes induced by impact events; nevertheless, the model still reproduces the overall trend and envelope. These results illustrate that the Hybrid PINN-LSTM learns a physically consistent mapping between the input ground motion and the nonlinear, non-smooth rocking dynamics.

Tables 3 and 4 summarize the model's accuracy using RMSE, MAE, R^2, Pearson correlation, and error quantiles across all prediction tests. Rotation θ achieves excellent accuracy with RMSE = 9.08×10^{-3}, MAE = 3.73×10^{-3}, and $R^2 = 0.704$, demonstrating that the physics constraints effectively regularize the model. Angular velocity $\dot{\theta}$ shows moderate accuracy (RMSE $\approx 9.33 \times 10^{-2}$, $R^2 = 0.325$), reflecting stronger sensitivity to small timing errors around impact events. As expected, angular acceleration $\ddot{\theta}$ exhibits the highest dispersion (RMSE ≈ 4.25, negative R^2), but the median per-record RMSE remains reasonable, indicating that most records are predicted with moderate accuracy and that the large overall error is dominated by a subset of high-intensity tests with strong impacts.

Moreover, the dense scatter plots and residual histograms presented in Figure 12 for all data, provide additional insight. Rotation and angular velocity predictions cluster tightly around the idealized line, with residuals strongly concentrated near

10969845, 2026, 12, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/eqe.70239 by Imperial College London, Wiley Online Library on [12/10/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License



(a) Rotation $\theta(t)$

(b) Angular velocity $\dot{\theta}(t)$

(c) Angular acceleration $\ddot{\theta}(t)$

FIGURE 11 | Comparison of measured and Hybrid PI-LSTM-predicted rotation $\theta(t)$, angular velocity $\dot{\theta}(t)$, and angular acceleration $\ddot{\theta}(t)$ for Specimen 4 subjected to Record 4.

TABLE 3 | Summary of Hybrid PINN-LSTM prediction accuracy over all records.

Target	RMSE	MAE	R^2	Pearson r	Bias (mean error)	median $\lvert e \rvert$	$p_{90}(\lvert e \rvert)$
Specimen 3							
θ	0.00391	0.00202	0.9382	0.9690	-9.53×10^{-6}	0.00095	0.00483
$\dot{\theta}$	0.17465	0.08014	0.3858	0.6228	3.73×10^{-4}	0.02685	0.18477
$\ddot{\theta}$	8.19822	3.65622	0.0580	0.2510	5.28×10^{-2}	1.11836	8.87240
Specimen 4							
θ	0.00908	0.00373	0.704	0.840	−0.00026	0.00132	0.00762
$\dot{\theta}$	0.0933	0.0547	0.325	0.573	−0.00283	0.0262	0.1425
$\ddot{\theta}$	4.2449	2.6029	−0.177	0.115	−0.0905	1.2490	7.0899

TABLE 4 | Per-record RMSE statistics for each target.

Target	mean RMSE_{rec}	$p_{50}(\text{RMSE}_{\text{rec}})$	$p_{90}(\text{RMSE}_{\text{rec}})$
Specimen 3			
θ	0.00355	0.00314	0.00602
$\dot{\theta}$	0.15326	0.12616	0.25663
$\ddot{\theta}$	7.05088	5.80090	12.19415
Specimen 4			
θ	0.00629	0.00302	0.01565
$\dot{\theta}$	0.09371	0.08050	0.15042
$\ddot{\theta}$	4.31524	4.20221	5.46029

zero. Acceleration predictions exhibit broader tails, but remain centered around zero bias. The best agreement is observed for the rotation $\theta(t)$, for which the PI-LSTM achieves an RMSE of 9.08×10^{-3}, negligible bias, and a correlation coefficient of $r = 0.84$. Record-wise statistics confirm that the model reproduces the rotational predictions consistently across the full dataset, with a median RMSE_{rec} of only 0.003. This strong performance reflects the ability of the physics-informed loss to enforce the governing equation of motion and geometric relations, ensuring that the learned solution remains consistent with the low-frequency dynamics that dominate rocking behavior between impacts. The results are of comparable accuracy with high-fidelity postdiction finite element (FE) simulations employing detailed contact modelling [49], yet our model arrives at these predictions without the heavy computational cost typically required by such models and, importantly, without enforcing a particular coefficient of restitution modelling assumption.

6 | Discussion

We have seen that the predictive performance of the newly proposed PINN-LSTM model partially decreases when examining the angular velocity response $\dot{\theta}(t)$ and deteriorates further for angular accelerations $\ddot{\theta}(t)$. Although the PINN-LSTM captures the overall angular velocity evolution and the trends associated

10969845, 2026, 12, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/eqe.70239 by Imperial College London, Wiley Online Library on [12/10/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

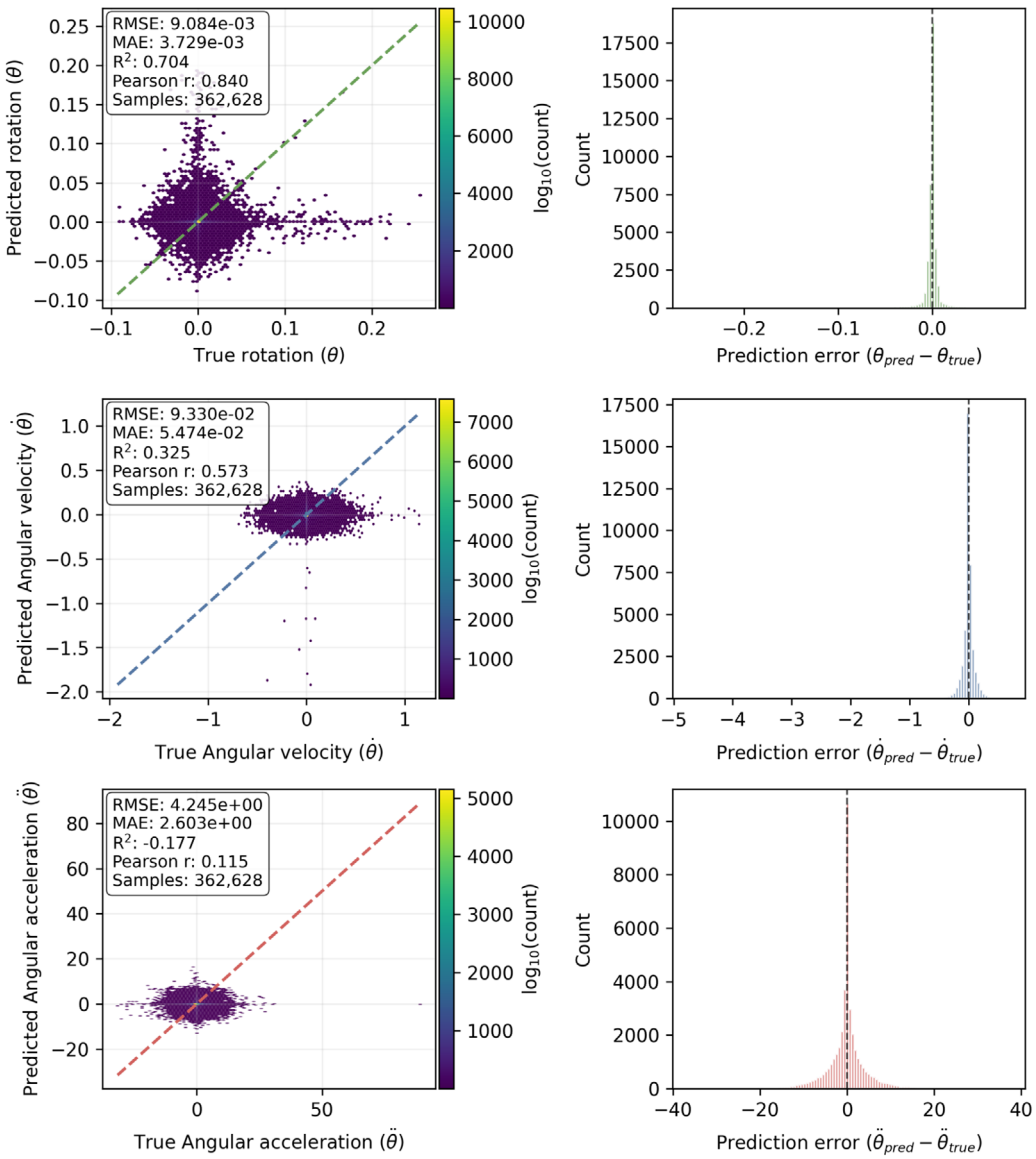


FIGURE 12 | Comparison of predicted and true rotation, angular velocity, and angular acceleration over the evaluation split. Each row shows the scatter (left) and log-count residual histogram (right) for one head, using consistent axis limits.

with each rocking cycle pretty well, the RMSE increases to 0.093 and the correlation coefficient decreases to $r = 0.573$ for $\dot{\theta}(t)$ in Specimen 4, for example. This behavior is expected because $\dot{\theta}(t)$ is substantially more sensitive than $\theta(t)$ to both experimental noise and the smoothing inherent to the neural network architecture. Besides, differentiation amplifies high-frequency noise in the measured rotation signals, and the physics-informed loss tends to regularize the solution by enforcing derivative consistency through automatic differentiation, thereby reducing the sharp velocity discontinuities that occur at lift-off, re-contact, and impact. Similar difficulties occur in finite element simulations, where predicted velocity jumps depend strongly on the chosen time-step, contact regularization strategy, and penalty stiffness, often requiring extremely small time increments to capture the discontinuity accurately. Analytical impact models also rely on semi-empirical restitution coefficients to match the experimentally observed loss in angular velocity, reflecting the fact that impact dynamics cannot be learned from first principles alone.

In this context, the Hybrid PINN-LSTM framework predicts angular velocities that remain consistent with the expected physical behavior during the smooth nonlinear dynamics between impacts. While the sharpness of the experimental velocity jumps is not perfectly reproduced, the model effectively precludes spurious oscillations, artificial energy gain, and cumulative drift. Hence, the resulting fidelity represents a computational efficient alternative to established analytical and numerical approaches, offering a physically plausible approximation of the discontinuous dynamics inherent to rigid body rocking.

However, the largest discrepancies arise for the rotational acceleration $\ddot{\theta}(t)$. Here, the RMSE exceeds 4.24 and the coefficient of determination becomes negative, indicating limited ability to replicate the impulsive, high-frequency components associated with impacts. This behavior is fully consistent with the mathematics of the PINN formulation: standard activation functions enforce continuous and differentiable network outputs,

10969845, 2026, 12, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/eqe.70239 by Imperial College London, Wiley Online Library on [12/10/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

and the physics residual is integrated over smooth domains. Consequently, the model captures the overall acceleration trend (median absolute error ≈ 1.25), but cannot reproduce the large, discontinuous angular acceleration. In contrast, FE models with explicit contact constraints can represent these discontinuities but require extremely small time increments, while analytical rocking models represent impulses via discrete impact maps. The limitations observed in $\ddot{\theta}(t)$ therefore arise from the inherent smoothness of conventional PINN approximations rather than from data scarcity.

In fact, it is expected that the predictive accuracy naturally degrades for higher-order derivatives. While the angular velocity predictions capture global oscillatory trends, the model tends to smooth the sharp discontinuities characteristic of impact events, resulting in a reduced R^2 of 0.325. The angular acceleration predictions exhibit the largest discrepancies, systematically underestimating peak spikes. To interpret these results rigorously, it is necessary to decouple the epistemic uncertainty inherent in the neural network from the measurement uncertainty present in the experimental database.

A primary source of model-induced uncertainty is the phenomenon of spectral bias, or the Frequency Principle, a well-documented characteristic of fully connected neural networks wherein low-frequency components of the target function are learned significantly faster than high-frequency components. Rigid-block rocking is mathematically characterized by non-smooth dynamics: $\theta(t)$ is continuous, but $\dot{\theta}(t)$ exhibits step discontinuities, and $\ddot{\theta}(t)$ theoretically approaches a Dirac delta function at impact [37]. These features lie precisely in the portion of the spectrum that neural networks learn most slowly, making them especially difficult to capture. A PINN-LSTM, utilizing standard continuous activation functions would be inherently biased towards smooth solutions. Consequently, the NN effectively identifies the timing and global intensity of impacts but lacks the high-frequency expressivity to replicate the singularity of the physical impact event.

On the other hand, the error metrics for acceleration must be examined within the significant uncertainty of an experimental database. The measured rotation and input ground motion were derived from optical tracking. Deriving angular acceleration necessitates double numerical differentiation of this discrete position signal, an operation that amplifies quadratically high-frequency measurement noise. As a result, the experimental angular acceleration contains substantial noise that does not reflect the true motion. The physics-informed model, governed by the analytical Equation of Motion, rightfully rejects this nonphysical noise. Therefore, the negative R^2 observed for acceleration is partly an artifact of comparing a spectrally smoothed, physics-compliant prediction against a noisy, numerically differentiated baseline. In this light, the Hybrid PINN-LSTM provides a smooth estimate of rigid-body forces that may be more representative of global stability mechanics than the raw experimental data implies.

In summary, the PINN-LSTM successfully reproduces the essential nonlinear rocking dynamics. For engineering applications governed by rotation or drift demands, such as overturning stability assessment or rotation-based fragility modelling, the framework offers a computationally efficient alternative to complex numerical simulations. However, for acceleration histories, the results suggest that targeted extensions, such as the addition of impact-aware terms or discontinuous neural operators, currently under investigation, may be requisite to capture the impulsive nature of rocking impacts with higher precision.

The present study is primarily concerned with the prediction of stable rocking response within the range represented by the available experimental dataset. The training set used for the main regression model was restricted to non-overturning cases. Although this provides a meaningful validation basis for predicting stable rocking motion, overturning and instability thresholds remain outside the current validated scope of the model. The model should therefore not be interpreted as an overturning classifier or as a direct predictor of failure-regime response.

It should also be noted that the ground-motion records were pre-processed to satisfy the physical constraints of the shaking-table tests. High-pass filtering and the imposed PGD limit were required to avoid excessive table displacement. These preprocessing steps may reduce the representation of long-period pulse-type motions, which are often relevant to overturning response. Consequently, the current dataset may not fully cover instability-dominated cases associated with strong long-period pulse content. Although we expect a good degree of generalization, extending the framework to include overturning cases, pulse-type ground motions, and a dedicated instability or classification branch should be treated with care in future work.

Our framework is different from conventional PINN formulations. Standard PINNs usually use feed-forward neural networks to approximate continuous solution fields and impose the governing equations at collocation points, often through automatic differentiation. These formulations are generally designed for problems where the solution is represented as a continuous function of space and time. In the present work, the PINN-LSTM framework is instead developed for discrete experimental time-series prediction. The input is the measured ground acceleration history, and the output is the corresponding rocking response history. The LSTM structure provides temporal memory, allowing the model to learn how the current response depends on the previous excitation and response states. The governing rocking equation is not imposed as a hard constraint. Instead, it is introduced as a soft physics-informed loss applied to the predicted response histories. This formulation is therefore more suitable for noisy experimental data, while still encouraging consistency with the underlying rocking mechanics.

Several improvements could be considered in future work to tackle the difficulties arising from the non-smooth dynamics of rocking blocks. For example, impact transitions may be represented more explicitly, for example through impact-aware loss terms, event-detection modules, or hybrid formulations that distinguish between free-rocking motion and impact instants. The estimation of velocity and acceleration could also be improved by using more robust derivative reconstruction or smoothing strategies, or by adopting a multi-task formulation in which rotation, velocity, and acceleration are constrained jointly. In addition, adaptive or uncertainty-based loss weighting could be used to balance the supervised data loss and the physics-informed

10969845, 2026, 12, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/eqe.70239 by Imperial College London, Wiley Online Library on [12/10/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

residual during training. More expressive temporal models, such as attention-based sequence models, temporal convolutional networks, or hybrid LSTM-attention architectures, may also help the model capture the short-duration, high-frequency features associated with impacts.

7 | Model Interpretation

To ensure that our Hybrid Physics-Informed Long Short-Term Memory (PINN-LSTM) model operates as a transparent physical engine rather than a black box, a comprehensive interpretability analysis is performed in this section using SHAP (SHapley Additive exPlanations). This game-theoretic framework is widely used in engineering applications of AI since it enables the decomposition of the model's predictions and the attribution of predictive contributions to specific input features [32]. Given the inherently complex nature of seismic data, which includes both high-dimensional histories and derived engineering IMs, SHAP analysis is carried out at two complementary levels: the Sequence Level, assessing feature importance across the time domain, and the Surrogate Level, evaluating contributions within the IM domain.

SHAP values attribute the prediction of a ML model to its input features. For a given model $f(x)$ and a specific input instance x, the SHAP value ϕ_i assigned to the i-th feature is the average marginal contribution of that feature across all possible subsets (coalitions) of inputs. This is mathematically expressed as:

$$\phi_i = \sum_{S \subseteq F \setminus \{i\}} \frac{|S|!(|F| - |S| - 1)!}{|F|!} \tag{22}$$

In this formulation, F denotes the set of all features, S represents a subset of features excluding feature i, and $|S|$ is the size of the subset. The term $f(S)$ represents the model's expected output conditioned on the subset S. A defining characteristic of SHAP is the additivity axiom, which ensures that the sum of the SHAP values for all features equals the deviation of the specific prediction $f(x)$ from the baseline average prediction ϕ_0. This property, expressed in Equation 23, provides a rigorous accounting of the factors driving the structural response prediction.

$$f(x) = \phi_0 + \sum_{i=1}^{M} \phi_i \tag{23}$$

This formulation ensures consistency, local accuracy, and additivity, meaning that the sum of individual SHAP values exactly reconstructs the model's prediction relative to a specified baseline. For time-series and structural dynamics applications, SHAP provides a principled mechanism for identifying which components of an input signal or IM exert the greatest influence on a predicted structural response. This makes the method particularly valuable for analyzing physics-informed models, where interpretability and physical traceability are essential.

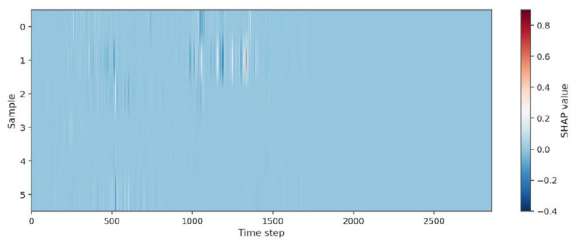
To interpret the behavior of our newly proposed PINN-LSTM) model, SHAP was first applied at the time-step level, treating each sample of $a_g(t)$ as an input feature. This sequence-level SHAP analysis highlights the portions of the excitation in time that most strongly influence the predicted rocking response (e.g., $\theta(t)$, $\theta_{\max}$, or overturning state). This analysis reveals that the PINN-LSTM has learned to function as a highly selective non-linear filter.

As illustrated in the SHAP heatmap (Figure 13a), the model's attention is not uniformly distributed across the duration of the seismic record. Instead, the importance scores are sparse and localized, forming distinct vertical bands that correspond to specific high-energy excursions in the ground motion. This observation is further quantified in Figure 13b, which plots the mean absolute SHAP value across all samples. The network exhibits negligible sensitivity to the low-amplitude phases of the excitation (e.g., $t < 5$s or $t > 30$s), focusing instead on sharp peaks of importance confined to narrow time windows (e.g., around $t \approx 9$s and $t \approx 22$s). These windows typically align with the arrival of coherent velocity pulses, confirming that the model has correctly identified the impulsive nature of the ground motion as the primary mechanism initiating and sustaining the rocking motion.

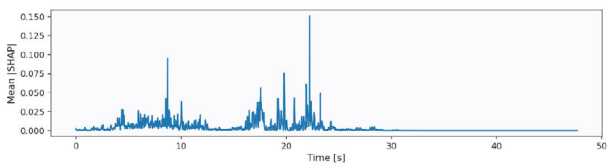
Transforming the SHAP importance values into the frequency domain provides perhaps the strongest validation of the model's physical consistency. The spectral importance profile (Figure 13c) demonstrates a dramatic decay in sensitivity as frequency increases. The vast majority of the model's decision-making power is concentrated in the ultra-low frequency range ($f <$ 2.0 Hz), with the mean importance effectively vanishing for higher frequencies. This behavior is examined further in the frequency-band beeswarm plot (Figure 13d). Here, the 0.00–1.50 Hz band emerges as the dominant predictor; high spectral amplitudes in this range (red dots) overwhelmingly drive the model towards predicting large rocking rotations or overturning. Conversely, frequency bands above 3.0 Hz show a mixed or negligible contribution, clustering tightly around zero. This hierarchy aligns perfectly with the theoretical understanding of rigid blocks as long-period structures that are susceptible to low-frequency resonance and displacement demands, while being largely immune to high-frequency, high-acceleration noise.

In addition to the sequence-level analysis, we also performed an IM-Surrogate SHAP analysis. To this end, a lightweight tree-based surrogate model that maps the same IM vectors to the PINN-LSTM predictions (rather than to the experimental response) was constructed. This IM-surrogate does not replace the PINN-LSTM; instead, it serves as a lens that reveals which conventional IMs best emulate the behavior of the hybrid network. When the surrogate reproduces the PINN-LSTM outputs with high accuracy, the corresponding SHAP beeswarm plots can be interpreted as an IM-level summary of what the PINN-LSTM has learned from the full waveforms. This enables statements such as PGA and PGD provide a compact summary of the features the PINN-LSTM uses to predict $\theta_{\max}$, and facilitates a direct comparison between the sequence-based hybrid model and traditional IM-based approaches. Two IM formulations are examined in this study. The first corresponds to a broad set of enhanced, feature-engineered predictors specifically developed for rocking dynamics and introduced in [50]. This set was designed to capture nonlinear causal pathways between ground-motion characteristics and rigid-block response and incorporates displacement, velocity, spectral, and pulse-related measures. The second formulation adopts a more compact, physics-inspired selection of IMs based on classical rigid-body mechanics, primarily comprising peak PGD, PGV, PGA, and Cumulative Absolute

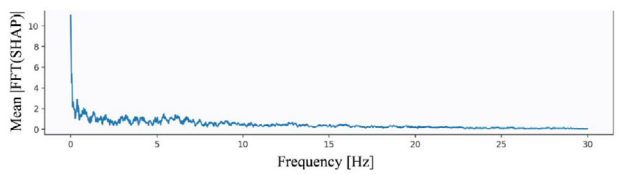
10969845, 2026, 12, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/eqe.70239 by Imperial College London, Wiley Online Library on [12/10/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License



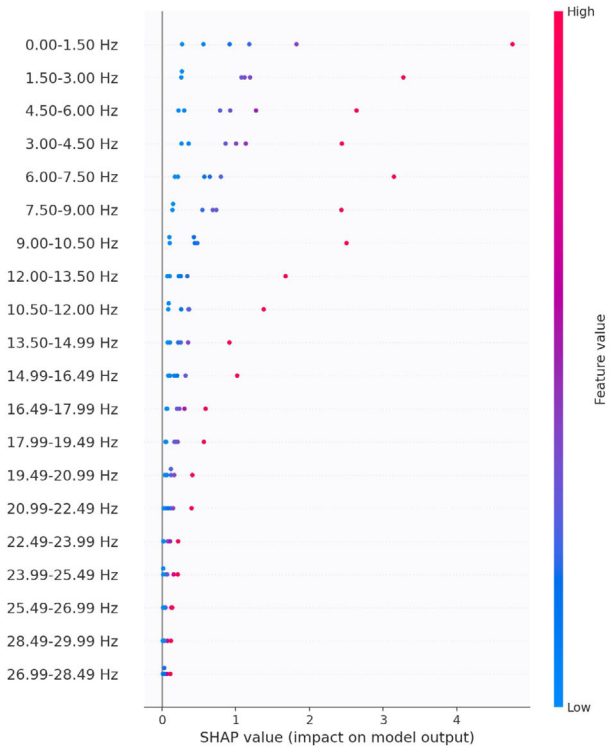
(a) Time-domain SHAP heatmap



(b) Mean temporal importance profile

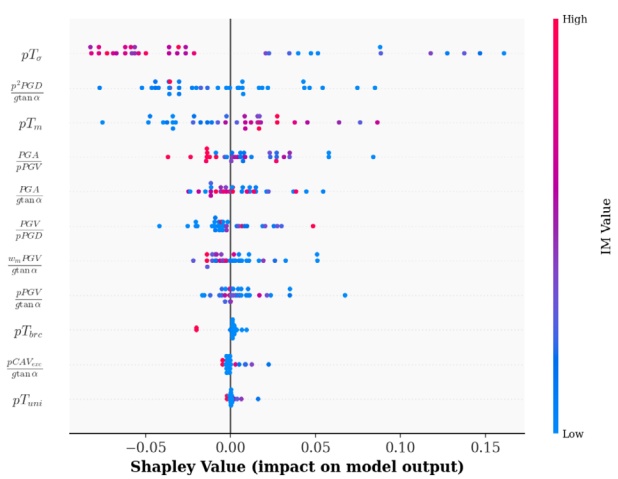


(c) Mean spectral importance (FFT of SHAP)

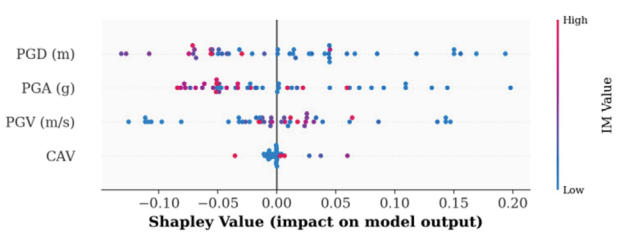


(d) Frequency band importance beeswarm

FIGURE 13 | Sequence-level SHAP analysis of the Hybrid PINN-LSTM model. Panels (a), (b), and (c) describe the model's temporal and spectral attribution patterns, demonstrating temporal sparsity and a focus on coherent pulses. Panel (d) illustrates the distribution of Shapley contributions across frequency bands, confirming that the model behaves as a low-pass filter emphasizing the 0.0–1.5 Hz band consistent with rigid-body rocking dynamics.



(a) SHAP summary for basic intensity measures.



(b) SHAP summary for physics-inspired intensity measures.

FIGURE 14 | SHAP-based interpretation of the model using two IM feature sets: (a) enhanced IMs proposed in [50], and (b) physics-inspired IMs based on classical rocking mechanics.

Velocity (CAV), which represent the dominant kinematic drivers of rocking initiation and subsequent motion.

Figure 14 presents the SHAP summary for both formulations. Several clear patterns emerge. Spectral displacement at the system period, pT_0, and its higher-order variants exhibit the highest positive Shapley values, indicating that long-period demand strongly increases rocking rotation, as expected. Likewise, velocity-related IMs such as $p\text{PGV}_{\text{ghar}}$ and PGV also rank highly, consistent with the established understanding that rocking initiation is governed primarily by velocity pulses rather than high-frequency acceleration components. In contrast, high-frequency IMs such

10969845, 2026, 12, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/eqe.70239 by Imperial College London, Wiley Online Library on [12/10/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

as PGA show limited importance, confirming their weak causal influence on free rocking amplitude.

The SHAP analysis for the physics-based IM set shown in Figure 14b produces a more compact but consistent picture. PGD and PGV emerge again as the dominant predictors, reflecting the fundamental role of displacement and velocity demand in rocking dynamics. CAV contributes moderately, while PGA remains of minor importance. These patterns mirror both theoretical models and the causal structure identified in [50].

Overall, the SHAP interpretation demonstrates that our Hybrid PINN-LSTM not only achieves high predictive accuracy but also learns a physically interpretable hierarchy of ground-motion features. The alignment between the model's inferred importance ranking and prior causal analyzes of rocking behavior provides additional confidence that the network is capturing meaningful mechanics rather than relying on black-box correlations.

8 | Conclusion

A novel Hybrid Physics-Informed Neural Network (Hybrid PINN-LSTM) framework enhanced with Explainable AI (XAI) for the seismic response prediction of rocking structures has been presented. By integrating governing physical laws, in the form of the non-linear equations of rigid rocking motion, with limited experimental data from shake table tests (in the order of a few tens only), the proposed model addresses the limitations of current purely data-driven approaches and purely analytical methods.

The Hybrid PINN-LSTM architecture demonstrated superior predictive capabilities compared to standard Deep Neural Networks and purely data-driven LSTMs. By embedding the physics of rocking motion directly into the loss function, the model successfully captured complex non-linear behaviors even with limited training data. The incorporation of experimental data proved critical for calibrating these physics-based components. Importantly, this hybrid approach allowed the model to learn real-world rocking dynamics without making any common a priori assumption about abstractions like the coefficient of restitution, which are often oversimplified in analytical rigid-block formulations.

Furthermore, the application of SHAP (SHapley Additive exPlanations) enhanced the transparency of our model in contrast with the typically opaque nature of neural networks. This XAI analysis revealed that the Hybrid PINN-LSTM model correctly identifies specific IMs, such as peak ground velocity (PGV) and acceleration (PGA), as the dominant features driving the rocking response confirming that the network has effectively learned the underlying mechanics of the system.

The proposed Hybrid-PINN framework provides significant computational advantages, achieving fast prediction of nonlinear rocking response once trained on a relatively small number of experimental outputs. This efficiency represents a substantial reduction in computational cost compared with postdiction high-fidelity finite element models or extensive nonlinear response-history analyzes, while maintaining robustness to noise and variability in the input ground motion. Although the present study focuses on single rigid blocks, the framework is expected to be readily extensible to more complex systems, such as flexible uplift systems and rocking frames. Besides, our integrated explainable-AI methodology further enables a systematic identification of the frequency components and time segments most responsible for controlling rocking behavior, offering new insight into the network's predictions and how the underlying physics are encoded. Future work will develop this framework into parametric design tools and extend the probabilistic assessment of overturning under real ground motions. The final aim is a tool that could estimate the rocking response of critical equipment on the basis of limited experimentation, free of complex impact and restitution modelling assumptions and with a physically informed basis.

In summary, this study advances the modelling of rocking structures by bridging classical mechanics, modern deep learning techniques, and explainable AI. It provides a path towards a physically grounded, data-enhanced, and interpretable solution to one of the most challenging nonlinear problems in earthquake engineering, with potential implications for seismic design, vulnerability assessment, and performance-based engineering of rocking structures.

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

The data that support the findings of this study are available from the corresponding author upon reasonable request.

References

1. G. W. Housner, "The Behavior of Inverted Pendulum Structures During Earthquakes," *Bulletin of the Seismological Society of America* 53 (1963): 403–417.

2. J. Zhang and N. Makris, "Rocking Response of Free-Standing Blocks Under Cycloidal Pulses," *Journal of Engineering Mechanics* 127, no. 5 (2001): 473–483.

3. N. Makris and Y. Roussos, "Rocking Response of Rigid Blocks Under Near-Source Ground Motions," *Geotechnique* 50, no. 3 (2000): 243–262.

4. N. Kalapodis, G. Kampas, and O.-J. Ktenidou, "Revisiting the Fundamental Structural Dynamic Systems: The Effect of Low Gravity," *Archive of Applied Mechanics* 89, no. 9 (2019): 1861–1884.

5. M. J. DeJong and E. G. Dimitrakopoulos, "Dynamically Equivalent Rocking Structures," *Earthquake Engineering & Structural Dynamics* 43 (2014): 1543–1563.

6. M. A. ElGawady, Q. Ma, J. Butterworth, and J. Ingham, "Effects of Interface Material on the Performance of Free Rocking Blocks," *Earthquake Engineering & Structural Dynamics* 40 (2011): 375–392.

7. D. Kalliontzis, S. Sritharan, and A. Schultz, "Improved Coefficient of Restitution Estimation for Free Rocking Members," *Journal of Structural Engineering* 142, no. 12 (2016): 06016002.

8. A. I. Giouvanidis, E. G. Dimitrakopoulos, and P. B. Lourénço, "Chattering: An Overlooked Peculiarity of Rocking Motion," *Nonlinear Dynamics* 109 (2022): 459–477.

9. J. A. Bachmann, M. K. Strand, M. F. Vassiliou, M. Broccardo, and B. Stojadinović, "Is Rocking Motion Predictable?" *Earthquake Engineering & Structural Dynamics* 47 (2017): 535–552.

10969845, 2026, 12, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/eqe.70239 by Imperial College London, Wiley Online Library on [12/10/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

10. D. Kalliontzis and S. Sritharan,, "Characterizing Dynamic Decay of Motion of Free-Standing Rocking Members," *Earthquake Spectra* 34 (2018): 843–866.

11. M. F. Vassiliou, C. Cengiz, M. Dietz, et al., "Data Set From Shake Table Tests of Free-Standing Rocking Bodies," *Earthquake Spectra* 37, no. 4 (2021): 2971–2987.

12. C. Zhong and C. Christopoulos, "Data Set From Shaking Table Tests of a Slender MDOF Structure With Base-Rocking Mechanisms," *Earthquake Spectra* 39, no. 3 (2023): 1930–1944.

13. B. Chiou, R. Darragh, N. Gregor, and W. Silva, "An Overview of the NGA Project Ground-Motion Database," *Earthquake Spectra* 24, no. 1 (2008): 23–44.

14. S. A. Burton, N. Makris, I. Konstantopoulos, and P. Antsaklis, "Modeling the Response of ER Damper: Phenomenology and Emulation," *Journal of Engineering Mechanics* 122, no. 9 (1996): 897–906.

15. S. Shen and C. Málaga-Chuquitaype, "Physics-Informed Artificial Intelligence Models for the Seismic Response Prediction of Rocking Structures," *Data-Centric Engineering* 5 (2024): e1.

16. Z. Achmet, S. Diamantopoulos, and M. Fragiadakis, "Rapid Seismic Response Prediction of Rocking Blocks Using Machine Learning," *Bulletin of Earthquake Engineering* 22 (2023): 3471–3489.

17. S. K. W. Chu, A. I. Giouvanidis, C. N. Loong, and E. G. Dimitrakopoulos, "New Perspectives in Causal Relationships Between the Response of a Rocking Block and Intensity Measures via Ensemble Machine Learning Methodologies," *Earthquake Engineering & Structural Dynamics* 54, no. 14 (2025): 3576–3593.

18. I. Karampinis, K. E. Bantilas, I. E. Kavvadias, L. Iliadis, and A. Elenas, "Seismic Response Prediction of Rigid Rocking Structures Using Explainable Lightgbm Models," *Mathematics* 12 (2024): 2280.

19. C. G. Lachanas, D. Vamvatsikos, and E. G. Dimitrakopoulos, "Intensity Measures as Interfacing Variables versus Response Proxies: The Case of Rigid Rocking Blocks," *Earthquake Engineering & Structural Dynamics* 52 (2023): 1722–1739.

20. T. Chen and H. Chen, "Approximations of Continuous Functionals by Neural Networks With Application to Dynamic Systems," *IEEE Transactions on Neural Networks* 4, no. 6 (1993): 910–918.

21. X. Guan, H. Burton, M. Shokrabadi, and Z. Yi, "Seismic Drift Demand Estimation for Steel Moment Frame Buildings: From Mechanics-Based to Data-Driven Models," *Journal of Structural Engineering* 147, no. 6 (2021): 04021058.

22. 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," *Journal of Computational Physics* 378 (2019): 686–707.

23. S. Cuomo, V. S. D. Cola, F. Giampaolo, G. Rozza, M. Raissi, and F. Piccialli, "Scientific Machine Learning Through Physics–Informed Neural Networks: Where We are and What's Next," *Journal of Scientific Computing* 92 (2022).

24. R. Zhang, Y. Liu, and H. Sun, "Physics-Informed Multi-LSTM Networks for Metamodeling of Nonlinear Structures," *Computer Methods in Applied Mechanics and Engineering* 369 (2020): 113226.

25. R. Zhang, Y. Liu, and H. Sun, "Physics-Guided Convolutional Neural Network (PHYCNN) for Data-Driven Seismic Response Modeling," *Engineering Structures* 215 (2020): 110704.

26. Y. Wu, Z. Yin, Y. Gao, S. Yang, and Y. Hou, "Constitutive Model-Constrained Physics-Informed Neural Networks Framework for Nonlinear Structural Seismic Response Prediction," *Computer Methods in Applied Mechanics and Engineering* 443 (2025): 118079.

27. S. Chen and S. Billings, "Neural Networks for Nonlinear Dynamic System Modelling and Identification," *International Journal of Control* 56 (1992): 319–346.

28. S. Lawrence, C. Giles, A. C. Tsoi, and A. Back, "Face Recognition: A Convolutional Neural-Network Approach," *IEEE Transactions on Neural Networks* 8, no. 1 (1997): 98–113.

29. Y. LeCun, Y. Bengio, and G. E. Hinton, "Deep Learning," *Nature* 521 (2015): 436–444.

30. R.-T. Wu and M. R. Jahanshahi, "Deep Convolutional Neural Network for Structural Dynamic Response Estimation and System Identification," *Journal of Engineering Mechanics* 145, no. 1 (2019): 04018125.

31. E. Junda and C. Malaga-Chuquitaype, "Seismic Acceleration Demands in Tall CLT Buildings, Predictive Models and Intensity Measures," *Engineering Structures* 298 (2024): 117024.

32. E. Junda, C. Malaga-Chuquitaype, and K. Chawgien, "Interpretable Machine Learning Models for the Estimation of Seismic Drifts in CLT Buildings," *Journal of Building Engineering* 70 (2023): 106365.

33. J. Pizarroso, J. Portela, and A. M. S. Roque, "Neuralsens: Sensitivity Analysis of Neural Networks," *Journal of Statistical Software* 102 (2022).

34. N. Makris and G. Kampas, "Size versus Slenderness: Two Competing Parameters in the Seismic Stability of Free-Standing Rocking Columns," *Bulletin of the Seismological Society of America* 106, no. 1 (2016): 104–122.

35. M. F. Vassiliou, R. Truniger, and B. Stojadinović, "An Analytical Model of a Deformable Cantilever Structure Rocking on a Rigid Surface: Development and Verification," *Earthquake Engineering & Structural Dynamics* 44, no. 15 (2015): 2775–2794.

36. C. Málaga-Chuquitaype, C. Menendez-Vicente, and R. Thiers-Moggia, "Experimental and Numerical Assessment of the Seismic Response of Steel Structures With Clutched Inerters," *Soil Dynamics and Earthquake Engineering* 121 (2019): 200–211.

37. Y. Zhang, R. Thiers-Moggia, and C. Málaga-Chuquitaype, "Impact and Clutch Nonlinearities in the Seismic Response of Inerto-Rocking Systems," *Bulletin of Earthquake Engineering* 21, no. 3 (2023): 1713–1745.

38. B. Chiou, R. Darragh, N. Gregor, and W. Silva, "NGA Project Strong-Motion Database," *Earthquake Spectra* 24, no. 1 (2008): 23–44.

39. L. WitMotion Shenzhen Co., "WT9011DCL-BT50 Bluetooth 5.0 9-Axis imu Sensor – Accelerometer, Gyroscope, Magnetometer," https://www.wit-motion.com/BLE/17.html, 2024, accessed: 2025-11-05.

40. S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation* 9, no. 8 (1997): 1735–1780.

41. Y. Bengio, P. Simard, and P. Frasconi, "Learning Long-Term Dependencies With Gradient Descent is Difficult," *IEEE Transactions on Neural Networks* 5 (1994): 157–166.

42. F. Gers, J. Schmidhuber, and F. Cummins, "Learning to Forget: Continual Prediction With LSTM," *Neural Computation* 12 (2000): 2451–2471.

43. F. Gers, N. Schraudolph, and J. Schmidhuber, "Learning Precise Timing With LSTM Recurrent Networks," *Journal of Machine Learning Research* 3 (2002): 115–143.

44. A. Graves, A. rahman Mohamed, and G. Hinton, "Speech Recognition With Deep Recurrent Neural Networks," 2013.

45. M. Sundermeyer, R. Schlüter, and H. Ney, "LSTM Neural Networks for Language Modeling," *Interspeech 2012*, 2012.

46. J. P. Ribeiro, S. M. Tavares, and M. Parente, "Stress-Strain Evaluation of Structural Parts Using Artificial Neural Networks," *Proceedings of the Institution of Mechanical Engineers, Part L: Journal of Materials: Design and Applications* 235, no. 6 (2021): 1271–1286.

47. E. G. Dimitrakopoulos and T. S. Paraskeva, "Dimensionless Fragility Curves for Rocking Response to Near-Fault Excitations," *Earthquake Engineering & Structural Dynamics* 44, no. 12 (2015): 2015–2033.

48. R. Thiers-Moggia and C. Málaga-Chuquitaype, "Performance-Based Seismic Design and Assessment of Rocking Timber Buildings Equipped With Inerters," *Engineering Structures* 248 (2021): 113164.

10969845, 2026, 12, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/eqe.70239 by Imperial College London, Wiley Online Library on [12/10/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

49. M. F. Vassiliou, K. R. Mackie, and B. Stojadinović, "A Finite Element Model for Seismic Response Analysis of Deformable Rocking Frames," *Earthquake Engineering & Structural Dynamics* 46 (2016): 447–466.

50. S. K. W. Chu, A. I. Giouvanidis, C. N. Loong, and E. G. Dimitrakopoulos, "New Perspectives in Causal Relationships Between the Response of a Rocking Block and Intensity Measures via Ensemble Machine Learning Methodologies," *Earthquake Engineering & Structural Dynamics* 54, no. 14 (2025): 3576–3593.

Appendix A: Implementation Details

This appendix complements Section 4.2 and Figure 5 of the manuscript. It is included to provide the implementation-level information required for a technically competent researcher to reproduce the reported model, retrain it for a new rocking specimen, or assess the assumptions required for such retraining. The appendix therefore focuses on the network architecture, physics-informed loss terms, loss weighting and curriculum, smoothing parameters, training protocol, data preprocessing, normalization, and dataset splitting (Tables A1–A6).

Scope of the Reported Implementation

The framework predicts rocking rotation histories for measured ground accelerations, without requiring a prescribed coefficient of restitution — impact and ene…. The framework predicts rocking rotation histories for measured ground accelerations, without requiring a prescribed coefficient of restitution – impact and energy-loss effects are instead learned from the response data, with rocking mechanics enforced via a soft physics-informed residual. Validation is limited to the experimental dataset's range; the reported regression results target non-overturning response, and use as an overturning-threshold or failure-regime predictor would require further validation with overturning cases and, where appropriate, a dedicated instability-classification branch.

A.1 | Input, Output, and Required Physical Parameters

Each training sample is one complete experimental record; training occurs at the record level, not by mixing individual time steps across records.

A.2 | Network Architecture

The model uses a fully connected pre-projection, stacked LSTM layers, and a multi-head prediction layer. This architecture applies to all reported results unless stated otherwise.

The LSTM provides temporal memory, so the predicted response at each time step depends on prior excitation and the learned hidden state – essential for rocking response, where the current state depends on the previous rocking phase, contact condition, and impact history.

A.3 | Regularized Rocking Equation and Smoothing Parameter

The physics-informed residual uses the regularized rocking equation of motion, replacing the discontinuous pivot-switching sign function with a smooth approximation:

$$\mathrm{sgn}(\theta) \approx \tanh(k\theta). \tag{A1}$$

The smoothing parameter $k = 30$ follows the reference formulation, balancing a sufficiently sharp transition near zero rotation against smooth gradients for stable optimization; smaller k over-smooths the switching behavior, while much larger k sharpens the residual but risks numerical stiffness during training.

Finite-difference derivatives in the kinematic losses and physics residual use central-difference operators, with light moving-average smoothing (kernel size 3) applied where needed to suppress spike artefacts without removing post-impact response characteristics.

A.4 | Loss Function and Weighting Strategy

The total objective combines supervised response matching, kinematic consistency, physics enforcement, integration consistency, and regularization:

$$\mathcal{L}_{\text{total}} = \alpha\mathcal{L}_{\text{data}} + \beta_1\mathcal{L}_{pv} + \beta_2\mathcal{L}_{dd} + \gamma\mathcal{L}_{PE} + \delta\mathcal{L}_{\text{int}} + \epsilon\mathcal{L}_{\text{reg}}. \tag{A2}$$

All losses use sequence masks so zero-padded time steps do not contribute to gradients; boundary masks are also applied to finite-difference losses to avoid derivative artefacts at the start and end of each sequence.

A.5 | Curriculum Schedule

A staged curriculum is used because imposing the equation-of-motion residual too strongly early in training, while predictions remain inaccurate, can destabilise optimization. The supervised data term therefore dominates initially, with physics-informed and kinematic terms increased gradually – following the general principle of prior physics-informed LSTM studies, adapted here to the non-smooth, impact-driven nature of rocking response.

Where acceleration supervision is included, the acceleration used in the physics residual is gradually blended from finite-difference to

TABLE A1 | Network input/output definitions and parameters.

Item	Description used in this study
Network input	Horizontal ground acceleration history, $a_g(t)$, sampled at the same time points as the measured response.
Primary output	Rocking rotation history, $\theta(t)$, in radians.
Auxiliary outputs	Angular velocity and acceleration heads for kinematic consistency; measured or reconstructed from $\theta(t)$ via finite differences and smoothing.
Physical parameters	Block mass m, pivot-to-CoM distance R, slenderness angle α, gravitational acceleration g, sampling interval Δt, moment of inertia I, and $\Gamma = I/(mR)$.
Dataset unit convention	Ground acceleration in physical units for the EOM residual; network input normalized for training. Rotation normalized for supervised learning, recovered in physical units for physics-loss evaluation.

10969845, 2026, 12, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/eqe.70239 by Imperial College London, Wiley Online Library on [12/10/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

TABLE A2 | Network architecture configuration.

Component	Configuration
Input projection	Fully connected layer mapping the input ground-acceleration channel to hidden dimension H.
Activation	SiLU, applied after input projection.
Sequence model	Two-layer unidirectional LSTM.
Hidden dimension	$H = 128$ hidden units.
Dropout	$p = 0.05$, applied between stacked LSTM layers.
Output layer	Fully connected multi-head output layer.
Output heads	Standardized rotation, angular velocity, and angular acceleration; rotation is the primary response output.
Residual skip connection	Small linear skip from input to output heads, preserving low-order trends and aiding optimization; small initial weight limits direct bypass of the recurrent representation.

TABLE A3 | Loss components and final-stage weights.

Term	Purpose	Final-stage weight
$\mathcal{L}_{\text{data}}$	Supervised MSE between predicted and measured standardized rotation.	$\alpha = 10.0$
$\mathcal{L}_{pv}$	Velocity kinematic consistency, direct prediction vs. finite-difference reconstruction.	$\beta_1 = 1.0$
$\mathcal{L}_{dd}$	Acceleration kinematic consistency, direct prediction vs. finite-difference reconstruction.	$\beta_2 = 1.0$ if acceleration head used; otherwise disabled.
$\mathcal{L}_{PE}$	Residual of the regularized rocking equation of motion.	$\gamma = 1.0$
$\mathcal{L}_{\text{int}}$	Integration consistency between predicted rotation and rotation reconstructed from predicted velocity.	$\delta = 0.05$
$\mathcal{L}_{\text{reg}}$	Spectral, temporal-gradient, and energy regularization.	ϵ ramped from 0 to small final values (representative: 0.05, 0.10, 0.10).

TABLE A4 | Curriculum stages and loss weighting.

Training stage	Epochs	Weighting strategy
1: data-dominated fitting	1–150	$\alpha = 20.0$, $\beta_1 = 0.05$, $\gamma = 0.0$, $\delta = 0.01$; regularization inactive.
2: transition	151–600	α: 20.0 → 10.0; β_1: 0.05 → 1.0; γ: 0.0 → 1.0; δ: 0.01 → 0.05.
3: steady training	601–end	Final weights retained: $\alpha = 10.0$, $\beta_1 = 1.0$, $\gamma = 1.0$, $\delta = 0.05$; regularization active.

TABLE A5 | Training protocol and hyperparameters.

Item	Value used in this study
Optimizer	Adam.
Initial learning rate	2×10^{-4}.
LR schedule	Cosine annealing to a minimum of 1×10^{-6}.
Batch size	16 records per mini-batch, with sequence padding and masking.
Training epochs	8000.
Gradient clipping	Max norm 2.0 initially, reduced to 1.0 after early data-fitting phase.
Random seed	Three random seeds, same record-level split; mean and range of test RMSE reported to assess initialisation sensitivity.
Precision	Mixed precision where supported; LSTM recurrence kept in full precision for numerical stability.

10969845, 2026, 12, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/eqe.70239 by Imperial College London, Wiley Online Library on [12/10/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

TABLE A6 | Data preprocessing steps.

Step	Procedure
Signal alignment & sampling	Ground acceleration and rotation placed on a common time vector and sampling interval.
Baseline correction	Pre-event baseline subtracted from rotation so initial rotation is near zero.
Record validation	Checked for required fields, consistent array lengths, finite values, positive time step, sufficient sequence length.
Ground-motion constraints	Accelerograms preprocessed to satisfy shaking-table displacement limits, affecting the excitation range represented; discussed as a manuscript limitation.
Input normalisation	Acceleration standardized using training-set statistics; normalized acceleration used as network input, physical-unit acceleration retained for the EOM residual.
Output normalisation	Rotation standardized using training-set statistics for supervised learning; de-standardized for the physics residual and reported predictions.
Derivative quantities	Velocity and acceleration measured directly or reconstructed via finite differences and smoothing; more sensitive to noise, filtering, and differentiation than rotation.

mixed finite-difference/network-predicted acceleration, preventing the acceleration head from dominating the residual before learning a stable representation.

A.6 | Training Protocol and Hyperparameters

Adam was used, consistent with prior physics-informed LSTM studies for nonlinear structural response prediction. The initial learning rate (2×10^{-4}) is more conservative than the 10^{-3} commonly used in earlier PI-LSTM work, reflecting the sensitivity introduced by abrupt impact-related transitions and derivative-based physics residuals. The longer schedule, combined with cosine annealing, allows gradual transition from data-dominated fitting to physics-informed refinement.

A.7 | Data Preprocessing and Normalisation

The same preprocessing pipeline is applied to training, validation, and test records. Normalisation parameters are fitted on the training set only and applied unchanged elsewhere. Network input and the supervised rotation loss use normalized quantities; the physics-informed residual uses physical-unit acceleration and de-normalized rotation.

A.8 | Dataset Splitting and Leakage Control

The train-validation-test split is performed at the record level, with all time steps from a given record assigned to the same subset, preventing leakage between training and testing. The split precedes normalisation and training; normalisation statistics are fitted on the training set only.

For retraining on a new specimen, the same record-level rule should be followed. Where multiple tests share a specimen, repeated run, or ground-motion family, group-based splitting is recommended when assessing generalization to unseen specimens, tests, or excitation classes.

A.9 | Information Required to Retrain the Model for a New Rocking System

To retrain the framework for a new configuration, the following should be provided:

- Measured ground acceleration and corresponding rocking rotation histories.
- Sampling interval and record duration per test.
- Specimen mass, geometry, slenderness angle, pivot-to-center-of-mass distance, and moment of inertia.
- Preprocessing choices, including filtering, baseline correction, and displacement/PGD restrictions.
- Training-set normalisation statistics for acceleration and rotation.
- Record-level train-validation-test split.
- Network hyperparameters, smoothing parameter k, loss weights, and curriculum schedule.
- Definition of the target application: stable response prediction, overturning classification, or both.

If the intended application includes overturning or instability-threshold prediction, the training dataset should include sufficient overturning and near-overturning cases, and the response-regression model should be supplemented by an explicitly validated instability-classification or threshold-prediction component. Without such validation, the framework should be interpreted as a stable-response prediction model within the range of the available experimental dataset.

10969845, 2026, 12, Downloaded from https://onlinelibrary.wiley.com/doi/10.1002/eqe.70239 by Imperial College London, Wiley Online Library on [12/10/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License